FORTINET

FortiOS Handbook
Security Profiles for FortiOS 5.0

FortiOS Handbook - Security Profiles for FortiOS 5.0

May 9, 2014

01-500-108920-20140403

Technical Documentation	docs.fortinet.com
Knowledge Base	kb.fortinet.com
Customer Service & Support	support.fortinet.com
Training Services	training.fortinet.com
FortiGuard	fortiguard.com
Document Feedback	techdocs@fortinet.com

Change Log

Date	Change Description
2014-05-09	Updated Antivirus flowcharts and minor fixes
2014-04-03	Updated description of FortiCloud Sandbox
2014-03-18	Multiple minor fixes and updates.
2013-12-05	Updated IPS signature rate count threshold
2013-07-25	Added topic: Profile Groups, Added Client Reputation Added Web Filter Overrides Replaced URL Filter with Web Site Filter Renamed Title of Book to match new Web-Based Manager section heading.
2013-04-24	Changes made according to fixes 155712, 121776 and expanded the Watermarking section to include new version of software.
2013-02-04	Initial Release.

Table of Contents

This FortiOS Handbook chapter contains the following sections:

Security Profiles overview describes Security Profiles components and their relation to firewall policies, as well as SSL content scanning and inspection. We recommend starting with this section to become familiar with the different features in your FortiGate unit.

Client Reputation explains how to track client behavior and report on activities that you determine are risky or otherwise noteworthy.

AntiVirus explains how the FortiGate unit scans files for viruses and describes how to configure the antivirus options.

Email filter explains how the FortiGate unit filters email, describes how to configure the filtering options and the action to take with email detected as spam.

Intrusion protection explains basic Intrusion Protection System (IPS) concepts and how to configure IPS options; includes guidance and a detailed table for creating custom signatures as well as several examples.

Web filter and FortiGuard Web Filter The first of these sections describes basic web filtering concepts, the order in which the FortiGate unit performs web filtering, and configuration. The second section describes enhanced features of the subscription-based FortiGuard Web Filtering service and explains how to configure them. We recommend reading both sections if you are using FortiGuard Web Filtering because settings you configure in one feature may affect the other.

Data leak prevention describes the DLP features that allow you to prevent sensitive data from leaving your network and explains how to configure the DLP rules, compound rules, and sensors.

Application control describes how your FortiGate unit can detect and take action against network traffic based on the application generating the traffic.

Security Profiles overview

Ranging from the FortiGate®-30 series for small businesses to the FortiGate-5000 series for large enterprises, service providers and carriers, the FortiGate line combines a number of security features to protect your network from threats. As a whole, these features, when included in a single Fortinet security appliance, are referred to as Security Profiles. The Security Profiles features your FortiGate model includes are:

- AntiVirus
- Intrusion Prevention System (IPS)
- Web filtering
- E-mail filtering, including protection against spam and grayware
- Data Leak Prevention (DLP)
- Application Control
- ICAP

Firewall policies limit access, and while this and similar features are a vital part of securing your network, they are not covered in this document.

The following topics are included in this section:

- Traffic inspection
- Content inspection and filtering
- Security Profiles components
- Security Profiles/lists/sensors

Traffic inspection

When the FortiGate unit examines network traffic one packet at a time for IPS signatures, it is performing traffic analysis. This is unlike content analysis where the traffic is buffered until files, email messages, web pages, and other files are assembled and examined as a whole.

DoS policies use traffic analysis by keeping track of the type and quantity of packets, as well as their source and destination addresses.

Application control uses traffic analysis to determine which application generated the packet.

Although traffic inspection doesn't involve taking packets and assembling files they are carrying, the packets themselves can be split into fragments as they pass from network to network. These fragments are reassembled by the FortiGate unit before examination.

No two networks are the same and few recommendations apply to all networks. This topic offers suggestions on how you can use the FortiGate unit to help secure your network against content threats.

IPS signatures

IPS signatures can detect malicious network traffic. For example, the Code Red worm attacked a vulnerability in the Microsoft IIS web server. Your FortiGate's IPS system can detect traffic attempting to exploit this vulnerability. IPS may also detect when infected systems communicate with servers to receive instructions.

IPS recommendations

- Enable IPS scanning at the network edge for all services.
- Use FortiClient endpoint IPS scanning for protection against threats that get into your network.
- Subscribe to FortiGuard IPS Updates and configure your FortiGate unit to receive push updates. This will ensure you receive new IPS signatures as soon as they are available.
- Your FortiGate unit includes IPS signatures written to protect specific software titles from DoS attacks. Enable the signatures for the software you have installed and set the signature action to *Block*.
- You can view these signatures by going to *Security Profiles > Intrusion Protection > Predefined* and sorting by, or applying a filter to, the *Group* column.
- Because it is critical to guard against attacks on services that you make available to the public, configure IPS signatures to block matching signatures. For example, if you have a web server, configure the action of web server signatures to *Block*.

Suspicious traffic attributes

Network traffic itself can be used as an attack vector or a means to probe a network before an attack. For example, SYN and FIN flags should never appear together in the same TCP packet. The SYN flag is used to initiate a TCP session while the FIN flag indicates the end of data transmission at the end of a TCP session.

The FortiGate unit has IPS signatures that recognize abnormal and suspicious traffic attributes. The SYN/FIN combination is one of the suspicious flag combinations detected in TCP traffic by the `TCP.BAD.FLAGS` signature.

The signatures that are created specifically to examine traffic options and settings, begin with the name of the traffic type they are associated with. For example, signatures created to examine TCP traffic have signature names starting with TCP.

Application control

While applications can often be blocked by the ports they use, application control allows convenient management of all supported applications, including those that do not use set ports.

Application control recommendations

- Some applications behave in an unusual manner in regards to application control. For more information, see "Application considerations" on page 144.
- By default, application control allows the applications not specified in the application control list. For high security networks, you may want to change this behavior so that only the explicitly allowed applications are permitted.

Content inspection and filtering

When the FortiGate unit buffers the packets containing files, email messages, web pages, and other similar files for reassembly before examining them, it is performing content inspection. Traffic inspection, on the other hand, is accomplished by the FortiGate unit examining individual packets of network traffic as they are received.

No two networks are the same and few recommendations apply to all networks. This topic offers suggestions on how you can use the FortiGate unit to help secure your network against content threats. Be sure to understand the effects of the changes before using the suggestions.

AntiVirus

The FortiGate antivirus scanner can detect viruses and other malicious payloads used to infect machines. The FortiGate unit performs deep content inspection. To prevent attempts to disguise viruses, the antivirus scanner will reassemble fragmented files and uncompress content that has been compressed. Patented Compact Pattern Recognition Language (CPRL) allows further inspection for common patterns, increasing detection rates of virus variations in the future.

AntiVirus recommendations

- Enable antivirus scanning at the network edge for all services.
- Use FortiClient endpoint antivirus scanning for protection against threats that get into your network.
- Subscribe to FortiGuard AntiVirus Updates and configure your FortiGate unit to receive push updates. This will ensure you receive new antivirus signatures as soon as they are available.
- Enable the Extended Virus Database if your FortiGate unit supports it.
- Examine antivirus logs periodically. Take particular notice of repeated detections. For example, repeated virus detection in SMTP traffic could indicate a system on your network is infected and is attempting to contact other systems to spread the infection using a mass mailer.
- The *builtin-patterns* file filter list contains nearly 20 file patterns. Many of the represented files can be executed or opened with a double-click. If any of these file patterns are not received as a part of your normal traffic, blocking them may help protect your network. This also saves resources since files blocked in this way do not need to be scanned for viruses.
- To conserve system resources, avoid scanning email messages twice. Scan messages as they enter and leave your network or when clients send and retrieve them, rather than both.

FortiGuard Web Filtering

The web is the most popular part of the Internet and, as a consequence, virtually every computer connected to the Internet is able to communicate using port 80, HTTP. Botnet communications take advantage of this open port and use it to communicate with infected computers. FortiGuard Web Filtering can help stop infections from malware sites and help prevent communication if an infection occurs.

FortiGuard Web Filtering recommendations

- Enable FortiGuard Web Filtering at the network edge.
- Install the FortiClient application and use FortiGuard Web Filtering on any systems that bypass your FortiGate unit.
- Block categories such as Pornography, Malware, Spyware, and Phishing. These categories are more likely to be dangerous.
- In the email filter profile, enable *IP Address Check* in *FortiGuard Email Filtering*. Many IP addresses used in spam messages lead to malicious sites; checking them will protect your users and your network.

Email filter

Spam is a common means by which attacks are delivered. Users often open email attachments they should not, and infect their own machine. The FortiGate email filter can detect harmful spam and mark it, alerting the user to the potential danger.

Email filter recommendations

- Enable email filtering at the network edge for all types of email traffic.
- Use FortiClient endpoint scanning for protection against threats that get into your network.
- Subscribe to the FortiGuard AntiSpam Service.

DLP

Most security features on the FortiGate unit are designed to keep unwanted traffic out of your network while DLP can help you keep sensitive information from leaving your network. For example, credit card numbers and social security numbers can be detected by DLP sensors.

DLP recommendations

- Rules related to HTTP posts can be created, but if the requirement is to block all HTTP posts, a better solution is to use application control or the *HTTP POST Action* option in the web filter profile.
- While DLP can detect sensitive data, it is more efficient to block unnecessary communication channels than to use DLP to examine it. If you don't use instant messaging or peer-to-peer communication in your organization, for example, use application control to block them entirely.

Security Profiles components

AntiVirus

Your FortiGate unit stores a virus signature database that can identify more than 15,000 individual viruses. FortiGate models that support additional virus databases are able to identify hundreds of thousands of viruses. With a FortiGuard AntiVirus subscription, the signature databases are updated whenever a new threat is discovered.

AntiVirus also includes file filtering. When you specify files by type or by file name, the FortiGate unit will stop the matching files from reaching your users.

FortiGate units with a hard drive or configured to use a FortiAnalyzer unit can store infected and blocked files for that you can examine later.

Intrusion Protection System (IPS)

The FortiGate Intrusion Protection System (IPS) protects your network against hacking and other attempts to exploit vulnerabilities of your systems. More than 3,000 signatures are able to detect exploits against various operating systems, host types, protocols, and applications. These exploits can be stopped before they reach your internal network.

You can also write custom signatures, tailored to your network.

Web filtering

Web filtering includes a number of features you can use to protect or limit your users' activity on the web.

FortiGuard Web Filtering is a subscription service that allows you to limit access to web sites. More than 60 million web sites and two billion web pages are rated by category. You can choose to allow or block each of the 77 categories.

URL filtering can block your network users from access to URLs that you specify.

Web content filtering can restrict access to web pages based on words and phrases appearing on the web page itself. You can build lists of words and phrases, each with a score. When a web content list is selected in a web filter profile, you can specify a threshold. If a user attempts to load a web page and the score of the words on the page exceeds the threshold, the web page is blocked.

Email filtering

FortiGuard AntiSpam is a subscription service that includes an IP address black list, a URL black list, and an email checksum database. These resources are updated whenever new spam messages are received, so you do not need to maintain any lists or databases to ensure accurate spam detection.

You can use your own IP address lists and email address lists to allow or deny addresses, based on your own needs and circumstances.

Data Leak Prevention (DLP)

Data leak prevention allows you to define the format of sensitive data. The FortiGate unit can then monitor network traffic and stop sensitive information from leaving your network. Rules for U.S. social security numbers, Canadian social insurance numbers, as well as Visa, Mastercard, and American Express card numbers are included.

Application Control

Although you can block the use of some applications by blocking the ports they use for communications, many applications do not use standard ports to communicate. Application control can detect the network traffic of more than 1000 applications, improving your control over application communication.

ICAP

This module allows for the offloading of certain processes to a separate server so that your FortiGate firewall can optimize its resources and maintain the best level of performance possible.

Security Profiles/lists/sensors

A profile is a group of settings that you can apply to one or more firewall policies. Each Security Profile feature is enabled and configured in a profile, list, or sensor. These are then selected in a security policy and the settings apply to all traffic matching the policy. For example, if you create an antivirus profile that enables antivirus scanning of HTTP traffic, and select the antivirus profile in the security policy that allows your users to access the World Wide Web, all of their web browsing traffic will be scanned for viruses.

Because you can use profiles in more than one security policy, you can configure one profile for the traffic types handled by a set of firewall policies requiring identical protection levels and types, rather than repeatedly configuring those same profile settings for each individual security policy.

For example, while traffic between trusted and untrusted networks might need strict protection, traffic between trusted internal addresses might need moderate protection. To provide the different levels of protection, you might configure two separate sets of profiles: one for traffic between trusted networks, and one for traffic between trusted and untrusted networks.

The Security Profiles include:

- antivirus profile
- IPS sensor
- Web filter profile
- Email filter profile
- Data Leak Prevention profile
- Application Control list
- VoIP profile

Although they're called profiles, sensors, and lists, they're functionally equivalent. Each is used to configure how the feature works.

Client Reputation

The Security scan types available on FortiGate units are varied and tailored to detect specific attacks. However, sometimes user/client behavior can increase the risk of attack or infection. For example, if one of your network clients receives email viruses on a daily basis while no other clients receive these attachments, extra measures may be required to protect the client, or a discussion with the user about this issue may be worthwhile.

Before you can decide on a course of action, you need to know the problem is occurring. Client reputation can provide this information by tracking client behavior and reporting on activities that you determine are risky or otherwise noteworthy.

To learn how to enable and customize Client Reputation on your FortiGate go to the following FortiGate Cookbook video:

Client Reputation Video

Summary of the Client Reputation features

Activities you can track include:

- Bad Connection Attempts: A typical BOT behavior is to connect to some hosts that do not exist on the Internet. This is because the BOT home needs to constantly change itself to dodge legislative enforcement or to hide from AV vendors. Bad connection attempts are tracked by:
  - Look ups for a DNS name that does not exist.
  - Connection attempts to an IP address that has no route.
  - HTTP 404 errors
- Packets that are blocked by security policies.
- Intrusion protection: Attack detected. The effect on reputation increases with severity of attack. A subscription to FortiGuard IPS updates is required.
- Malware protection: Malware detected. This requires a subscription to FortiGuard Antivirus updates.
- Web activity: Visit to web site in risky categories, including Potentially Liable, Adult/Mature Content, Bandwidth Consuming and Security Risk. A subscription to FortiGuard Web Filtering is required.
- Application protection: Client uses software in risky categories, including Botnet, P2P, Proxy, and Games applications. A subscription to FortiGuard IPS updates is required.
- Geographical locations that clients are communicating with. Access to the FortiGuard geographic database and a valid Fortinet support contract is required.

You can configure how severely each type of tracked activity will impact the reputation of the client in a sliding scale of Low, Medium, High or Critical. You can also choose to ignore an activity by setting it to Off. When an activity is turned off, it will have no effect on reputation.

You can enable client reputation tracking for your FortiGate unit by going to *Security Profiles > Client Reputation > Threat Level Definition*. Turning on client reputation tracking turns on traffic logging for all security policies, for all DoS policies and for all sniffer policies. While client

reputation is enabled, logging cannot be turned off for these policies. Traffic logging must be enabled for data to be added to the client reputation database.

Client reputation only highlights risky activity and does not include tools to stop it. Instead, client reputation is a tool that exposes risky behavior. When you uncover risky behavior that you are concerned about, you can take additional action to stop it. That action could include adding more restrictive security policies to block the activity or increase Security Profiles protection. You can also taking other measures outside your FortiGate unit to stop the activity.

To support client reputation your FortiGate unit must be registered, have a valid support contract and be licensed for FortiGuard antivirus, IPS and Web Filtering.

After client reputation is turned on, the FortiGate unit tracks recent behavior using a sliding window and displays current data for this window. The client reputation monitor displays clients and their activities in charts ordered according to how risky the behavior exhibited by the client is.

Client Reputation data is stored in traffic log messages in the newly added client reputation fields (crscore and craction). When you enable client reputation *Log Security Events* or *Log all Sessions* is enabled in all security policies. *Log Security Events* records traffic log messages for Security Profile sessions and *Log all Sessions* records traffic logs for all sessions. When Client Reputation is enabled you cannot select *No Log* in a security policy. Using client reputation data in log messages, you can configure FortiAnalyzer to produce a client reputation report.

Enabling client reputation can affect system performance if you had not been using traffic logging.

This chapter describes:

- Applying client reputation monitoring to your network
- Viewing client reputation results
- Setting the client reputation profile/definition
- Expanding client reputation to include more types of behavior
- Client reputation execute commands
- Client reputation diagnose commands

Applying client reputation monitoring to your network

Client reputation monitoring is applied to network traffic by going to *Security Profiles > Client Reputation > Threat Level Definition* turning on *Client Reputation Tracking* and selecting *Apply*.

You can then either change the client reputation profile used by your FortiGate unit or you can accept the default profile. The client reputation profile indicates how risky you consider different types of client behavior to be. See for details.

Viewing client reputation results

To view Client Reputation results go to *Security Profiles > Client Reputation > Reputation Score* to view the client reputation monitor. The monitor displays information about risky behavior as it was found. You can drill down into individual items to get more information about the behavior found and the client that caused it.

The client reputation monitor updates every 2 minutes. You can also select *Refresh* to manually update the display.

Select *Reset* to clear all client reputation data and restart the reporting window.

Figure 1 shows example client reputation results that shows activity from for different IP addresses that matched the kinds of traffic to be monitored according to the client reputation profile. You can see the IP address or name of each client and the amount of risky activity detected. The list at the bottom of the display shows more information about each device. The device information is gathered from enabling device monitoring by going to *User & Device > Device > Device Definition*.

Figure 1: Example client reputation results

You can select any of the bars in the graph to view information for each time the risky behavior was detected during the past 7 days (or whatever the Client Reputation window is). Information for each event detected includes the date and time the event was detected, the destination address, the application, and the client reputation score.

Changing the client reputation reporting window and database size

By default, client reputation reports on activity for the last seven days. You can change this reporting window using the following command:

```
config client-reputation profile
   set window-size <interval_int>
end
```

Where <interval-int> is the reporting window in days. Range 1 to 30 days, default 7 days.

Enter the following command to set the client reputation report size:

```
config client-reputation profile
   set max-rep-db-size <size>
end
```

Where <size> can be from 10 to 2000 MBytes (2 GBytes). The default size is 100 MBytes.

Client reputation data update and maintenance intervals

Client reputation updates its database every 2 minutes by querying the log database for client reputation information. This means that data displayed in the client reputation monitor is very current, at the most 2 minutes old.

Client reputation includes a data maintenance routine that runs every 12 hours to perform maintenance functions on the client reputation database. This routine:

- Checks the number of tracked hosts. If the number is at the maximum of 5000, the maintenance routing removes the oldest ten percent (500) of hosts from the list. If the number is less than the maximum, nothing changes.

- Deletes any reputation data associated with a host that is not in the tracking list (usually this only occurs if hosts are removed).

- Deletes any reputation data that is older than the current time minus the window-size in days.

Setting the client reputation profile/definition

Configure the client reputation profile by going to *Security Profiles > Client Reputation > Threat Level Definition*. You configure one client reputation profile for all of the activity monitored by the FortiGate unit. The profile sets the risk levels for the types of behavior that client reputation monitors. You can set the risk to off, low, medium, high and critical for the following types of behavior:

- Application Protection
  - Botnet applications
  - P2P applications
  - Proxy applications
  - Games applications
- Intrusion protection (IPS)
  - Critical severity attack detected
  - High severity attack detected
  - Medium severity attack detected
  - Low severity attack detected
  - Informational severity attack detected
- Malware Protection
  - Malware detected
  - Botnet connection detected
- Packet based inspection
  - Blocked by firewall policy
  - Failed connection attempts
- Web Activity
  - All blocked URLs
  - Visit to security risk sites
  - Visit to potentially liable sites
  - Visit to adult/mature content sites
  - Visit to bandwidth consuming sites

Figure 2: Default client reputation profile

To configure the profile, decide how risky or dangerous each of the types of behavior are to your network and rate them accordingly. The higher you rate a type of behavior, the more visible clients engaging in this behavior will become in the client reputation monitor and the more easily you can detect this behavior.

For example, if you consider malware a high risk for your network, you can set the client reputation profile for malware to high or critical (as it is in the default client reputation profile). Then, whenever any amount of malware is detected, clients that originated the malware will be very visible in the client reputation monitor.

Set the risk to off for types of activity that you do not want client reputation to report on. This does not reduce the performance requirements or the amount of data gathered by client reputation, just the report output.

You can change a profile setting at any time and data that has already been collected will be used.

It is normally not necessary to change the *Risk Level Values* but it can be done if you need to alter the relative importance of the risk settings.

Expanding client reputation to include more types of behavior

You can use the following command to change the client reputation profile from the CLI to include client reputation reporting about more settings:

```
config client-reputation profile
```

In addition to the settings configurable from the web-based manager, you can also set the following options:

- `geolocation` to enable reporting on connections to and from different countries (geographical locations). For example, use the following command to indicate that you consider communication with Aruba to be medium risk:

```
config client-reputation profile
   config geolocation
      edit 0
         set country AW
         set level medium
      end
   end
```

- `url-block-detected` to report on connections blocked by web filtering. Use the following command to enable reporting about blocked URLs and set the risk level to medium:

```
config client-reputation profile
   set url-block-detected medium
end
```

From the CLI you can configure client reputation to report more FortiGuard web filtering categories and more types of applications. For example, to report on social network activity (application control category 23):

```
config client-reputation-profile
   config application
      edit 0
         set category 23
         set level medium
      end
   end
```

To report on the local web filtering category (category 22):

```
config client-reputation-profile
   config web
      edit 0
         set group 22
         set level medium
      end
   end
```

Client reputation execute commands

The `execute client-reputation` command includes the following options:

- `erase`, deletes all client reputation data.
- `host-count`, lists the clients that started sessions recorded by client reputation
- `host-detail`, for a specified client's IP address, displays the client reputation traffic log messages saved for that client.
- `host-summary`, for a specified client's IP address, displays the client's IP address, total entries, and total score.
- `purge`, deletes all data from the client reputation database.
- `topN`, display the top N clients identified by client reputation.

Client reputation diagnose commands

The `diagnose client-reputation` command includes the following options

- `convert-timestamp` convert a client reputation database timestamp to date and time
- `test-all` adds log messages from multiple sources to the client reputation database for testing
- `test-app` adds application control log messages to the client reputation database for testing
- `test-ips` adds Intrusion Protection log messages to the client reputation database for testing
- `test-webfilter` adds webfilter log messages to the client reputation database for testing

AntiVirus

This section describes how to configure the antivirus options. From an antivirus profile you can configure the FortiGate unit to apply antivirus protection to HTTP, FTP, IMAP, POP3, SMTP, IM, and NNTP sessions. If your FortiGate unit supports SSL content scanning and inspection, you can also configure antivirus protection for HTTPS, IMAPS, POP3S, SMTPS, and FTPS sessions. You can also apply flow-based antivirus projection to SMB or CIFS (Microsoft file sharing) traffic.

In many cases you can just customize the default antivirus profile and apply it to the security policy that accepts the traffic to be virus scanned. You can also create custom antivirus profiles if want to apply different types of virus protection to different traffic.

The following topics are included in this section:

- Antivirus concepts
- Enable antivirus scanning
- Grayware scanning
- Windows file sharing (CIFS) flow-based antivirus scanning
- Advanced Persistent Threat (APT) protection
- Testing your antivirus configuration
- Antivirus examples

Antivirus concepts

The word "antivirus" refers to a group of features that are designed to prevent unwanted and potentially malicious files from entering your network. These features all work in different ways, which include checking for a file size, name, or type, or for the presence of a virus or grayware signature.

The antivirus scanning routines your FortiGate unit uses are designed to share access to the network traffic. This way, each individual feature does not have to examine the network traffic as a separate operation, and the overhead is reduced significantly. For example, if you enable file filtering and virus scanning, the resources used to complete these tasks are only slightly greater than enabling virus scanning alone. Two features do not require twice the resources.

How antivirus scanning works

Antivirus scanning examines files for viruses, worms, trojans, and malware. The antivirus scan engine has a database of virus signatures it uses to identify infections. If the scanner finds a signature in a file, it determines that the file is infected and takes the appropriate action.

The most thorough scan requires that the FortiGate unit have the whole file for the scanning procedure. To achieve this, the antivirus proxy buffers the file as it arrives. Once the transmission is complete, the virus scanner examines the file. If no infection is present, it is sent to the destination. If an infection is present, a replacement message is set to the destination.

During the buffering and scanning procedure, the client must wait. With a default configuration, the file is released to the client only after it is scanned. You can enable client comforting in the Proxy Options profile to feed the client a trickle of data to prevent them from thinking the transfer is stalled, and possibly cancelling the download.

Buffering the entire file allows the FortiGate unit to eliminate the danger of missing an infection due to fragmentation because the file is reassembled before examination. Archives can also be expanded and the contents scanned, even if archives are nested.

Since the FortiGate unit has a limited amount of memory, files larger than a certain size do not fit within the memory buffer. The default buffer size is 10 MB. You can use the `uncompsizelimit` CLI command to adjust the size of this memory buffer.

Files larger than the buffer are passed to the destination without scanning. You can use the *Oversize File/Email* setting to block files larger than the antivirus buffer if allowing files that are too large to be scanned is an unacceptable security risk.

Flow-based antivirus scanning

If your FortiGate unit supports flow-based antivirus scanning, you can choose to select it instead of proxy-based antivirus scanning. Flow-based antivirus scanning uses the FortiGate IPS engine to examine network traffic for viruses, worms, trojans, and malware, without the need to buffer the file being checked.

The advantages of flow-based scanning include faster scanning and no maximum file size. Flow-based scanning doesn't require the file be buffered so it is scanned as it passes through the FortiGate unit, packet-by-packet. This eliminates the maximum file size limit and the client begins receiving the file data immediately. Also, flow-based scanning does not change packets as they pass through the FortiGate unit, while proxy-based scanning can change packet details such as sequence numbers. The changes made by proxy-based scanning do not affect most networks.

The trade-off for these advantages is that flow-based scans detect a smaller number of infections. Viruses in documents, packed files, and some archives are less likely to be detected because the scanner can only examine a small portion of the file at any moment. Also, the file archive formats flow-based scanning will examine are limited to ZIP and GZIP.

Antivirus scanning order

The antivirus scanning function includes various modules and engines that perform separate tasks.

Proxy-based antivirus scanning order

Figure 3 on page 28 illustrates the antivirus scanning order when using proxy-based scanning. The first check for oversized files/email is to determine whether the file exceeds the configured size threshold. The `uncompsizelimit` check is to determine if the file can be buffered for file type and antivirus scanning. If the file is too large for the buffer, it is allowed to pass without being scanned. For more information, see the `config antivirus service` command. The antivirus scan includes scanning for viruses, as well as for grayware and heuristics if they are enabled.

File filtering includes file pattern and file type scans which are applied at different stages in the antivirus process.

Figure 3: Antivirus scanning order when using the normal, extended, or extreme database

Start

FTP, NNTP, SMTP, POP3, IMAP, HTTP traffic

Buffering Stage
File message is buffered

Oversized threshold checking stage
Does file/ email exceed threshold? — Yes → Oversized file/ email action — Block
No

Pass

Pattern matching stage
Does file pattern match? — Yes → Matching file pattern action — Block
No → Allow

Does file pattern match? — Yes → Matching file pattern action — Block
No → Allow

Uncompsizelimit threshold checking stage
Does file/ email exceed threshold?
Yes | No

Antivirus profile filtering stage
Does scan detect infection? — Yes
No

Type matching stage
Does file type match? — Yes → Matching file type action — Block
No → Allow

Block

Pass file/ email

Passes to next process

If a file fails any of the tasks of the antivirus scan, no further scans are performed. For example, if the file `fakefile.EXE` is recognized as a blocked file pattern, the FortiGate unit will send the end user a replacement message, and delete or quarantine the file. The unit will not perform virus scan, grayware, heuristics, and file type scans because the previous checks have already determined that the file is a threat and have dealt with it.

Flow-based antivirus scanning order

illustrates the antivirus scanning order when using flow-based scanning (i.e. the flow-based database). The antivirus scan takes place before any other antivirus-related

scan. If file filter is not enabled, the file is not buffered. The antivirus scan includes scanning for viruses, as well as for grayware and heuristics if they are enabled.

File filtering includes file pattern and file type scans which are applied at different stages in the antivirus process.

Figure 4: Antivirus scanning order when using flow-based scanning
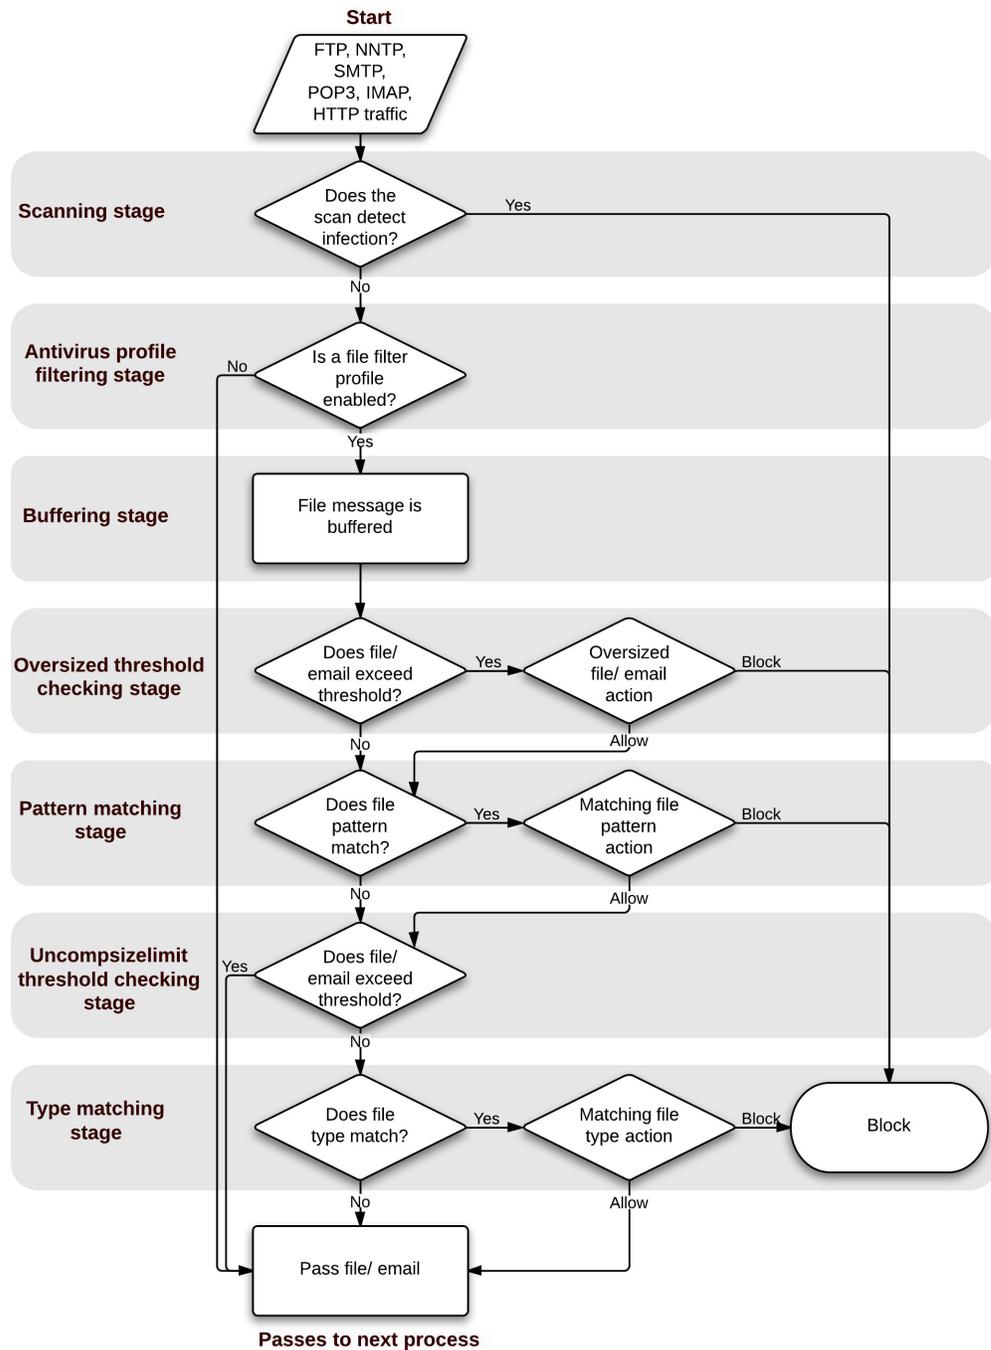
Antivirus databases

The antivirus scanning engine relies on a database of virus signatures to detail the unique attributes of each infection. The antivirus scan searches for these signatures, and when one is discovered, the FortiGate unit determines the file is infected and takes action.

All FortiGate units have the normal antivirus signature database but some models have additional databases you can select for use. Which you choose depends on your network and security needs.

Normal	Includes viruses currently spreading as determined by the FortiGuard Global Security Research Team. These viruses are the greatest threat. The Normal database is the default selection and it is available on every FortiGate unit.
Extended	Includes the normal database in addition to recent viruses that are no-longer active. These viruses may have been spreading within the last year but have since nearly or completely disappeared.
Extreme	Includes the extended database in addition to a large collection of 'zoo' viruses. These are viruses that have not spread in a long time and are largely dormant today. Some zoo viruses may rely on operating systems and hardware that are no longer widely used.

Antivirus techniques

The antivirus features work in sequence to efficiently scan incoming files and offer your network optimum antivirus protection. The first four features have specific functions, the fifth, heuristics, protects against new, or previously unknown virus threats. To ensure that your system is providing the most protection available, all virus definitions and signatures are updated regularly through the FortiGuard antivirus services. The features are discussed in the order that they are applied, followed by FortiGuard antivirus.

Virus scan

If the file passes the file pattern scan, the FortiGate unit applies a virus scan to it. The virus definitions are kept up-to-date through the FortiGuard Distribution Network (FDN). For more information, see "FortiGuard Antivirus" on page 30.

Grayware

If the file passes the virus scan, it will be checked for grayware. Grayware configurations can be turned on and off as required and are kept up to date in the same manner as the antivirus definitions. For more information, see "Grayware scanning" on page 35.

Heuristics

After an incoming file has passed the grayware scan, it is subjected to the heuristics scan. The FortiGate heuristic antivirus engine, if enabled, performs tests on the file to detect virus-like behavior or known virus indicators. In this way, heuristic scanning may detect new viruses, but may also produce some false positive results. You configure heuristics from the CLI.

FortiGuard Antivirus

FortiGuard Antivirus services are an excellent resource which includes automatic updates of virus and IPS (attack) engines and definitions, as well as the local spam DNS black list (DNSBL),

through the FDN. The FortiGuard Center web site also provides the FortiGuard Antivirus virus and attack encyclopedia.

The connection between the FortiGate unit and FortiGuard Center is configured in *System > Config > FortiGuard*.

Enable antivirus scanning

Antivirus scanning is configured in an antivirus profile, but it is enabled in a firewall policy. Once the use of an antivirus profile is enabled and selected in one or more firewall policies, all the traffic controlled by those firewall policies will be scanned according to your settings.

Antivirus Profiles

From *Security Profiles > Antivirus > Profile* you can configure antivirus profiles that are then applied to firewall policies. A profile is specific configuration information that defines how the traffic within a policy is examined and what action may be taken based on the examination.

You can create multiple antivirus profiles for different antivirus scanning requirements. For example, you create an antivirus profile that specifies only virus scanning for POP3 which you then apply to the out-going firewall policy. You can also choose specific protocols, such as POP3, that will be blocked and then archived by the unit. This option is available only in the CLI.

Within antivirus profiles, you can also choose specific protocols to be blocked and then archive them. This is available only in the CLI.

To enable antivirus scanning — web-based manager

1. Go to *Security Profiles > AntiVirus > Profile*.

2. View and optionally change the *default* antivirus profile.

• You can also select *Create New* to create a new antivirus profile, or select an existing antivirus profile and choose *Edit*.

3. Select the inspection and the traffic you want scanned for viruses.

4. Select *OK*.

5. Go to *Policy > Policy > Policy* and either add or select the security policy that accepts the traffic to be virus scanned.

6. Turn on antivirus and select the profile that you configured.

7. Select OK to save the security policy.

To enable antivirus scanning — CLI

You need to configure the scan option for each type of traffic you want scanned. In this example, antivirus scanning of HTTP traffic is enabled in the profile.

```
config antivirus profile
  edit default
    config http
      set options scan
  end
end
```

Then enter a command similar to the following to add the default antivirus profile to a security policy.

```
config firewall policy
   edit 0
      set srcintf internal
      set dstintf wan1
      set srcaddr all
      set dstaddr all
      set schedule always
      set service ALL
      set action allow
      set utm-status enable
      set av-profile default
   end
```

Changing the default antivirus database

If your FortiGate unit supports extended, extreme, or flow-based virus database definitions, you can select the virus database most suited to your needs.

In most circumstances, the regular virus database provides sufficient protection. Viruses known to be active are included in the regular virus database. The extended database includes signatures of the viruses that have become rare within the last year in addition to those in the normal database. The extreme database includes legacy viruses that have not been seen in the wild in a long time in addition to those in the extended database.

The flow-based database contains a subset of the virus signatures in the extreme database. Unlike the other databases, selecting the flow-based database also changes the way the FortiGate unit scans your network traffic for viruses. Instead of the standard proxy-based scan, network traffic is scanned as it streams thought the FortiGate unit. For more information on the differences between flow-based and proxy-based antivirus scanning, see "How antivirus scanning works" on page 26.

If you require the most comprehensive antivirus protection, enable the extended virus database. The additional coverage comes at a cost, however, because the extra processing requires additional resources.

To change the antivirus database

```
config antivirus settings
   set default-db extended
end
```

Configuring the scan buffer size

When checking files for viruses using the proxy-based scanning method, there is a maximum file size that can be buffered. Files larger than this size are passed without scanning. The default size for all FortiGate models is 10 megabytes.

Archived files are extracted and email attachments are decoded before the FortiGate unit determines if they can fit in the scan buffer. For example, a 7 megabyte ZIP file containing a 12 megabyte EXE file will be passed without scanning with the default buffer size. Although the archive would fit within the buffer, the uncompressed file size will not.

In this example, the `uncompsizelimit` CLI command is used to change the scan buffer size to 20 megabytes for files found in HTTP traffic:

```
config antivirus service http
   set uncompsizelimit 20
end
```

The maximum buffer size varies by model. Enter `set uncompsizelimit?` to display the buffer size range for your FortiGate unit.

Flow-based scanning does not use a buffer and therefore has no file-size limit. File data is scanned as it passes through the FortiGate unit. The `uncompsizelimit` setting has no effect for flow-based scanning.

Configuring archive scan depth

The antivirus scanner will open archives and scan the files inside. Archives within other archives, or nested archives, are also scanned to a default depth of twelve nestings. You can adjust the number of nested archives to which the FortiGate unit will scan with the `uncompnestlimit` CLI command. Further, the limit is configured separately for each traffic type.

For example, this CLI command sets the archive scan depth for SMTP traffic to 5. That is, archives within archives will be scanned five levels deep.

```
config antivirus service smtp
   set uncompnestlimit 5
end
```

You can set the nesting limit from 2 to 100.

Configuring a maximum allowed file size

Proxy options allow you to enforce a maximum allowed file size for each of the network protocols in the profile. They are HTTP, FTP, IMAP, POP3, SMTP, IM, and NNTP. If your FortiGate unit supports SSL content scanning and inspection, you can also configure a maximum file size for HTTPS, IMAPS, POP3S, SMTPS, and FTPS.

The action you set determines what the FortiGate unit does with a file that exceeds the oversized file threshold. Two actions are available:

Block	Files that exceed the oversize threshold are dropped and a replacement message is sent to the user instead of the file.
Pass	Files exceed the oversized threshold are allowed through the FortiGate unit to their destination. Note that passed files are not scanned for viruses. File Filtering, both file pattern and file type, are applied, however.

You can also use the maximum file size to help secure your network. If you're using a proxy-based virus scan, the proxy scan buffer size limits the size of the files that can be scanned for infection. Files larger than this limit are passed without scanning. If you configure the maximum file size to block files larger than the scan buffer size, large infected files will not by-pass antivirus scanning.

In this example, the maximum file size will be configured to block files larger than 10 megabytes, the largest file that can be antivirus scanned with the default settings. You will need to configure a proxy options profile and add it to a security policy.

Set proxy options profile to block files larger than 10 MB

1. Go to *Policy > Policy > Proxy Options*.

2. Edit the default or select *Create New* to add a new one.

3. Scroll down to the common Options Section and place a check in the box next to *BlockOversized File/Email*

4. The sub line *Threshold (MB)* will appear with a value field. Enter 10.

5. Select *OK* or *Apply*.

The proxy options profile is configured, but to block files, you must select it in the firewall policies handling the traffic that contains the files you want blocked.

To select the Proxy Options profile in a security policy

1. Go to *Policy > Policy > Policy*.

2. Edit or create a security policy.

3. Select a proxy-based security profile. You will know that there is a proxy component to the Security Profile because when a Security Profile is Proxy based the *Proxy Options* field will be visible (for example, select an Antivirus profile that includes proxy scanning).

4. Beside *Proxy Options* select the name of the MTU proxy options protocol.

5. Select *OK* to save the security policy.

Once you complete these steps, any files in the traffic subject to Security Profile scanning handled by this policy that are larger than 10MB will be blocked. If you have multiple firewall policies, examine each to determine if you want to apply similar file blocking the them as well.

Configuring client comforting

When proxy-based antivirus scanning is enabled, the FortiGate unit buffers files as they are downloaded. Once the entire file is captured, the FortiGate unit scans it. If no infection is found, the file is sent along to the client. The client initiates the file transfer and nothing happens until the FortiGate finds the file clean, and releases it. Users can be impatient, and if the file is large or the download slow, they may cancel the download, not realizing that the transfer is in progress.

The client comforting feature solves this problem by allowing a trickle of data to flow to the client so they can see the file is being transferred. The default client comforting transfer rate sends one byte of data to the client every ten seconds. This slow transfer continues while the FortiGate unit buffers the file and scans it. If the file is infection-free, it is released and the client will receive the remainder of the transfer at full speed. If the file is infected, the FortiGate unit caches the URL and drops the connection. The client does not receive any notification of what happened because the download to the client had already started. Instead, the download stops and the user is left with a partially downloaded file.

If the user tries to download the same file again within a short period of time, the cached URL is matched and the download is blocked. The client receives the Infection cache message

replacement message as a notification that the download has been blocked. The number of URLs in the cache is limited by the size of the cache.

> Client comforting can send unscanned and therefore potentially infected content to the client. You should only enable client comforting if you are prepared to accept this risk. Keeping the client comforting interval high and the amount low will reduce the amount of potentially infected data that is downloaded.

Client comforting is available for HTTP and FTP traffic. If your FortiGate unit supports SSL content scanning and inspection, you can also configure client comforting for HTTPS and FTPS traffic.

Enable and configure client comforting

1. Go to *Policy > Policy > Proxy Options*.

2. Select a Proxy Options profile and choose *Edit*, or select *Create New* to make a new one.

3. Scroll down to the *Common Options* section and check the box next to *Comfort Clients*. This will set the option on all of the applicable protocols. The ability to set this feature on a protocol by protocol basis exists in the CLI

4. Select *OK or Apply* to save the changes.

5. Select this Proxy Options profile in any security policy for it to take effect on all traffic handled by the policy.

The default values for Interval and Amount are 10 and 1, respectively. This means that when client comforting takes effect, 1 byte of the file is sent to the client every 10 seconds. You can change these values to vary the amount and frequency of the data transferred by client comforting.

Grayware scanning

Grayware programs are unsolicited software programs installed on computers, often without the user's consent or knowledge. Grayware programs are generally considered an annoyance, but they can also cause system performance problems or be used for malicious purposes.

To allow the FortiGate unit to scan for known grayware programs, you must enable both antivirus scanning and grayware detection.

Enter the following command to enable grayware detection:

```
config antivirus settings
   set grayware enable
end
```

With grayware detection enabled, the FortiGate unit will scan for grayware any time it checks for viruses.

Windows file sharing (CIFS) flow-based antivirus scanning

FortiOS 5.0 now supports virus scanning of Windows file sharing traffic. This includes CIFS, SMB, and SAMBA traffic. This feature is applied by enabling SMB scanning in an antivirus profile and then adding this profile to a security policy that accepts CIFS traffic. CIFS virus scanning is available only through flow-based antivirus scanning.

FortiOS 5.0 flow-based virus scanning can detect the same number of viruses in CIFS/SMB/SAMBA traffic as it can for all supported content protocols.

Figure 5: Configuring CIFS/SMB/SAMBA virus scanning

Use the following command to enable CIFS/SMB/SAMBA virus scanning in an antivirus profile:

```
config antivirus profile
  edit smb-profile
    config smb
      set options scan
      set avdb flow-based
    end
```

Then add this antivirus profile to a security policy that accepts the traffic to be virus scanned. In the security policy the service can be set to ANY, SAMBA, or SMB.

```
config firewall policy
  edit 0
    set service ANY
    ...
    set utm-status enable
    set av-profile smb-profile
  end
```

Note the following about CFIS/SMB/SAMBA virus scanning:

- Some newer version of SAMBA clients and SMB2 can spread one file across multiple sessions, preventing some viruses from being detected if this occurs.
- Enabling CIFS/SMB/SAMBA virus scanning can affect FortiGate performance.
- SMB2 is a new version of SMB that was first partially implemented in Windows Vista.
- Currently SMB2 is supported by Windows Vista or later, and partly supported by Samba 3.5 and fully support by Samba 3.6.
- The latest version of SMB2.2 will be introduced with Windows 8.
- Most clients still use SMB as default setting.

Advanced Persistent Threat (APT) protection

New advanced persistent threat (APT) protection features in FortiOS 5.0 include botnet protection, phishing protection, and zero-day threat protection using FortiGuard Analytics for sandboxing.

Botnet and phishing protection

In an antivirus profile you can configure the FortiGate unit to detect and block botnet connection attempts. This feature also blocks attempted access to phishing URLs.

The antivirus database includes a constantly updated database of known command and control (C&C) sites that Botnet clients attempt to connect too as well as a database of phishing URLs.

To enable Botnet and phishing protection in an antivirus profile select Block Connections to Botnet Servers. Botnet protection is available for proxy and flow-based antivirus profiles.

Figure 6: Adding Botnet and phishing protection.

FortiGuard Sandbox (in the cloud sandboxing, zero day threat analysis and submission)

In a Proxy Mode antivirus profile, enabling Send Files to *FortiGuard Sandbox for Inspection* to causes your FortiGate unit to upload files to FortiGuard where the file will be executed and the resulting behavior analyzed for risk. You have the choice of uploading all files or only the suspicious ones. If the file exhibits risky behavior or is found to contain a virus, a new virus signature is created and added to the FortiGuard antivirus signature database. The next time your FortiGate unit updates its antivirus database it will have the new signature.

Currently, a file is considered suspicious if it does not contain a known virus and if it has some suspicious characteristics. The suspicious characteristics can change depending on the current threat climate and other factors. Fortinet optimizes how files are uploaded as required.

The FortiGuard Sandbox feature is available if you have a valid FortiCloud subscription. To verify whether or not a subscription is associated with your FortiGate go to *System > Dashboard > Status* and check the License Information widget in the FortiCloud subsection.

Figure 7: Enabling FortiGuard Sandbox in an Antivirus Profile

On the FortiGate, there are two ways to verify that files are being uploaded to the FortiCloud Sandbox. The first is to go to *System > Config > FortiSandbox*. The window is for configuring whether or not the FortiGate unit is to use the FortiCloud Sandbox or a FortiSandbox Appliance but it also shows the statistics of files submitted to the Sandbox over the last seven days.

The second method is to got System > Dashboard > Status and view the *Advanced Threat Protection Statistics* dashboard widget. This widget will show essentially the same information.This widget is not one of the default ones so you will have to add it to the Dashboard.

Figure 8: Example Advanced Threat Protection Statistics widget showing Sandbox submissions

To view information relating to the Antivirus function from the FortiCloud side, go to *System > Dashboard > Status* and look at the *License Information* widget. In the *FortiCloud* subsection in the *Account* line, select the *Launch Portal* link. Once at the portal select the icon for the specific FortiGate that you view the information for.

Under the Logs & Archives tab of the menu bar you will find the UTM option. Once this option is selected, you will have the option of choosing AntiVirus. The site will display records within the designated time frame that refer to AntiVirus events recorded by the logs.

Figure 9: Example view of FortiCloud's AntiVirus logs

In addition to the normal UTM logs, there is a new menu item in that top menu bar that appears when your FortiGate is configured to submit files to the FortiSandbox. This page on the site will

display more granular information on files with viruses that are submitted by your FortiGate unit.This information will include:

- Date and Time
- File Name
- User Name
- Service
- Source IP
- Destination IP
- Vdom
- Analysis
- URL

Testing your antivirus configuration

You have configured your FortiGate unit to stop viruses, but you'd like to confirm your settings are correct. Even if you have a real virus, it would be dangerous to use for this purpose. An incorrect configuration will allow the virus to infect your network.

To solve this problem, the European Institute of Computer Anti-virus Research has developed a test file that allows you to test your antivirus configuration. The EICAR test file is not a virus. It can not infect computers, nor can it spread or cause any damage. It's a very small file that contains a sequence of characters. Your FortiGate unit recognizes the EICAR test file as a virus so you can safely test your FortiGate unit antivirus configuration.

Go to http://www.fortiguard.com/antivirus/eicartest.html to download the test file (eicar.com) or the test file in a ZIP archive (eicar.zip).

If the antivirus profile applied to the security policy that allows you access to the Web is configured to scan HTTP traffic for viruses, any attempt to download the test file will be blocked. This indicates that you are protected.

Antivirus examples

The following examples provide a sample antivirus configuration scenario for a fictitious company.

Configuring simple antivirus protection

Small offices, whether they are small companies, home offices, or satellite offices, often have very simple needs. This example details how to enable antivirus protection on a FortiGate unit located in a satellite office. The satellite office does not have an internal email server. To send and retrieve email, the employees connect to an external mail server.

Creating an antivirus profile

Most antivirus settings are configured in an antivirus profile. Antivirus profiles are selected in firewall policies. This way, you can create multiple antivirus profiles, and tailor them to the traffic controlled by the security policy in which they are selected. In this example, you will create one antivirus profile.

To create an antivirus profile — web-based manager

1. Go to *Security Profiles > AntiVirus > Profiles*.

2. Select *Create New*.

3. In the *Name* field, enter `basic_antivirus`.

4. In the *Comments* field, enter `Antivirus protection for web and email traffic`.

5. Select the *Virus Scan* check boxes for the *HTTP*, *IMAP*, *POP3*, and *SMTP* traffic types.

6. Select *OK* to save the antivirus profile.

To create an antivirus profile — CLI

```
config antivirus profile
  edit basic_antivirus
    set comment "Antivirus protection for web and email traffic"
    config http
      set options scan
    end
    config imap
      set options scan
    end
    config pop3
      set options scan
    end
    config smtp
      set options scan
    end
  end
```

Selecting the antivirus profile in a security policy

An antivirus profile directs the FortiGate unit to scan network traffic only when it is selected in a security policy. When an antivirus profile is selected in a security policy, its settings are applied to all the traffic the security policy handles.

To select the antivirus profile in a security policy — web-based manager

1. Go to *Policy > Policy > Policy*.

2. Create a new or edit a security policy.

3. Turn on Antivirus.

4. Select an antivirus profile.

5. Select *OK* to save the security policy.

To select the antivirus profile in a security policy — CLI

```
config firewall policy
  edit 1
    set utm-status enable
    set profile-protocol-options default
    set av-profile basic_antivirus
  end
```

HTTP, IMAP, POP3, and SMTP traffic handled by the security policy you modified will be scanned for viruses. A small office may have only one security policy configured. If you have multiple policies, consider enabling antivirus scanning for all of them.

Protecting your network against malicious email attachments

Grayware is commonly delivered by email or the web. The Example.com corporation has been the victim of multiple greyware infections in the past. Now that the company has a FortiGate unit protecting its network, you (Example.com's system administrator) can configure the unit to scan email and web traffic to filter out greyware attachments.

Enabling antivirus scanning in the antivirus profile

The primary means to avoid viruses is to configure the FortiGate unit to scan email and web traffic for virus signatures. You enable virus scanning in the antivirus profile and then select the antivirus profile in firewall policies that control email traffic.

To enable antivirus scanning in the antivirus profile

1. Go to *Security Profiles > AntiVirus > Profiles*.
2. Create a new or edit an antivirus profile.
3. Select *Virus Scan and Removal* for *HTTP* to scan web traffic for viruses.
4. Select the *Virus Scan* check box for *IMAP*, *POP3*, and *SMTP* to scan all email protocols for viruses.
5. Select *OK* or *Apply* to save the antivirus profile.

Selecting the antivirus profile in a security policy

An antivirus profile directs the FortiGate unit to scan network traffic only when it is selected in a security policy. When an antivirus profile is selected in a security policy, its settings are applied to all the traffic the security policy handles.

To select the antivirus profile in a security policy

1. Go to *Policy > Policy > Policy*.
2. Create or edit a policy that controls the network traffic controlling email traffic.
3. Turn on Antivirus.
4. Select an antivirus profile.
5. Select *OK* to save the security policy.

Email filter

This section describes how to configure FortiGate email filtering for IMAP, POP3, and SMTP email. Email filtering includes both spam filtering and filtering for any words or files you want to disallow in email messages. If your FortiGate unit supports SSL content scanning and inspection, you can also configure spam filtering for IMAPS, POP3S, and SMTPS email traffic.

The following topics are included in this section:

- Email filter concepts
- Enable email filtering
- Configure email traffic types to inspect
- Configure the spam action
- Configure the tag location
- Configure the tag format
- Configure FortiGuard email filters
- Configure local email filters
- Email filter examples

Email filter concepts

You can configure the FortiGate unit to manage unsolicited commercial email by detecting and identifying spam messages from known or suspected spam servers.

The FortiGuard Antispam Service uses both a sender IP reputation database and a spam signature database, along with sophisticated spam filtering tools, to detect and block a wide range of spam messages. Using FortiGuard Antispam email filter profile settings, you can enable IP address checking, URL checking, email checksum checking, and spam submission. Updates to the IP reputation and spam signature databases are provided continuously via the global FortiGuard Distribution Network.

From the FortiGuard Antispam Service page in the FortiGuard Center, you can find out whether an IP address is blacklisted in the FortiGuard antispam IP reputation database, or whether a URL or email address is in the signature database.

Email filter techniques

The FortiGate unit has a number of techniques available to help detect spam. Some use the FortiGuard Antispam Service and require a subscription. The remainder use your DNS servers or use lists that you must maintain.

FortiGuard IP address check

The FortiGate unit queries the FortiGuard Antispam Service to determine if the IP address of the client delivering the email is blacklisted. A match will cause the FortiGate unit to treat delivered messages as spam.

The default setting of the `smtp-spamhdrip` CLI command is `disable`. If enabled, the FortiGate unit will check all the IP addresses in the header of SMTP email against the FortiGuard Antispam Service. For more information, see the *FortiGate CLI Reference*.

FortiGuard URL check

The FortiGate unit queries the FortiGuard Antispam service to determine if any URL in the message body is associated with spam. If any URL is blacklisted, the FortiGate unit determines that the email message is spam.

Detect phishing URLs in email

The FortiGate unit sends the URL links in email messages to FortiGuard to determine if the links are associated with a known phishing site. If such a link is detected, the link is removed from the message. The URL remains, but it is no longer a selectable hyperlink.

FortiGuard email checksum check

The FortiGate unit sends a hash of an email to the FortiGuard Antispam server, which compares the hash to hashes of known spam messages stored in the FortiGuard Antispam database. If the hash results match, the email is flagged as spam.

FortiGuard spam submission

Spam submission is a way you can inform the FortiGuard AntiSpam service of non-spam messages incorrectly marked as spam. When you enable this setting, the FortiGate unit adds a link to the end of every message marked as spam. You then select this link to inform the FortiGuard AntiSpam service when a message is incorrectly marked.

IP address black/white list check

The FortiGate unit compares the IP address of the client delivering the email to the addresses in the IP address black/white list specified in the email filter profile. If a match is found, the FortiGate unit will take the action configured for the matching black/white list entry against all delivered email.

The default setting of the `smtp-spamhdrip` CLI command is `disable`. If enabled, the FortiGate unit will check all the IP addresses in the header of SMTP email against the specified IP address black/white list. For more information, see the *FortiGate CLI Reference*.

HELO DNS lookup

The FortiGate unit takes the domain name specified by the client in the HELO greeting sent when starting the SMTP session and does a DNS lookup to determine if the domain exists. If the lookup fails, the FortiGate unit determines that any messages delivered during the SMTP session are spam.

Email address black/white list check

The FortiGate unit compares the sender email address, as shown in the message envelope MAIL FROM, to the addresses in the email address black/white list specified in the email filter profile. If a match is found, the FortiGate unit will take the action configured for the matching black/white list entry.

Return email DNS check

The FortiGate unit performs a DNS lookup on the reply-to domain to see if there is an A or MX record. If no such record exists, the message is treated as spam.

Banned word check

The FortiGate unit blocks email messages based on matching the content of the message with the words or patterns in the selected spam filter banned word list. This feature is only available in the CLI.

Order of spam filtering

The FortiGate unit checks for spam using various filtering techniques. The order in which the FortiGate unit uses these filters depends on the mail protocol used.

Filters requiring a query to a server and a reply (FortiGuard Antispam Service and DNSBL/ORDBL) are run simultaneously. To avoid delays, queries are sent while other filters are running. The first reply to trigger a spam action takes effect as soon as the reply is received.

Each spam filter passes the email to the next if no matches or problems are found. If the action in the filter is *Mark as Spam*, the FortiGate unit tags the email as spam according to the settings in the email filter profile.

For SMTP and SMTPS, if the action is discard, the email message is discarded or dropped.

If the action in the filter is *Mark as Clear*, the email is exempt from any remaining filters. If the action in the filter is *Mark as Reject*, the email session is dropped. Rejected SMTP or SMTPS email messages are substituted with a configurable replacement message.

Order of SMTP and SMTPS spam filtering

The FortiGate unit scans SMTP and SMTPS email for spam in the order given below. SMTPS spam filtering is available on FortiGate units that support SSL content scanning and inspection.

1. IP address black/white list (BWL) check on last hop IP
2. DNSBL & ORDBL check on last hop IP, FortiGuard Antispam IP check on last hop IP, HELO DNS lookup
3. MIME headers check, E-mail address BWL check
4. Banned word check on email subject
5. IP address BWL check (for IPs extracted from "Received" headers)
6. Banned word check on email body
7. Return email DNS check, FortiGuard Antispam email checksum check, FortiGuard Antispam URL check, DNSBL & ORDBL check on public IP extracted from header.

Order of IMAP, POP3, IMAPS and POP3S spam filtering

The FortiGate unit scans IMAP, POP3, IMAPS and POP3S email for spam in the order given below. IMAPS and POP3S spam filtering is available on FortiGate units that support SSL content scanning and inspection.

1. MIME headers check, E-mail address BWL check
2. Banned word check on email subject
3. IP BWL check
4. Banned word check on email body
5. Return email DNS check, FortiGuard Antispam email checksum check, FortiGuard Antispam URL check, DNSBL & ORDBL check.

Enable email filtering

Unlike antivirus protection, no single control enables all email filtering. Your FortiGate unit uses many techniques to detect spam; some may not be appropriate for every situation. For this reason, when you enable email filtering, you must then choose when techniques are applied to email traffic.

To enable email filtering

1. Go to *Security Profiles > Email Filter > Profile*.

   The default email filter profile is presented. You can edit this profile or create a new one.

2. Select the *Inspection Mode*.

   Proxy detection involves buffering the file and examining it as a whole. Advantages of proxy-based detection include a more thorough examination of attachments, especially archive formats and nesting.

   Flow-based detection examines the file as it passes through the FortiGate unit without any buffering. Advantages of flow-based detection include speed and no interruption of detection during conserve mode.

3. Select *Enable Spam Detection and Filtering*.

4. If you wish to leave everything in it's default setting you can select *OK* or *Apply.*

Once you have enabled the email filter you can further specify what protocols to inspect.

Configure email traffic types to inspect

The FortiGate unit examines IMAP, POP3, and SMTP email traffic. If your FortiGate unit supports content inspection, it can also examine IMAPS, POP3S, and SMTPS traffic. The options that you will see in the profile window are IMAP, POP3 and SMTP

To select the email traffic types to inspect

1. Go to *Security Profiles > Email Filter > Profile*.

2. The default email filter profile is presented. To edit another profile, select it from the drop down in the *Edit Email Filter Profile* title bar.

3. Select *Enable Spam Detection and Filtering*.

4. Select the types of email that you want the FortiGate unit to examine when using this email filter profile.

5. Select *Apply*.

The traffic types you enable will be examined according to the settings in the email filter profile.

Configure the spam action

When spam is detected, the FortiGate unit will deal with it according to the *Spam Action* setting in the email filter profile. Note that POP3S, IMAPS and SMTPS spam filtering is available only on FortiGate units that support SSL content scanning and inspection. POP3, IMAP, POP3S and IMAPS mail can only be tagged. SMTP and SMTPS mail can be set to *Discard* or *Tagged*:

- **Discard**: When the spam action is set to *Discard*, messages detected as spam are deleted. No notification is sent to the sender or recipient.

- **Tagged**: When the spam action is set to *Tagged*, messages detected as spam are labelled and delivered normally. The text used for the label is set in the *Tag Format* field and the label is placed in the subject or the message header, as set with the *Tag Location* option.

To configure the spam action

1. Go to *Security Profiles > Email Filter > Profile*.

2. The default email filter profile is presented. To edit another profile, select it from the drop down in the *Edit Email Filter Profile* title bar.

3. Select *Enable Spam Detection and Filtering*.

4. Select the types of email that you want the FortiGate unit to examine when using this email filter profile.

5. Select *Apply*.

6. The *Spam Action* row has a drop-down selection under the SMTP traffic type. Select *Discard* or *Tagged*.

   No selection is available for POP3 or IMAP traffic. *Tagged* is the only applicable action for those traffic types.

   By default, the tag location for any traffic set to *Tagged* is *Subject* and the tag format is `Spam`. If you want to change these settings, continue with .

7. Select *Apply*.

Select the edited email filter profile in a security policy, and the traffic controlled by the security policy will be scanned according to the settings you configured. You may select the email filter profile in more than one security policy if required.

Configure the tag location

When the spam action is set to *Tagged*, the *Tag Location* setting determines where the tag is applied in the message.

To configure the tag location

1. Go to *Security Profiles > Email Filter > Profile*.

2. The default email filter profile is presented. To edit another profile, select it from the drop down in the *Edit Email Filter Profile* title bar.

3. Select *Enable Spam Detection and Filtering*.

4. Select the types of email that you want the FortiGate unit to examine when using this email filter profile.

5. Select *Apply*.

6. The *Tag Location* row has two options for each traffic type. Note that if the spam action for SMTP traffic is set to discard, the tag location will not be available. Select the tag location:

• *Subject*: The FortiGate unit inserts the tag at the beginning of the message subject. For example, if the message subject is "Buy stuff!" and the tag is "[spam]", the new message subject is "[spam] Buy stuff!" if the message is detected as spam.

• *MIME*: The FortiGate unit inserts the tag into the message header. With most mail readers and web-based mail services, the tag will not be visible. Despite this, you can still set up a rule based on the presence or absence of the tag.

7. Select *Apply*.

Configure the tag format

When the spam action is set to *Tagged*, the *Tag Format* setting determines what text is used as the tag applied to the message.

To configure the tag format

1. Go to *Security Profiles > Email Filter > Profile*.

2. The default email filter profile is presented. To edit another profile, select it from the drop down in the *Edit Email Filter Profile* title bar.

3. Select *Enable Spam Detection and Filtering*.

4. Select the types of email that you want the FortiGate unit to examine when using this email filter profile.

5. Select *Apply*.

6. The *Tag Format* row has a field for each traffic type. Note that if the spam action for SMTP traffic is set to discard, the tag format will not be available.Enter the text the FortiGate unit will use as the tag for each traffic type.

7. Select *Apply*.

Configure FortiGuard email filters

FortiGuard email filtering techniques us FortiGuard services to detect the presence of spam among your email. A FortiGuard subscription is required to use the FortiGuard email filters. You can enable the following types of FortiGuard email filtering:

FortiGuard IP address checking	When you enable FortiGuard IP address checking, your FortiGate unit will submit the IP address of the client to the FortiGuard service for checking. If the IP address exists in the FortiGuard IP address black list, your FortiGate unit will treat the message as spam.
FortiGuard URL checking	When you enable FortiGuard URL checking, your FortiGate unit will submit all URLs appearing in the email message body to the FortiGuard service for checking. If a URL exists in the FortiGuard URL black list, your FortiGate unit will treat the message as spam.
FortiGuard phishing URL detection	When you enable FortiGuard phishing URL detection, your FortiGate unit will submit all URL hyperlinks appearing in the email message body to the FortiGuard service for checking. If a URL exists in the FortiGuard URL phishing list, your FortiGate unit will remove the hyperlink from the message. The URL will remain in place, but it will no longer be a selectable hyperlink.
FortiGuard email checksum checking	When you enable FortiGuard email checksum checking, your FortiGate unit will submit a checksum of each email message to the FortiGuard service for checking. If a checksum exists in the FortiGuard checksum black list, your FortiGate unit will treat the message as spam.
FortiGuard spam submission	When you enable FortiGuard email checksum checking, your FortiGate unit will append a link to the end of every message detected as spam. This link allows email users to "correct" the FortiGuard service by informing it that the message is not spam.

Carefully consider the use of the *Spam submission* option on email leaving your network. Users not familiar with the feature may click the link on spam messages because they are curious. This will reduce the accuracy of the feature.

To enable FortiGuard email filtering

1. Go to *Security Profiles > Email Filter > Profile*.

2. The default email filter profile is presented. To edit another profile, select it from the drop down in the *Edit Email Filter Profile* title bar.

3. Select *Enable Spam Detection and Filtering*.

4. Select the types of email that you want the FortiGate unit to examine when using this email filter profile.

5. Select *Apply*.

6. Under the heading *FortiGuard Spam Filtering*, select one or more of the following options:

• *IP Address Check*.

• *URL Check*.

• *Detect Phishing URLs in Email*.

• *E-mail Checksum Check*.

• *Spam Submission*.

7. Select *Apply*.

Select the edited email filter profile in a security policy, and the traffic controlled by the security policy will be scanned according to the settings you configured. You may select the email filter profile in more than one security policy if required.

Configure local email filters

Local email filtering techniques us your own resources, whether DNS checks or IP address and email address lists that you maintain. You can enable three types of local filtering:

• Black and white list (BWL) checking (includes email addresses and IP addresses)

• HELO DNS lookup

• Return email DNS checking

Enabling IP address and email address black/white list checking

When you email enable black/white list (BWL) checking, your FortiGate unit will perform IP address BWL checking and email address BWL checking.

IP address BWL checking matches client IP addresses with IP addresses in the selected email BWL list and acts according to the action configured for the IP address in the list: allow the message, reject it, or mark it as spam.

Email address BWL checking matches sender email addresses with email addresses in the selected email BWL list acts according to the action configured for the email address in the list: allow the message or mark it as spam.

Before you can enable IP address and email address black/white list spam filtering you must create an email black/white list.

To create an email black/white list

1. Go to *Security Profiles > Email Filter > Email List*.

2. Select *Create New*.

3. Enter a name for the BWL list.

4. Optionally, enter a description or comments about the list.

5. Select *OK* to save the list.

When a new back/white list is created, it is empty. To perform any actions, you must add IP and email addresses to the list.

To add an IP address to an email black/white list

1. Go to *Security Profiles > Email Filter > Email List*.

2. Edit a list.

3. Select *Create New*.

4. Select *IP/Netmask*.

5. Enter the IP address or netmask in the IP/netmask field.

6. Select the action:

- *Mark as Clear*: Messages from clients with matching IP addresses will be allowed, bypassing further email filtering.

- *Mark as Reject*: Messages from clients with matching IP addresses will be rejected. The FortiGate unit will return a reject message to the client. *Mark as Reject* only applies to mail delivered by SMTP. If an IP address black/white list is used with POP3 or IMAP mail, addresses configured with the *Mark as Reject* action will be marked as spam.

- *Mark as Spam*: Messages from clients with matching IP addresses will be treated as spam, subject to the action configured in the applicable email filter profile. For more information, see "Configure the spam action" on page 46.

7. By default, the address is enabled and the FortiGate unit will perform the action if the address is detected. To disable checking for the address, clear the *Enable* check box.

8. Select *OK*.

To add an email address to an email black/white list

1. Go to *Security Profiles > Email Filter > Email List*.

2. Edit a list.

3. Select *Create New*.

4. Select *Email Address*.

5. Enter the email address in the *Email Address* field.

6. If you need to enter a pattern in the *Email Address* field, select whether to use wildcards or regular expressions to specify the pattern.

   Wildcard uses an asterisk ("*") to match any number of any character. For example, *@example.com will match all addresses ending in @example.com.

   Regular expressions use Perl regular expression syntax. See http://perldoc.perl.org/perlretut.html for detailed information about using Perl regular expressions.

7. Select the action:

- *Mark as Spam*: Messages with matching reply-to email addresses will be treated as spam, subject to the action configured in the applicable email filter profile. For more information, see "Configure the spam action" on page 46.

- *Mark as Clear*: Messages with matching reply-to addresses will be allowed, bypassing further email filtering.

8. By default, the address is enabled and the FortiGate unit will perform the action if the address is detected. To disable checking for the address, clear the *Enable* check box.

9. Select *OK* to save the address.

To enable IP address black/white list checking

1. Go to *Security Profiles > Email Filter > Profile*.

2. The default email filter profile is presented. To edit another profile, select it from the drop down in the *Edit Email Filter Profile* title bar.

3. Select *Enable Spam Detection and Filtering* and select Apply.

4. Under the heading *Local Spam Filtering*, select *BWL Check*.

5. Select the IP address black/white list to use from the drop-down list.

6. Select *Apply*.

Select the email filter profile in a security policy, and the traffic accepted by the security policy will be scanned according to the settings you configured.

Enabling HELO DNS lookup

Whenever a client opens an SMTP session with a server, the client sends a HELO command with the client domain name. When you enable HELO DNS lookup, your FortiGate unit will take the domain the client submits as part of the HELO greeting and send it to the configured DNS. If the domain does not exist, your FortiGate unit will treat all messages the client delivers as spam.

The HELO DNS lookup is available only for SMTP traffic.

To enable HELO DNS lookup

1. Go to *Security Profiles > Email Filter > Profile*.

2. The default email filter profile is presented. To edit another profile, select it from the drop down in the *Edit Email Filter Profile* title bar.

3. Select *Enable Spam Detection and Filtering* and select Apply

4. Under the heading *Local Spam Filtering*, select *HELO DNS Lookup*.

5. Select *Apply*.

Select the edited email filter profile in a security policy, and the traffic controlled by the security policy will be scanned according to the settings you configured. You may select the email filter profile in more than one security policy if required.

Enabling return email DNS checking

When you enable return email DNS checking, your FortiGate unit will take the domain in the reply-to email address and send it to the configured DNS. If the domain does not exist, your FortiGate unit will treat the message as spam.

To enable return email DNS check

1. Go to *Security Profiles > Email Filter > Profile*.

2. The default email filter profile is presented. To edit another profile, select it from the drop down in the *Edit Email Filter Profile* title bar.

3. Select *Enable Spam Detection and Filtering* and select Apply

4. Under the heading *Local Spam Filtering*, select *Return E-mail DNS Check*.

5. Select *Apply*.

Select the edited email filter profile in a security policy, and the traffic controlled by the security policy will be scanned according to the settings you configured. You may select the email filter profile in more than one security policy if required.

Enabling banned word checking

When you enable banned word checking, your FortiGate unit will examine the email message for words appearing in the banned word list specified in the email filter profile. If the total score of the banned word discovered in the email message exceeds the threshold value set in the email filter profile, your FortiGate unit will treat the message as spam.

When determining the banned word score total for an email message, each banned word score is added once no matter how many times the word appears in the message. Use the command `config spamfilter bword` to add an email banned word list. Use the command `config spamfilter profile` to add a banned word list to an email filtering profile.

How content is evaluated

Every time the banned word filter detects a pattern in an email message, it adds the pattern score to the sum of scores for the message. You set this score when you create a new pattern to block content. The score can be any number from zero to 99999. Higher scores indicate more offensive content. When the total score equals or exceeds the threshold, the email message is considered as spam and treated according to the spam action configured in the email filter profile. The score for each pattern is counted only once, even if that pattern appears many times in the email message. The default score for banned word patterns is 10 and the default threshold is 10. This means that by default, an email message is blocked by a single match.

A pattern can be part of a word, a whole word, or a phrase. Multiple words entered as a pattern are treated as a phrase. The phrase must appear as entered to match. You can also use wildcards or regular expressions to have a pattern match multiple words or phrases.

For example, the FortiGate unit scans an email message that contains only this sentence: "The score for each word or phrase is counted only once, even if that word or phrase appears many times in the email message."

Banned word pattern	Pattern type	Assigned score	Score added to the sum for the entire page	Comment
word	Wildcard	20	20	The pattern appears twice but multiple occurrences are only counted once.
word phrase	Wildcard	20	0	Although each word in the phrase appears in the message, the words do not appear together as they do in the pattern. There are no matches.
word*phrase	Wildcard	20	20	The wildcard represents any number of any character. A match occurs as long as "word" appears before "phrase" regardless of what is in between them.
mail*age	Wildcard	20	20	Since the wildcard character can represent any characters, this pattern is a match because "email message" appears in the message.

In this example, the message is treated as spam if the banned word threshold is set to 60 or less.

Adding words to a banned word list

Each banned word list contains a number of words, each having a `score`, and specifying where the FortiGate unit will search for the word (in the message subject, message `body`, or `all` which means both)

When the FortiGate unit accepts an email message containing one or more words in the banned word list specified in the active email filter profile, it totals the scores of the banned words in the email message. If the total is higher than the threshold set in the email filter profile, the email message will be detected as spam. If the total score is lower than the threshold, the message will be allowed to pass as normal.

The score of a banned word present in the message will be counted toward the score total only once, regardless of how many times the word appears in the message.

When you enter a word, set the `Pattern-type` to wildcards or regular expressions.

Wildcard uses an asterisk ("*") to match any number of any character. For example, re* will match all words starting with "re".

Regular expression uses Perl regular expression syntax. See http://perldoc.perl.org/perlretut.html for detailed information about using Perl regular expressions.

Email filter examples

Configuring simple antispam protection

Small offices, whether they are small companies, home offices, or satellite offices, often have very simple needs. This example details how to enable antispam protection on a FortiGate unit located in a satellite office.

Creating an email filter profile

Most email filter settings are configured in an email filter profile. Email filter profiles are selected in firewall policies. This way, you can create multiple email filter profiles, and tailor them to the traffic controlled by the security policy in which they are selected. In this example, you will create one email filter profile.

To create an email filter profile — web-based manager

1. Go to *Security Profiles > Email Filter > Profile*.
2. Select the *Create New* icon in the Edit Email Filter Profile window title.
3. In the *Name* field, enter `basic_emailfilter`.
4. Select *Enable Spam Detection and Filtering*.
5. Ensure that *IMAP*, *POP3*, and *SMTP* are selected in the header row.

   These header row selections enable or disable examination of each email traffic type. When disabled, the email traffic of that type is ignored by the FortiGate unit and no email filtering options are available.
6. Under *FortiGuard Spam Filtering*, enable *IP Address Check*.
7. Under *FortiGuard Spam Filtering*, enable *URL Check*.
8. Under *FortiGuard Spam Filtering*, enable *E-mail Checksum Check*.
9. Select *OK* to save the email filter profile.

To create an email filter profile — CLI

```
config spamfilter profile
  edit basic_emailfilter
    set options spamfsip spamfsurl spamfschksum
  end
```

Selecting the email filter profile in a security policy

An email filter profile directs the FortiGate unit to scan network traffic only when it is selected in a security policy. When an email filter profile is selected in a security policy, its settings are applied to all the traffic the security policy handles.

To select the email filter profile in a security policy — web-based manager

1. Go to *Policy > Policy > Policy*.

2. Create a new or edit a policy.

3. Turn on email filtering.

4. Select the `basic_emailfilter` profile from the list.

5. Select *OK* to save the security policy.

To select the email filter profile in a security policy — CLI

```
config firewall policy
  edit 1
    set utm-status enable
    set profile-protocol-options default
    set spamfilter-profile basic_emailfilter
  end
```

IMAP, POP3, and SMTP email traffic handled by the security policy you modified will be scanned for spam. Spam messages have the text "Spam" added to their subject lines. A small office may have only one security policy configured. If you have multiple policies, consider enabling spam scanning for all of them.

Blocking email from a user

Employees of the Example.com corporation have been receiving unwanted email messages from a former client at a company called example.net. The client's email address is client@example.net. All ties between the company and the client have been severed, but the messages continue. The FortiGate unit can be configured to prevent these messages from being delivered.

To create the email address list

1. Go to *Security Profiles > Email Filter > Email List*.

2. Select *Create New*.

3. Enter a name for the new email address list.

4. Optionally, enter a descriptive comment for the email address list.

5. Select *OK* to create the list.

6. Select *Create New* to add a new entry to the email address list.

7. Select *Email Address*.

8. Enter `client@example.net` in the *E-mail Address* field.

• If you wanted to prevent everyone's email from the client's company from getting through you could have used *@example.net instead.

9. Leave *Pattern Type* set to the default, *Wildcard*.

10. Leave *Action* as *Mark as Spam* to have the FortiGate unit mark all messages from example.net as spam.

Now that the email address list is created, you must enable the email filter in the email filter profile.

To enable Email Filter

1. Go to *Security Profiles > Email Filter > Profile*.

2. Select the email filter profile that is used by the firewall policies handling email traffic from the email filter profile drop down list.

3. In the row *Tag Location*, select *Subject* for all three mail protocols.

4. In the row *Tag Format*, enter `SPAM:` in all three fields.

5. Select *Enable Spam Detection and Filtering*.

6. Ensure that the check boxes labeled *IMAP*, *POP3*, and *SMTP* in the header row are selected.

7. Under *Local Spam Filtering*, enable *BWL Check* and select the email address list you created in the previous procedure from the drop down list.

8. Select *OK*.

When this email filter profile is selected in a security policy, the FortiGate unit will add "SPAM:" to the subject of any email message from an address ending with @example.net for all email traffic handled by the security policy. Recipients can ignore the message or they can configure their email clients to automatically delete messages with "SPAM:" in the subject.

Intrusion protection

The FortiGate Intrusion Protection system combines signature detection and prevention with low latency and excellent reliability. With intrusion protection, you can create multiple IPS sensors, each containing a complete configuration based on signatures. Then, you can apply any IPS sensor to any security policy.

This section describes how to configure the FortiGate Intrusion Protection settings.

If you enable virtual domains (VDOMs) on the FortiGate unit, intrusion protection is configured separately for each virtual domain.

The following topics are included:

- IPS concepts
- Enable IPS scanning
- Configure IPS options
- Enable IPS packet logging
- IPS examples

IPS concepts

The FortiGate intrusion protection system protects your network from outside attacks. Your FortiGate unit has two techniques to deal with these attacks: anomaly- and signature-based defense.

Anomaly-based defense

Anomaly-based defense is used when network traffic itself is used as a weapon. A host can be flooded with far more traffic than it can handle, making the host inaccessible. The most common example is the denial of service (DoS) attack, in which an attacker directs a large number of computers to attempt normal access of the target system. If enough access attempts are made, the target is overwhelmed and unable to service genuine users. The attacker does not gain access to the target system, but it is not accessible to anyone else.

The FortiGate DoS feature will block traffic above a certain threshold from the attacker and allow connections from other legitimate users. The DoS policy configuration information can be found in the Firewall Handbook.

Signature-based defense

Signature-based defense is used against known attacks or vulnerability exploits. These often involve an attacker attempting to gain access to your network. The attacker must communicate with the host in an attempt to gain access and this communication will include particular commands or sequences of commands and variables. The IPS signatures include these command sequences, allowing the FortiGate unit to detect and stop the attack.

Signatures

IPS signatures are the basis of signature-based intrusion protection. Every attack can be reduced to a particular string of commands or a sequence of commands and variables.

Signatures include this information so your FortiGate unit knows what to look for in network traffic.

Signatures also include characteristics about the attack they describe. These characteristics include the network protocol in which the attack will appear, the vulnerable operating system, and the vulnerable application.

To view the complete list of signatures, go to *Security Profiles > Intrusion Protection > IPS Signatures*. This will include the predefined signatures and any custom signatures that you may have created.

Protocol decoders

Before examining network traffic for attacks, the IPS engine uses protocol decoders to identify each protocol appearing in the traffic. Attacks are protocol-specific, so your FortiGate unit conserves resources by looking for attacks only in the protocols used to transmit them. For example, the FortiGate unit will only examine HTTP traffic for the presence of a signature describing an HTTP attack.

IPS engine

Once the protocol decoders separate the network traffic by protocol, the IPS engine examines the network traffic for the attack signatures.

IPS sensors

The IPS engine does not examine network traffic for all signatures, however. You must first create an IPS sensor and specify which signatures are included. Add signatures to sensors individually using signature entries, or in groups using IPS filters.

To view the IPS sensors, go to *Security Profiles > Intrusion Protection > IPS Sensor*.

IPS filters

IPS sensors contain one or more IPS filters. A filter is a collection of signature attributes that you specify. The signatures that have all of the attributes specified in a filter are included in the IPS filter.

For example, if your FortiGate unit protects a Linux server running the Apache web server software, you could create a new filter to protect it. By setting *OS* to *Linux*, and *Application* to *Apache,* the filter will include only the signatures that apply to both Linux and Apache. If you wanted to scan for all the Linux signatures and all the Apache signatures, you would create two filters, one for each.

To view the filters in an IPS sensor, go to *Security Profiles > Intrusion Protection > IPS Sensor*, select the IPS sensor containing the filters you want to view, and choose *Edit*.

Custom/predefined signature entries

Signature entries allow you to add an individual custom or predefined IPS signature. If you need only one signature, adding a signature entry to an IPS sensor is the easiest way. Signature entries are also the only way to include custom signatures in an IPS sensor.

Another use for signature entries are to change the settings of individual signatures that are already included in a filter within the same IPS sensor. Add a signature entry with the required settings above the filter, and the signature entry will take priority.

Policies

To use an IPS sensor, you must select it in a security policy or an interface policy. An IPS sensor that it not selected in a policy will have no effect on network traffic.

IPS is most often configured as part of a security policy. Unless stated otherwise, discussion of IPS sensor use will be in regards to firewall policies in this document.

Enable IPS scanning

Enabling IPS scanning involves two separate parts of the FortiGate unit:

- The security policy allows certain network traffic based on the sender, receiver, interface, traffic type, and time of day. Firewall policies can also be used to deny traffic, but those policies do not apply to IPS scanning.
- The IPS sensor contains filters, signature entries, or both. These specify which signatures are included in the IPS sensor.

When IPS is enabled, an IPS sensor is selected in a security policy, and all network traffic matching the policy will be checked for the signatures in the IPS sensor.

General configuration steps

For best results in configuring IPS scanning, follow the procedures in the order given. Also, note that if you perform any additional actions between procedures, your configuration may have different results.

1. Create an IPS sensor.
2. Add filters and/or predefined signatures and custom signatures to the sensor. The filters and signatures specify which signatures the IPS engine will look for in the network traffic.
3. Select a security policy or create a new one.
4. In the security policy, turn on *IPS*, and choose the IPS sensor from the list.

All the network traffic controlled by this security policy will be processed according to the settings in the policy. These settings include the IPS sensor you specify in the policy.

Creating an IPS sensor

You need to create an IPS sensor and save it before configuring it with filters and entries.

To create a new IPS sensor

1. Go to *Security Profiles > Intrusion Protection > IPS Sensors*.
2. Select the *Create New* icon in the top of the Edit IPS Sensor window.
3. Enter the name of the new IPS sensor.
4. Optionally, you may also enter a comment. The comment will appear in the IPS sensor list and serves to remind you of the details of the sensor.
5. Select *OK*.

The IPS sensor is created and the sensor configuration window appears. A newly created sensor is empty and contains no filters or signatures. You need to add one or more filters or signatures before the sensor can take effect.

Creating an IPS filter

While individual signatures can be added to a sensor, a filter allows you to add multiple signatures to a sensor by specifying the characteristics of the signatures to be added.

To create a new IPS filter

1. Go to *Security Profiles > Intrusion Protection > IPS Sensors*.

2. Select the IPS sensor to which you want to add the filter using the drop-down list in the top row of the Edit IPS Sensor window.

3. Select the *Create New* icon

4. For *Sensor Type* chose *Filter Based*.

5. Configure the filter that you require. Signatures matching all of the characteristics you specify in the filter will be included in the filter. Select Specify and choose the filter option that have the appropriate parameters.

Basic

Severity

Refers to the level of threat possed by the attack.

The options include:

- critical
- high
- medium
- low
- info

Target

Refers to the type of device targeted by the attack.

The options include:

- client
- server

OS

Refers to the Operating System affected by the attack.

The options include:

BSD	Linux	MacOS
Other	Solaris	Windows

Advanced

Application

Refers to the vendor or or type of application affected by the attack.

The options include:.

Adobe	Apache	Apple
CGI_app	Cisco	HP
IBM	IE	IIS
Mozilla	MS_Office	Novel
Oracle	PHP_app	Sun

This list can be expanded to include more options by selecting the [show more...] link. The additional options include:

ASP_app	CA	DB2
IM	Ipswitch	MailEnable
MediaPlayer	MS_Exchange	MSSQL
MySQL	Netscape	P2P
PostgreSQL	Real	Samba
SAP	SCADA	Sendmail
Veritas	Winamp	Other

Protocol

Refers to the protocol that is the vector for the attack.

The options include:

DNS	FTP	HTTP
ICMP	IMAP	LDAP
POP3	SCCP	SIP
SMTP	SNMP	SSH
SSL	TCP	UDP

This list can be expanded to include more options by selecting the [show more...] link. The additional options include:

BO	DCERPC	DHCP
DNP3	H323	IM
MSSQL	NBSS	NNTP
P2P	RADIUS	RDT
RPC	TRCP	RTP
RTSP	TELNET	TFN
Other		

6. Choose an action for when a signature is triggered.

Action	Description
Signature Default	All predefined signatures have an *Action* attribute that is set to Pass or Drop. This means that if a signature included in the filter has an *Action* setting of Pass, traffic matching the signature will be detected and then allowed to continue to its destination. Select *Accept signature defaults* use the default action for each included signature. Note: to see what the default for a signature is, go to the *IPS Signatures* page and enable the column *Action*, then find the row with the signature name in it.
Monitor All	Select *Monitor all* to pass all traffic matching the signatures included in the filter, regardless of their default *Action* setting.
Block All	Select Block all to drop traffic matching any the signatures included in the filter.
Reset	Select Reset to reset the session whenever the signature is triggered. In the CLI this action is referred to as Reject.
Quarantine	Has 2 fields the need to be configured: 1. Method: • Attacker's IP Address - Traffic from the Attacker's IP address is refused until the expiration time from the trigger is reached. • Attacker and Victim Address - All traffic from the Attacker's address to the Victim's address will be blocked. • Attack's incoming interface - the interface that experienced the attack will refuse further traffic. 2. Expires (time frame that the quarantine will be in effect): • 5 Minute(s) • 30 Minutes(s) • 1 Hour(s) • 1 Day(s) • Week(s) • Month(s) • Year(s)
Packet Logging	Select to enable packet logging for the filter. When you enable packet logging on a filter, the unit saves a copy of the packets that match any signatures included in the filter. The packets can be analyzed later. For more information about packet filtering, see "Monitoring Security Profiles activity" on page 169

7   Select *OK*.

The filter is created and added to the filter list.

Updating predefined IPS signatures

The FortiGuard Service periodically updates the pre-defined signatures and adds new signatures to counter emerging threats as they appear.

Because the signatures included in filters are defined by specifying signature attributes, new signatures matching existing filter specifications will automatically be included in those filters. For example, if you have a filter that includes all signatures for the Windows operating system, your filter will automatically incorporate new Windows signatures as they are added.

Viewing and searching predefined IPS signatures

Go to *Security Profiles > Intrusion Protection > IPS Signatures* to view the list of existing IPS signatures. You may find signatures by paging manually through the list, apply filters, or by using the search field.

Searching manually

Signatures are displayed in a paged list, with 50 signatures per page. The bottom of the screen shows the current page and the total number of pages. You can enter a page number and press enter, to skip directly to that page. Previous Page and Next Page buttons move you through the list, one page at a time. The First Page and Last Page button take you to the beginning or end of the list.

Applying filters

You can enter criteria for one of more columns, and only the signatures matching all the conditions you specify will be listed.

To apply filters

1. Go to *Security Profiles > Intrusion Protection > IPS Signatures.*

2. Select column by which to filter.

3. Select the funnel/filter icon and enter the value or values to filter by.

4. Use additional columns as needed to refine search.

The available options vary by column. For example, Enable allows you to choose between two options, while OS has multiple options, and you may select multiple items together. Filtering by name allows you to enter a text string and all signature names containing the string will be displayed.

IPS processing in an HA cluster

IPS processing in an HA cluster is no different than with a single FortiGate unit, from the point of view of the network user. The difference appears when a secondary unit takes over from the primary, and what happens depends on the HA mode.

Active-passive

In an active-passive HA cluster, the primary unit processes all traffic just as it would in a stand-alone configuration. Should the primary unit fail, a secondary unit will assume the role of the primary unit and begin to process network traffic. By default, the state of active communication sessions are not shared with secondary units and will not survive the fail-over condition. Once the sessions are reestablished however, traffic processing will continue as normal.

If your network requires that active sessions are taken over by the new primary unit, select *Enable Session Pick-up* in your HA configuration. Because session information must be sent to all subordinate units on a regular basis, session pick-up is a resource-intensive feature and is not enabled by default.

Active-active

The fail-over process in an active-active cluster is similar to an active-passive cluster. When the primary unit fails, a secondary unit takes over and traffic processing continues. The load-balancing schedule used to distribute sessions to the cluster members is used by the new primary unit to redistribute sessions among the remaining subordinate units. If session pick-up is not enabled, the sessions active on the failed primary are lost, and the sessions redistributed among the secondary units may also be lost. If session pick-up is enabled, all sessions are handled according to their last-known state.

Configure IPS options

There are a number of CLI commands that influence how IPS functions.

Hardware Acceleration

In order to provide control over the hardware's processing of IPS there are commands to configure and control the hardware accelleration of IPS. There are two settings that can be chosen, one for the network processor and one for the content processor.

Network processor acceleration can be disabled or set to enable basic acceleration.

Content processor acceleration can be disabled or set to either basic or advanced acceleration.

These Settings are only found in the CLI:

```
config ips global
   set np-accel-mode {none | basic}
   set cp-accel-mode {none | basic | advanced}
end
```

Extended IPS Database.

Some models have access to an extended IPS Database. The extended database may affect the performance of the FortiGate unit so depending on the model of the FortiGate unit the extended database package may not be enabled by default. For example, the D-series Desktop model have this option disabled by default.

This feature can only be enbled through the CLI.

```
config ips global
   set database extended
end
```

Configuring the IPS engine algorithm

The IPS engine is able to search for signature matches in two ways. One method is faster but uses more memory, the other uses less memory but is slower. Use the `algorithm` CLI command to select one method:

```
config ips global
    set algorithm {super | high | low | engine-pick}
end
```

Specify `high` to use the faster more memory intensive method or `low` for the slower memory efficient method. The setting `super` improves the performance for FortiGate units with more than 4GB of memory. The default setting is `engine-pick`, which allows the IPS engine to choose the best method on the fly.

Configuring the IPS engine-count

FortiGate units with multiple processors can run more than one IPS engine concurrently. The `engine-count` CLI command allows you to specify how many IPS engines are used at the same time:

```
config ips global
    set engine-count <int>
end
```

The recommended and default setting is `0`, which allows the FortiGate unit to determine the optimum number of IPS engines.

Configuring fail-open

If the IPS engine fails for any reason, it will fail open by default. This applies for inspection of all the protocols inspected by FortiOS IPS protocol decoders, including but not limited to HTTP, HTTPS, FTP, SMTP, POP3, IMAP, etc. This means that traffic continues to flow without IPS scanning. If IPS protection is more important to your network than the uninterrupted flow if network traffic, you can disable this behavior using the `fail-open` CLI command:

```
config ips global
    set fail-open {enable | disable}
end
```
The default setting is enable.

Configuring the session count accuracy

The IPS engine can keep track of the number of open session in two ways. An accurate count uses more resources than a less accurate heuristic count.

```
config ips global
    set session-limit-mode {accurate | heuristic}
end
```
The default is heuristic.

Configuring the IPS buffer size

Set the size of the IPS buffer.

```
config ips global
   set socket-size <int>
end
```

The acceptable range is from 1 to 64 megabytes. The default size varies by model.

Configuring protocol decoders

The FortiGate Intrusion Protection system uses protocol decoders to identify the abnormal traffic patterns that do not meet the protocol requirements and standards. For example, the HTTP decoder monitors traffic to identify any HTTP packets that do not meet the HTTP protocol standards.

To change the ports a decoder examines, you must use the CLI. In this example, the ports examined by the DNS decoder are changed from the default 53 to 100, 200, and 300.

```
config ips decoder dns_decoder
   set port_list "100,200,300"
end
```

You cannot assign specific ports to decoders that are set to *auto* by default. These decoders can detect their traffic on any port. Specifying individual ports is not necessary.

Configuring security processing modules

FortiGate Security Processing Modules, such as the CE4, XE2, and FE8, can increase overall system performance by accelerating some security and networking processing on the interfaces they provide. They also allow the FortiGate unit to offload the processing to the security module, thereby freeing up its own processor for other tasks. The security module performs its own IPS and firewall processing, but you can configure it to favor IPS in hostile high-traffic environments.

If you have a security processing module, use the following CLI commands to configure it to devote more resources to IPS than firewall. This example shows the CLI commands required to configure a security module in slot 1 for increased IPS performance.

```
config system amc-slot
   edit sw1
      set optimization-mode fw-ips
      set ips-weight balanced
      set ips-p2p disable
      set ips-fail-open enable
      set fp-disable none
      set ipsec-inb-optimization enable
      set syn-proxy-client-timer 3
      set syn-proxy-server-timer 3
   end
```

In addition to offloading IPS processing, security processing modules provide a hardware accelerated SYN proxy to defend against SYN flood denial of service attacks. When using a security module, configure your DoS anomaly check for `tcp_syn_flood` with the *Proxy* action. The *Proxy* action activates the hardware accelerated SYN proxy.

IPS signature rate count threshold

The IPS signature threshold can allow configuring a signature so that it will not be triggered until a rate count threshold is met. This provides a more controlled recording of attack activity. For example, if multiple login attempts produce a failed result over a short period of time then an alert would be sent and perhaps traffic blocked. This would be a more rational response than sending an alert every time a login failed.

The syntax for this configuration is as follows:

```
config ips sensor
  edit default
    config entries
        edit <Filter ID number>
        set rule <*id>
            set rate-count <integer between 1 - 65535>
            set rate-duration <integer between 1 - 65535>
```

The value of the rate-duration is an integer for the time in seconds.

```
        set rate-mode <continuous | periodical>
```

The rate-mode refers to how the count threshold is met.

If the setting is "`continuous`", and the action is set to block, as soon as the `rate-count` is reached the action is engaged. For example, if the count is 10, as soon as the signature is triggered 10 times the traffic would be blocked.

If the setting is "periodical", the FortiGate allows up to the value of the rate-count incidents where the signature is triggered during the rate-duration. For example, if the rate count is 100 and the duration is 60, the signature would need to be triggered 100 times in 60 seconds for the action to be engaged.

```
        set rate-track <dest-ip | dhcp-client-mac | dns-domain |
            none | src-ip>
```

This setting allow the tracking of one of the protocol fields within the packet.

```
      end
    end
```

Enable IPS packet logging

Packet logging saves the network packets containing the traffic matching an IPS signature to the attack log. The FortiGate unit will save the logged packets to wherever the logs are configured to be stored, whether memory, internal hard drive, a FortiAnalyzer unit, or the FortiGuard Analysis and Management Service.

You can enable packet logging in the filters. Use caution in enabling packet logging in a filter. Filters configured with few restrictions can contain thousands of signatures, potentially resulting in a flood of saved packets. This would take up a great deal of space, require time to sort through, and consume considerable system resources to process. Packet logging is designed as a focused diagnostic tool and is best used with a narrow scope.

Although logging to multiple FortiAnalyzer units is supported, packet logs are not sent to the secondary and tertiary FortiAnalyzer units. Only the primary unit receives packet logs.

To enable packet logging for a filter

1. Create a filter in an IPS sensor. For more information, see "Creating an IPS filter" on page 58.

2. Before saving the filter, select *Enable All* for *Packet Logging*.

3. Select the IPS sensor in the security policy that allows the network traffic the FortiGate unit will examine for the signature.

For information on viewing and saving logged packets, see "Monitoring Security Profiles activity" on page 169.

IPS examples

Configuring basic IPS protection

Small offices, whether they are small companies, home offices, or satellite offices, often have very simple needs. This example details how to enable IPS protection on a FortiGate unit located in a satellite office. The satellite office contains only Windows clients.

Creating an IPS sensor

Most IPS settings are configured in an IPS sensor. IPS sensors are selected in firewall policies. This way, you can create multiple IPS sensors, and tailor them to the traffic controlled by the security policy in which they are selected. In this example, you will create one IPS sensor.

To create an IPS sensor— web-based manager

1. Go to *Security Profiles > Intrusion Protection > IPS Sensors*.

2. Select the *Create New* icon in the top of the Edit IPS Sensor window.

3. In the *Name* field, enter `basic_ips`.

4. In the *Comments* field, enter `IPS protection for Windows clients.`

5. Select *OK*.

6. Select the *Create New* drop-down to add a new component to the sensor and for the *Sensor Type* choose *Filter* Based.

7. In the Filter Options choose the following:

   a. For *Severity*: select all of the options

   b. For *Target*: select *Client* only.

   c. For *OS*: select *Windows* only.

8. For the *Action* leave as the default.

9. Select *OK* to save the filter.

10. Select *OK* to save the IPS sensor.

To create an IPS sensor — CLI

```
config ips sensor
  edit basic_ips
    set comment "IPS protection for Windows clients"
      config entries
        edit 1
          set location client
          set os windows
    end
  end
```

Selecting the IPS sensor in a security policy

An IPS sensor directs the FortiGate unit to scan network traffic only when it is selected in a security policy. When an IPS sensor is selected in a security policy, its settings are applied to all the traffic the security policy handles.

To select the IPS sensor in a security policy — web-based manager

1. Go to *Policy > Policy > Policy*.
2. Select a policy.
3. Select the *Edit* icon.
4. Enable the *IPS* option.
5. Select the `basic_ips` profile from the list.
6. Select *OK* to save the security policy.

To select the IPS sensor in a security policy — CLI

```
config firewall policy
  edit 1
    set utm-status enable
    set ips-sensor basic_ips
  end
```

All traffic handled by the security policy you modified will be scanned for attacks against Windows clients. A small office may have only one security policy configured. If you have multiple policies, consider enabling IPS scanning for all of them.

Using IPS to protect your web server

Many companies have web servers and they must be protected from attack. Since web servers must be accessible, protection is not as simple as blocking access. IPS is one tool your FortiGate unit has to allow you to protect your network.

In this example, we will configure IPS to protect a web server. As shown in , a FortiGate unit protects a web server and an internal network. The internal network will have its own policies and configuration but we will concentrate on the web server in this example.

Figure 10:A simple network configuration
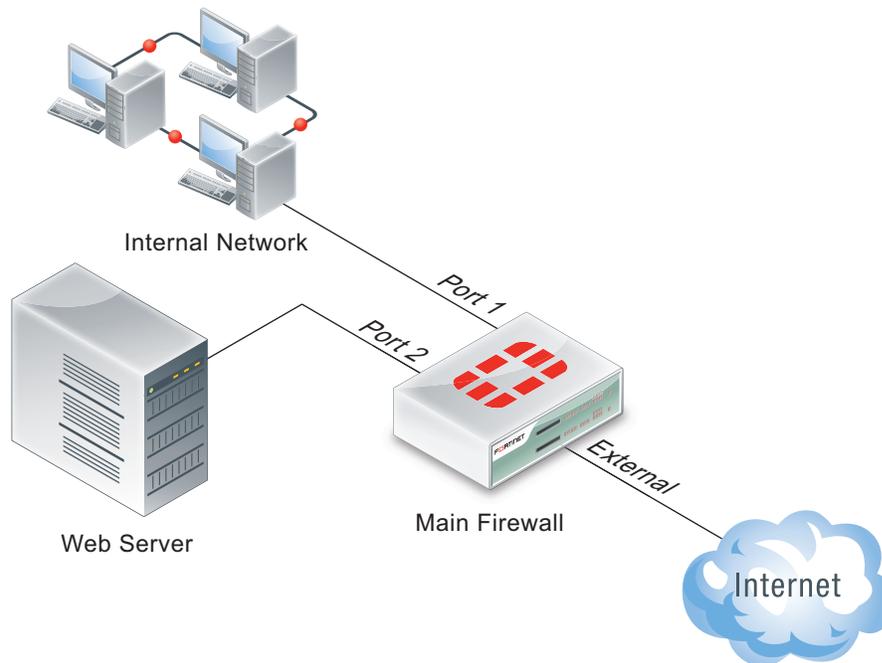
The FortiGate unit is configured with:

- a virtual IP to give the web server a unique address accessible from the Internet.
- a security policy to allow access to the web server from the Internet using the virtual IP.

To protect the web server using intrusion protection, you need to create an IPS sensor, populate it with filters, then enable IPS scanning in the security policy.

To create an IPS sensor

1. Go to *Security Profiles > Intrusion Protection > IPS Sensors* and select *Create New*.
2. Enter `web_server` as the name of the new IPS sensor.
3. Select *OK*.

The new IPS sensor is created but it has no filters, and therefore no signatures are included.

The web server operating system is Linux, so you need to create a filter for all Linux server signatures.

To create the Linux server filter

1. Go to *Security Profiles > Intrusion Protection > IPS Sensors* and select the `web_server` IPS sensor and select the *Edit* icon.
2. Select *Add Filter*.
3. Enter `Linux Server` as the name of the new filter.
4. For *Target*, select *Specify* and choose *server*.
5. In the Filter Options choose the following:
   a. For *Severity*: select all of the options
   b. For *Target*: select *server* only.
   c. For *OS*: select *Linux* only.
6. Select *OK*.

The filter is saved and the IPS sensor page reappears. In the filter list, find the *Linux Server* filter and look at the value in the *Count* column. This shows how many signatures match the current filter settings. You can select the *View Rules* icon to see a listing of the included signatures.

To edit the security policy

1. Go to *Policy > Policy > Policy*, select security policy that allows access to the web server, and select the *Edit* icon.
2. Enable IPS option and choose the `web_server` IPS sensor from the list.
3. Select *OK*.

Since IPS is enabled and the `web_server` IPS sensor is specified in the security policy controlling the web server traffic, the IPS sensor examines the web server traffic for matches to the signatures it contains.

Create and test a packet logging IPS sensor

In this example, you create a new IPS sensor and include a filter that detects the EICAR test file and saves a packet log when it is found. This is an ideal first experience with packet logging because the EICAR test file can cause no harm, and it is freely available for testing purposes.

Create an IPS senor

1. Go to *Security Profiles > Intrusion Protection > IPS Sensors*.
2. Select *Create New*.
3. Name the new IPS sensor `EICAR_test`.
4. Select *OK*.

Create an entry

1. Select the Create New drop down menu and for *Sensor Type* choose *Specify Signatures*.
2. Rather than search through the signature list, use the name filter by selecting the search icon over the header of the *Signature* column.
3. Enter `EICAR` in the Search field.
4. Highlight the `Eicar.Virus.Test.File` signature by clicking on it.
5. Select *Block All* as the *Action*.
6. Select *Enable*, *Packet Logging.*
7. Select OK to save the IPS sensor.

You are returned to the IPS sensor list. The `EICAR test` sensor appears in the list.

Add the IPS sensor to the security policy allowing Internet access

1. Go to *Policy > Policy > Policy*.
2. Select the security policy that allows you to access the Internet.
3. Select the *Edit* icon.
4. Enable *Log Allowed Traffic*.
5. Enable the *IPS* option.
6. Choose `EICAR test` from the available IPS sensors.
7. Select *OK*.

With the IPS sensor configured and selected in the security policy, the FortiGate unit blocks any attempt to download the EICAR test file.

Test the IPS sensor

1. Using your web browser, go to http://www.eicar.org/anti_virus_test_file.htm.

2. Scroll to the bottom of the page and select *eicar.com* from the row labeled as using the standard HTTP protocol.

3. The browser attempts to download the requested file and,

- If the file is successfully downloaded, the custom signature configuration failed at some point. Check the custom signature, the IPS sensor, and the firewall profile.

- If the download is blocked with a high security alert message explaining that you're not permitted to download the file, the EICAR test file was blocked by the FortiGate unit antivirus scanner before the IPS sensor could examine it. Disable antivirus scanning and try to download the EICAR test file again.

- If no file is downloaded and the browser eventually times out, the custom signature successfully detected the EICAR test file and blocked the download.

Viewing the packet log

1. Go to *Log&Report > Log & Archive Access > Security Log*.

2. Locate the log entry that recorded the blocking of the EICAR test file block. The Message field data will be `tools: EICAR.AV.Test.File.Download`.

3. Select the *View Packet Log* icon in the *Packet Log* column.

4. The packet log viewer is displayed.

Configuring a Fortinet Security Processing module

The Example Corporation has a web site that is the target of SYN floods. While they investigate the source of the attacks, it's very important that the web site remain accessible. To enhance the ability of the company's FortiGate-620B to deal with SYN floods, the administrator will install an ASM-CE4 Fortinet Security Processing module and have all external access to the web server come though it.

The security processing modules not only accelerate and offload network traffic from the FortiGate unit's processor, but they also accelerate and offload security and content scanning. The ability of the security module to accelerate IPS scanning and DoS protection greatly enhances the defense capabilities of the FortiGate-620B.

Assumptions

As shown in other examples and network diagrams throughout this document, the Example Corporation has a pair of FortiGate-620B units in an HA cluster. To simplify this example, the cluster is replaced with a single FortiGate-620B.

An ASM-CE4 is installed in the FortiGate-620B.

The network is configured as shown in Figure 11.

Network configuration

The Example Corporation network needs minimal changes to incorporate the ASM-CE4. Interface amc-sw1/1 of the ASM-CE4 is connected to the Internet and interface amc-sw1/1 is connected to the web server.

Since the main office network is connected to port2 and the Internet is connected to port1, a switch is installed to allow both port1 and amc-sw1/1 to be connected to the Internet.

Figure 11: The FortiGate-620B network configuration
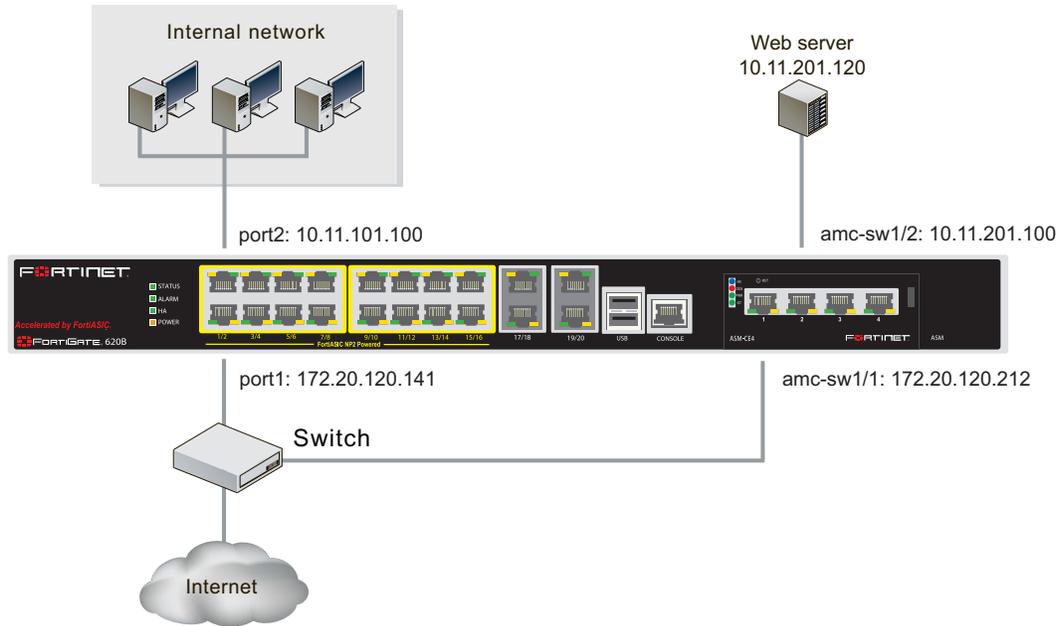
The switch used to connect port1 and amc-sw1/1 to the Internet must be able to handle any SYN flood, all of the legitimate traffic to the web site, and all of the traffic to and from the Example Corporation internal network. If the switch can not handle the bandwidth, or if the connection to the service provider can not provide the required bandwidth, traffic will be lost.

Security module configuration

The Fortinet security modules come configured to give equal priority to content inspection and firewall processing. The Example Corporation is using a ASM-CE4 module to defend its web server against SYN flood attacks so firewall processing is a secondary consideration.

Use these CLI commands to configure the security module in ASM slot 1 to devote more resources to content processing, including DoS and IPS, than to firewall processing.

```
config system amc-slot
   edit sw1
      set optimization-mode fw-ips
      set ips-weight balanced
      set ips-p2p disable
      set ips-fail-open enable
      set fp-disable none
      set ipsec-inb-optimization enable
      set syn-proxy-client-timer 3
      set syn-proxy-server-timer 3
   end
```

These settings do not disable firewall processing. Rather, when the security module nears its processing capacity, it will chose to service content inspection over firewall processing.

IPS Sensor

You can group signatures into IPS sensors for easy selection when applying to firewall policies. You can define signatures for specific types of traffic in separate IPS sensors, and then select those sensors in profiles designed to handle that type of traffic. For example, you can specify all

of the web-server related signatures in an IPS sensor, and that sensor can then be applied to a firewall policy that controls all of the traffic to and from a web server protected by the unit.

The FortiGuard Service periodically updates the pre-defined signatures, with signatures added to counter new threats. Since the signatures included in filters are defined by specifying signature attributes, new signatures matching existing filter specifications will automatically be included in those filters. For example, if you have a filter that includes all signatures for the Windows operating system, your filter will automatically incorporate new Windows signatures as they are added.

Each IPS sensor consists of two parts: filters and overrides. Overrides are always checked before filters.

Each filter consists of a number of signatures attributes. All of the signatures with those attributes, and only those attributes, are checked against traffic when the filter is run. If multiple filters are defined in an IPS Sensor, they are checked against the traffic one at a time, from top to bottom. If a match is found, the unit takes the appropriate action and stops further checking.

A signature override can modify the behavior of a signature specified in a filter. A signature override can also add a signature not specified in the sensor's filters. Custom signatures are included in an IPS sensor using overrides.

The signatures in the overrides are first compared to network traffic. If the IPS sensor does not find any matches, it then compares the signatures in each filter to network traffic, one filter at a time, from top to bottom. If no signature matches are found, the IPS sensor allows the network traffic.

The signatures included in the filter are only those matching every attribute specified. When created, a new filter has every attribute set to *all* which causes every signature to be included in the filter. If the severity is changed to high, and the target is changed to server, the filter includes only signatures checking for high priority attacks targeted at servers.

Custom Application & IPS Signatures

Creating a custom IPS signature

The FortiGate predefined signatures cover common attacks. If you use an unusual or specialized application or an uncommon platform, add custom signatures based on the security alerts released by the application and platform vendors.

You can add or edit custom signatures using the web-based manager or the CLI.

To create a custom signature

1. Go to *Security Profiles > Intrusion Protection > IPS Signatues*.
2. Select *Create New* to add a new custom signature.
3. Enter a *Name* for the custom signature.
4. Enter the *Signature*. For information about completing this field, see "Custom signature syntax and keywords".
5. Select *OK*.

Custom signature syntax and keywords

All custom signatures follow a particular syntax. Each begins with a header and is followed by one or more keywords. The syntax and keywords are detailed in the next two topics.

Custom signature syntax

A custom signature definition is limited to a maximum length of 512 characters. A definition can be a single line or span multiple lines connected by a backslash (\\) at the end of each line.

A custom signature definition begins with a header, followed by a set of keyword/value pairs enclosed by parenthesis [( )]. The keyword and value pairs are separated by a semi colon (;) and consist of a keyword and a value separated by a space. The basic format of a definition is HEADER (KEYWORD VALUE;)

You can use as many keyword/value pairs as required within the 512 character limit. To configure a custom signature, go to *Security Profiles > Intrusion Protection > IPS Signatues,* select *Create New* and enter the data directly into the *Signature* field, following the guidance in the next topics.

Table 1 shows the valid characters and basic structure. For details about each keyword and its associated values, see "Custom signature keywords" on page 76.

Table 1: Valid syntax for custom signature fields

Field	Valid Characters	Usage
HEADER	F-SBID	The header for an attack definition signature. Each custom signature must begin with this header.
KEYWORD	Each keyword must start with a pair of dashes (--), and consist of a string of 1 to 19 characters. Normally, keywords are an English word or English words connected by an underscore (_). Keywords are case insensitive.	The keyword is used to identify a parameter. See "Custom signature keywords" on page 76 for tables of supported keywords.
VALUE	Double quotes (") must be used around the value if it contains a space and/or a semicolon (;). If the value is NULL, the space between the KEYWORD and VALUE can be omitted. Values are case sensitive. Note: If double quotes are used for quoting the value, the double quotes are not considered as part of the value string.	The value is set specifically for a parameter identified by a keyword.

Custom signature keywords

Table 2: Information keywords

Keyword and value	Description
`--attack_id <id_int>;`	Use this optional value to identify the signature. It cannot be the same value as any other custom rules. If an attack ID is not specified, the FortiGate automatically assigns an attack ID to the signature. If you are using VDOMs, custom signatures appear only in the VDOM in which you create them. You can use the same attack ID for signatures in different VDOMs. An attack ID you assign must be between 1000 and 9999. Example: `--attack_id 1234;`
`--name <name_str>;`	Enter the name of the rule. A rule name must be unique. If you are using VDOMs, custom signatures appear only in the VDOM in which you create them. You can use the same rule name for signatures in different VDOMs. The name you assign must be a string greater than 0 and less than 64 characters in length. Example: `--name "Buffer_Overflow";`

Table 3: Session keywords

Keyword and value	Description
`--flow {from_client[,reversed]	from_server[,reversed]	bi_direction };`	Specify the traffic direction and state to be inspected. They can be used for all IP traffic. Example: --src_port 41523; --flow bi_direction; The signature checks traffic to and from port 41523. If you enable "quarantine attacker", the optional `reversed` keyword allows you to change the side of the connection to be quarantined when the signature is detected. For example, a custom signature written to detect a brute-force log in attack is triggered when "Login Failed" is detected `from_server` more than 10 times in 5 seconds. If the attacker is quarantined, it is the server that is quarantined in this instance. Adding `reversed` corrects this problem and quarantines the actual attacker. Previous FortiOS versions used `to_client` and `to_server` values. These are now deprecated, but still function for backwards compatibility.
`--service {HTTP	TELNET	FTP	DNS	SMTP	POP3	IMAP	SNMP	RADIUS	LDAP	MSSQL	RPC	SIP	H323	NBSS	DCERPC	SSH	SSL};`	Specify the protocol type to be inspected. This keyword allows you to specify the traffic type by protocol rather than by port. If the decoder has the capability to identify the protocol on any port, the signature can be used to detect the attack no matter what port the service is running on. Currently, HTTP, SIP, SSL, and SSH protocols can be identified on any port based on the content.

Table 4: UDP header keywords

`--dst_port`
`[!]{<port_int> |`
`:<port_int> |`
`<port_int>: |`
`<port_int>:<port_int>};`
	• `<port_int>` is a single port.
	• `:<port_int>` includes the specified port and all lower numbered ports.
	• `<port_int>:` includes the specified port and all higher numbered ports.
	• `<port_int>:<port_int>` includes the two specified ports and all ports in between.

`--src_port`
`[!]{<port_int> |`
`:<port_int> |`
`<port_int>: |`
`<port_int>:<port_int>};`
	• `<port_int>` is a single port.
	• `:<port_int>` includes the specified port and all lower numbered ports.
	• `<port_int>:` includes the specified port and all higher numbered ports.
	• `<port_int>:<port_int>` includes the two specified ports and all ports in between.

Table 5: ICMP keywords

Keyword and Value	Usage
`--icmp_code <code_int>;`	Specify the ICMP code to match.
`--icmp_id <id_int>;`	Check for the specified ICMP ID value.
`--icmp_seq <seq_int>;`	Check for the specified ICMP sequence value.
`--icmp_type <type_int>;`	Specify the ICMP type to match.

Table 6: Other keywords

Keyword and Value	Description
`--data_size {<size_int> | <<size_int> | ><size_int> | <port_int><><port_int>};`	Test the packet payload size. With `data_size` specified, packet reassembly is turned off automatically. So a signature with `data_size` and `only_stream` values set is wrong. • `<size_int>` is a particular packet size. • `<<size_int>` is a packet smaller than the specified size. • `><size_int>` is a packet larger than the specified size. • `<size_int><><size_int>` is a packet within the range between the specified sizes.
`--data_at <offset_int>[, relative];`	Verify that the payload has data at a specified offset, optionally looking for data relative to the end of the previous content match.
`--rate <matches_int>,<time_int>;`	Instead of generating log entries every time the signature is detected, use this keyword to generate a log entry only if the signature is detected a specified number of times within a specified time period. • `<matches_int>` is the number of times a signature must be detected. • `<time_int>` is the length of time in which the signature must be detected, in seconds. For example, if a custom signature detects a pattern, a log entry will be created every time the signature is detected. If `--rate 100,10;` is added to the signature, a log entry will be created if the signature is detected 100 times in the previous 10 seconds. Use this command with `--track` to further limit log entries to when the specified number of detections occur within a certain time period involving the same source or destination address rather than all addresses.
`--rpc_num <app_int>[, <ver_int> | *][, <proc_int> | *>];`	Check for RPC application, version, and procedure numbers in SUNRPC CALL requests. The `*` wildcard can be used for version and procedure numbers.

Table 6: Other keywords (continued)

Keyword and Value	Description
`--same_ip;`	Check that the source and the destination have the same IP addresses.
`--track {client | server}[,block_int];`	When used with `--rate`, this keyword narrows the custom signature rate totals to individual addresses. • `client` has the FortiGate unit maintain a separate count of signature matches for each source address. • `server` has the FortiGate unit maintain a separate count of signature matches for each destination address. • `block_int` has the FortiGate unit block connections for the specified number of seconds, from the client or to the server, depending on which is specified. For example, if `--rate 100,10` is added to the signature, a log entry will be created if the signature is detected 100 times in the previous 10 seconds. The FortiGate unit maintains a single total, regardless of source and destination address. If the same custom signature also includes `--track client;` matches are totalled separately for each source address. A log entry is added when the signature is detected 100 times in 10 seconds within traffic from the same source address. The `--track` keyword can also be used without `--rate`. If an integer is specified, the client or server will be blocked for the specified number of seconds every time the signature is detected.

Creating a custom signature to block access to example.com

In this first example, you will create a custom signature to block access to the example.com URL.

This example describes the use of the custom signature syntax to block access to a URL. To create the custom signature entry in the FortiGate unit web-based manager, see "Creating a custom IPS signature" on page 74.

1. Enter the custom signature basic format

All custom signatures have a header and at least one keyword/value pair. The header is always the same:

```
F-SBID( )
```

The keyword/value pairs appear within the parentheses and each pair is followed by a semicolon.

2. Choose a name for the custom signature

Every custom signature requires a name, so it is a good practice to assign a name before adding any other keywords.

Use the `--name` keyword to assign the custom signature a name. The name value follows the keyword after a space. Enclose the name value in double-quotes:

```
F-SBID( --name "Block.example.com"; )
```

The signature, as it appears here, will not do anything if you try to use it. It has a name, but does not look for any patterns in network traffic. You must specify a pattern that the FortiGate unit will search for.

3. Add a signature pattern

Use the `--pattern` keyword to specify what the FortiGate unit will search for:

```
F-SBID( --name "Block.example.com"; --pattern "example.com"; )
```

The signature will now detect the example.com URL appearing in network traffic. The custom signature should only detect the URL in HTTP traffic, however. Any other traffic with the URL should be allowed to pass. For example, an email message to or from example.com should not be stopped.

4. Specify the service

Use the `--service` keyword to limit the effect of the custom signature to only the HTTP protocol.

```
F-SBID( --name "Block.example.com"; --pattern "example.com";
    --service HTTP; )
```

The FortiGate unit will limit its search for the pattern to the HTTP protocol. Even though the HTTP protocol uses only TCP traffic, the FortiGate will search for HTTP protocol communication in TCP, UDP, and ICMP traffic. This is a waste of system resources that you can avoid by limiting the search further, as shown below.

5. Specify the traffic type.

Use the `--protocol tcp keyword to limit the effect of the custom signature to only TCP traffic. This will save system resources by not unnecessarily scanning UDP and ICMP traffic.`

```
F-SBID( --name "Block.example.com"; --pattern "example.com";
    --service HTTP; --protocol tcp; )
```

The FortiGate unit will limit its search for the pattern to TCP traffic and ignore UDP and ICMP network traffic.

6. Ignore case sensitivity

By default, patterns are case sensitive. If a user directed his or her browser to Example.com, the custom signature would not recognize the URL as a match.

Use the `--no_case` keyword to make the pattern matching case insensitive.

```
F-SBID( --name "Block.example.com"; --pattern "example.com";
    --service HTTP; --no_case; )
```

Unlike all of the other keywords in this example, the `--no_case` keyword has no value. Only the keyword is required.

7. Limit pattern scans to only traffic sent from the client

The `--flow` command can be used to further limit the network traffic being scanned to only that send by the client or by the server.

```
F-SBID( --name "Block.example.com"; --pattern "example.com";
    --service HTTP; --no_case; --flow from_client; )
```

Web servers do not contact clients until clients first open a communication session. Therefore, using the `--flow from_client` command will force the FortiGate to ignore all

traffic from the server. Since the majority of HTTP traffic flows from the server to the client, this will save considerable system resources and still maintain protection.

8. Specify the context

When the client browser tries to contact example.com, a DNS is first consulted to get the example.com server IP address. The IP address is then specified in the URL field of the HTTP communication. The domain name will still appear in the host field, so this custom signature will not function without the `--context host` keyword/value pair.

```
F-SBID( --name "Block.example.com"; --pattern "example.com";
    --service HTTP; --no_case; --flow from_client;
    --context host; )
```

Creating a custom signature to block the SMTP "vrfy" command

The SMTP "vrfy" command can be used to verify the existence of a single email address or to list all of the valid email accounts on an email server. A spammer could potentially use this command to obtain a list of all valid email users and direct spam to their inboxes.

In this example, you will create a custom signature to block the use of the vrfy command. Since the custom signature blocks the vrfy command from coming through the FortiGate unit, the administrator can still use the command on the internal network.

This example describes the use of the custom signature syntax to block the vrfy command. To create the custom signature entry in the FortiGate unit web-based manager, see "Creating a custom IPS signature" on page 74.

1. Enter the custom signature basic format

All custom signatures have a header and at least one keyword/value pair. The header is always the same:

```
F-SBID( )
```

The keyword/value pairs appear within the parentheses and each pair is followed by a semicolon.

2. Choose a name for the custom signature

Every custom signature requires a name, so it is a good practice to assign a name before you add any other keywords.

Use the `--name` keyword to assign the custom signature a name. The name value follows the keyword after a space. Enclose the name value in double-quotes:

```
F-SBID( --name "Block.SMTP.VRFY.CMD"; )
```

The signature, as it appears here, will not do anything if you try to use it. It has a name, but does not look for any patterns in network traffic. You must specify a pattern that the FortiGate unit will search for.

3. Add a signature pattern

Use the `--pattern` keyword to specify what the FortiGate unit will search for:

```
F-SBID( --name "Block.SMTP.VRFY.CMD"; --pattern "vrfy"; )
```

The signature will now detect the vrfy command appearing in network traffic. The custom signature should only detect the command in SMTP traffic, however. Any other traffic with the pattern should be allowed to pass. For example, an email message discussing the vrfy command should not be stopped.

4. Specify the service

Use the `--service` keyword to limit the effect of the custom signature to only the HTTP protocol.

```
F-SBID( --name "Block.SMTP.VRFY.CMD"; --pattern "vrfy";
    --service SMTP; )
```

The FortiGate unit will limit its search for the pattern to the SMTP protocol.

Even though the SMTP protocol uses only TCP traffic, the FortiGate will search for SMTP protocol communication in TCP, UDP, and ICMP traffic. This is a waste of system resources that you can avoid by limiting the search further, as shown below.

5. Specify the traffic type.

Use the `--protocol tcp keyword to limit the effect of the custom signature to only TCP traffic. This will save system resources by not unnecessarily scanning UDP and ICMP traffic.`

```
F-SBID( --name "Block.SMTP.VRFY.CMD"; --pattern "vrfy";
    --service SMTP; --protocol tcp; )
```

The FortiGate unit will limit its search for the pattern to TCP traffic and ignore the pattern in UDP and ICMP network traffic.

6. Ignore case sensitivity

By default, patterns are case sensitive. If a user directed his or her browser to Example.com, the custom signature would not recognize the URL as a match.

Use the `--no_case` keyword to make the pattern matching case insensitive.

```
F-SBID( --name "Block.SMTP.VRFY.CMD"; --pattern "vrfy";
    --service SMTP; --no_case; )
```

Unlike all of the other keywords in this example, the `--no_case` keyword has no value. Only the keyword is required.

7. Specify the context

The SMTP vrfy command will appear in the SMTP header. The `--context host` keyword/value pair allows you to limit the pattern search to only the header.

```
F-SBID( --name "Block.SMTP.VRFY.CMD"; --pattern "vrfy";
    --service SMTP; --no_case; --context header; )
```

Web filter

This section describes FortiGate web filtering for HTTP traffic. The three main parts of the web filtering function, the Web Content Filter, the URL Filter, and the FortiGuard Web Filtering Service interact with each other to provide maximum control over what the Internet user can view as well as protection to your network from many Internet content threats. Web Content Filter blocks web pages containing words or patterns that you specify. URL filtering uses URLs and URL patterns to block or exempt web pages from specific sources. FortiGuard Web Filtering provides many additional categories you can use to filter web traffic.

This section describes the Web Content Filter and URL Filter functions. For information on FortiGuard Web Filtering, see "FortiGuard Web Filter" on page 133

The following topics are included in this section:

- Web filter concepts
- Inspections Modes
- FortiGuard Web Filtering Service
- Overriding FortiGuard website categorization
- SafeSearch
- YouTube Education Filter
- Web Site Filter
- Web content filter
- Advanced web filter configurations
- Working with the Interface
- Web filtering example

Web filter concepts

Web filtering is a means of controlling the content that an Internet user is able to view. With the popularity of web applications, the need to monitor and control web access is becoming a key component of secure content management systems that employ antivirus, web filtering, and messaging security. Important reasons for controlling web content include:

- lost productivity because employees are accessing the web for non-business reasons
- network congestion — when valuable bandwidth is used for non-business purposes, legitimate business applications suffer
- loss or exposure of confidential information through chat sites, non-approved email systems, instant messaging, and peer-to-peer file sharing
- increased exposure to web-based threats as employees surf non-business-related web sites
- legal liability when employees access/download inappropriate and offensive material
- copyright infringement caused by employees downloading and/or distributing copyrighted material.

As the number and severity of threats increase on the World Wide Web, the risk potential increases within a company's network as well. Casual non-business related web surfing has caused many businesses countless hours of legal litigation as hostile environments have been created by employees who download and view offensive content. Web-based attacks and

threats are also becoming increasingly sophisticated. Threats and web-based applications that cause additional problems for corporations include:

- spyware/grayware
- phishing
- pharming
- instant messaging
- peer-to-peer file sharing
- streaming media
- blended network attacks.

Spyware, also known as grayware, is a type of computer program that attaches itself to a user's operating system. It does this without the user's consent or knowledge. It usually ends up on a computer because of something the user does such as clicking on a button in a pop-up window. Spyware can track the user's Internet usage, cause unwanted pop-up windows, and even direct the user to a host web site. For further information, visit the FortiGuard Center.

Some of the most common ways of grayware infection include:

- downloading shareware, freeware, or other forms of file-sharing services
- clicking on pop-up advertising
- visiting legitimate web sites infected with grayware.

Phishing is the term used to describe attacks that use web technology to trick users into revealing personal or financial information. Phishing attacks use web sites and email that claim to be from legitimate financial institutions to trick the viewer into believing that they are legitimate. Although phishing is initiated by spam email, getting the user to access the attacker's web site is always the next step.

Pharming is a next generation threat that is designed to identify and extract financial, and other key pieces of information for identity theft. Pharming is much more dangerous than phishing because it is designed to be completely hidden from the end user. Unlike phishing attacks that send out spam email requiring the user to click to a fraudulent URL, pharming attacks require no action from the user outside of their regular web surfing activities. Pharming attacks succeed by redirecting users from legitimate web sites to similar fraudulent web sites that have been created to look and feel like the authentic web site.

Instant messaging presents a number of problems. Instant messaging can be used to infect computers with spyware and viruses. Phishing attacks can be made using instant messaging. There is also a danger that employees may use instant messaging to release sensitive information to an outsider.

Peer-to-peer (P2P) networks are used for file sharing. Such files may contain viruses. Peer-to-peer applications take up valuable network resources and may lower employee productivity but also have legal implications with the downloading of copyrighted or sensitive company material.

Streaming media is a method of delivering multimedia, usually in the form of audio or video to Internet users. Viewing streaming media impacts legitimate business by using valuable bandwidth.

Blended network threats are rising and the sophistication of network threats is increasing with each new attack. Attackers learn from each previous successful attack and enhance and update attack code to become more dangerous and fast spreading. Blended attacks use a combination of methods to spread and cause damage. Using virus or network worm techniques combined with known system vulnerabilities, blended threats can quickly spread through email, web sites, and Trojan applications. Examples of blended threats include Nimda, Code Red, Slammer, and Blaster. Blended attacks can be designed to perform different types of attacks,

which include disrupting network services, destroying or stealing information, and installing stealthy backdoor applications to grant remote access.

Different ways of controlling access

The methods available for monitoring and controlling Internet access range from manual and educational methods to fully automated systems designed to scan, inspect, rate and control web activity.

Common web access control mechanisms include:

- establishing and implementing a well-written usage policy in the organization on proper Internet, email, and computer conduct
- installing monitoring tools that record and report on Internet usage
- implementing policy-based tools that capture, rate, and block URLs.

The final method is the focus of this topic. The following information shows how the filters interact and how to use them to your advantage.

Order of web filtering

The FortiGate unit applies web filters in a specific order:

1. URL filter
2. FortiGuard Web Filter
3. web content filter
4. web script filter
5. antivirus scanning.

If you have blocked a FortiGuard Web Filter category but want certain users to have access to URLs within that pattern, you can use the *Override* within the FortiGuard Web Filter. This will allow you to specify which users have access to which blocked URLs and how long they have that access. For example, if you want a user to be able to access www.example.com for one hour, you can use the override to set up the exemption. Any user listed in an override must fill out an online authentication form that is presented when they try to access a blocked URL before the FortiGate unit will grant access to it. For more information, see .

Inspections Modes

Proxy

Proxy-based inspection involves buffering the traffic and examining it as a whole before determining an action. The process of having the whole of the data to analyze allow this process to include more points of data to analyze than the flow-based or DNS methods.

The advantage of a proxy-based method is that the inspection can be more thorough than the other methods, resulting in fewer false positive or negative results in the analysis of the data.

Flow-based

The Flow-based inspection method examines the file as it passes through the FortiGate unit without any buffering. As each packet of the traffic arrives it is process and forwarded without waiting for the complete file or web page, etc.

The advantage of the flow-based method is that the user sees a faster response time for HTTP requests and there is less chance of a time-out error due to the server at the other end responding slowly.

The disadvantages of this method are that there is a higher probability of a false positive or negative in the analysis of the data and that a number of points of analysis that can be used in the proxy-based method are not available in the flow-based inspection method. There is also fewer actions available to choose from based on the categorization of the website by FortiGuard services.

DNS

The DNS inspection method uses the same categories as the FortiGuard Service. It is lightweight in terms of resource usage because it doesn't involve any proxy-based or flow-based inspection.

A DNS request is typically the first part of any new session to a new website. This inspection method takes advantage of that and places the results of the categorization of websites right on the FortiGuard DNS servers. When the FortiGate resolves a URL, in addition to the IP address of the website it also receives a domain rating.

In the same way that the flow-based inspection method had fewer filters and points of analysis than the proxy-based inspection method, DNS has fewer settings still. All of its inspection is based on the IP address, the domain name and the rating provided by the FortiGuard DNS server.

FortiGuard Web Filtering Service

FortiGuard Web Filter is a managed web filtering solution available by subscription from Fortinet. FortiGuard Web Filter enhances the web filtering features supplied with your FortiGate unit by sorting billions of web pages into a wide range of categories users can allow or block. The FortiGate unit accesses the nearest FortiGuard Web Filter Service Point to determine the category of a requested web page, and then applies the security policy configured for that user or interface.

FortiGuard Web Filter includes over 45 million individual ratings of web sites that apply to more than two billion pages. Pages are sorted and rated into several dozen categories administrators can allow or block. Categories may be added or updated as the Internet evolves. To make configuration simpler, you can also choose to allow or block entire groups of categories. Blocked pages are replaced with a message indicating that the page is not accessible according to the Internet usage policy.

FortiGuard Web Filter ratings are performed by a combination of proprietary methods including text analysis, exploitation of the web structure, and human raters. Users can notify the FortiGuard Web Filter Service Points if they feel a web page is not categorized correctly, so that the service can update the categories in a timely fashion.

Before you begin to use the FortiGuard Web Filter options you should verify that you have a valid subscription to the service for your FortiGate firewall.

FortiGuard Web Filter and your FortiGate unit

When FortiGuard Web Filter is enabled in a web filter profile, the setting is applied to all firewall policies that use this profile. When a request for a web page appears in traffic controlled by one of these firewall policies, the URL is sent to the nearest FortiGuard server. The URL category is returned. If the category is blocked, the FortiGate unit provides a replacement message in place of the requested page. If the category is not blocked, the page request is sent to the requested URL as normal.

Figure 12: Webfiltering flowchart
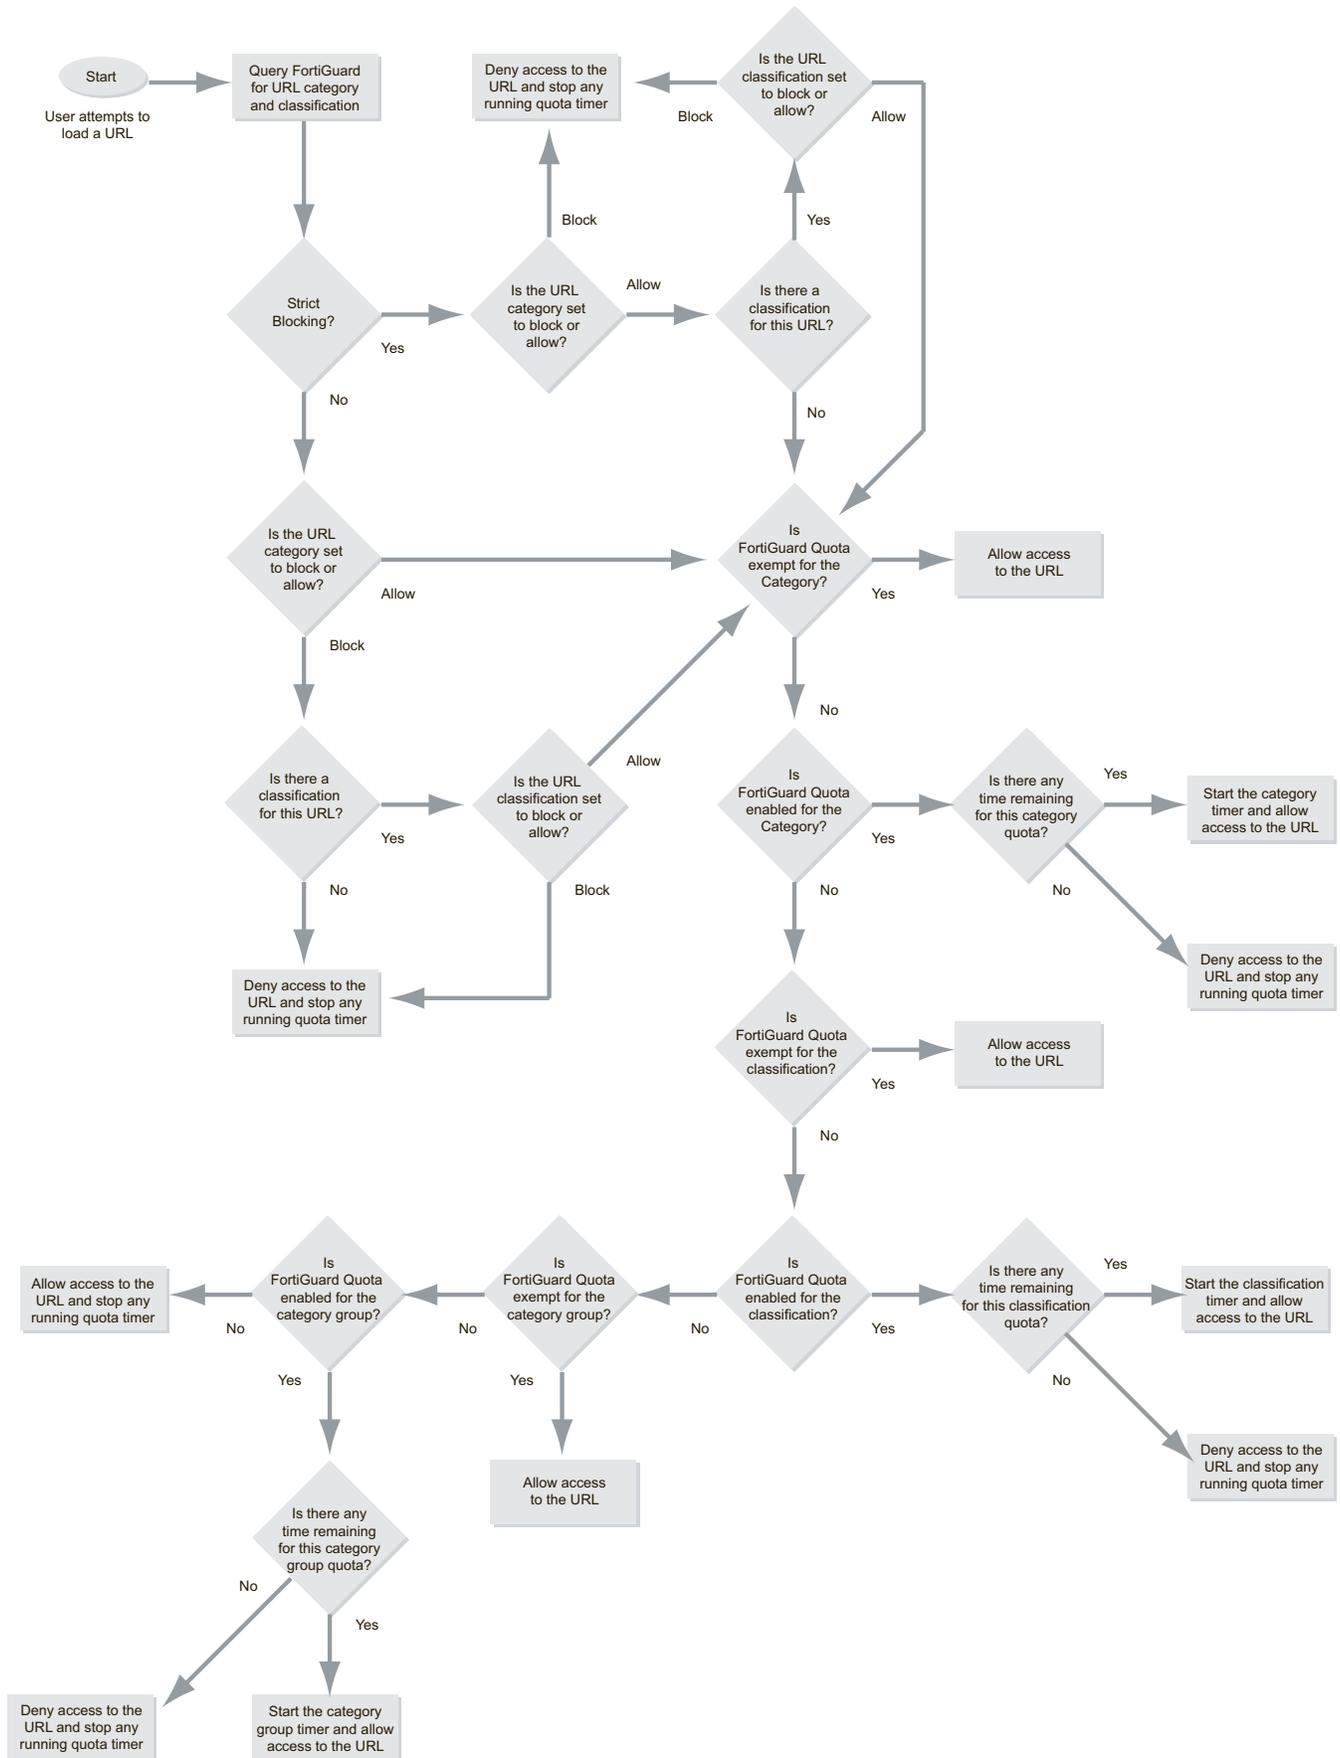
Enabling FortiGuard Web Filter

FortiGuard Web Filter is enabled and configured within web filter profiles by enabling FortiGuard Categories. The service is engaged by turning on the Web Filter profile and selecting a profile that has FortiGuard Categories enabled on one or more active policies being run by the firewall.

There is also a system wide setting for the enabling or disabling of FortiGuard Web Filter that is only in the CLI.

```
config system fortiguard
   set webfilter-force-off
```

The two options on this setting are enable or disable. The syntax of the settings name is "force-off" so in order to enable FortiGuard Webfilter you have to choose disable for the setting and enable if you want to turn it off.

General configuration steps

1. Go to Security Profiles > Web Filter > Profile.
2. Select the Edit icon of the web filter profile in which you want to enable FortiGuard Web Filter, or select Create New to add a new web filter profile.
3. Select an Inspection Mode.
4. If you are using FortiGuard Categories, enable the feature, select the categories and select the action to be performed.
5. The categories allow you to block or allow access to general or more specific web site categories. Configure access as required.
6. Save the filter and web filter profile.
7. To complete the configuration, you need to select the security policy controlling the network traffic you want to restrict. Then, in the security policy, enable Web Filter and select the appropriate web filter profile from the list.

Configuring FortiGuard Web Filter settings

FortiGuard Web Filter includes a number of settings that allow you to determine various aspects of the filtering behavior.

To configure FortiGuard Web Filter settings

1. Go to *Security Profiles > Web Filter > Profile*.
2. Select the web filter profile in which you want to enable FortiGuard Web Filter from the drop down list in the Edit Web Filter Profile window title bar, or select *Create New* to add a new web filter profile.
3. The category groups are listed in a table. You can expand each category group to view and configure every category within the groups. If you change the setting of a category group, all categories within the group inherit the change.
4. Select the category groups and categories to which you want to apply an action.

5. Select an action from the *Change Action for Selected Categories* drop-down list immediately below the category table. Five actions are available:

- *Allow* permits access to the sites within the category.
- *Monitor* permits and logs access to sites in the category. You may also enable user quotas when enabling the monitor action.
- *Warning* presents the user with a message, allowing them to continue if they choose.
- *Authenticate* requires a user authenticate with the FortiGate unit before being allowed access to the category or category group.
- *Block* prevents access to sites within the category. Users attempting to access a blocked site will receive a replacement message explaining that access to the site is blocked.

6. Select *OK*.

In older versions of FortiOS there was a character limitation for the URL of 2048 bytes or approximately 321 characters. If the URL you were trying to reach was longer the URL sent to FortiGuard would be truncated and the service would be unable to categorize the site. Starting in version 5 of the firmware the parsed URL has been increase to 4Kilobytes, effectively doubling the length of a URL capable of being categorized.

To configure the FortiGuard Web Filter categories

1. Go to *Security Profiles > Web Filter > Profiles*.
2. Select the web filter profile in which you want to enable FortiGuard Web Filter from the drop down list in the Edit Web Filter Profile window title bar, or select *Create New* to add a new web filter profile.
3. Select *Create New*.
4. Select a *Filter Type* of *Category*.
5. Select the required category groups. You may also expand the category groups to select individual categories.
6. Select the *Monitor* action.
7. Enable *Enforce Quota* to activate the quota for the selected categories and category groups.
8. Select *Hours*, *Minutes*, or *Seconds* and enter the number of hours, minutes, or seconds. This is the daily quota allowance for each user.
9. Select *OK*.
10. Select *Apply*.

Apply the web filter profile to an identity-based security policy. All the users subject to that policy are restricted by the quotas.

If you look at your logs carefully, you may notice that not every URL connection in the log shows a category. They are left blank. If you take one of those URL and enter it in the FortiGuard website designed to show the category for a URL it will successfully categorize it.

The reason for this is that to optimize speed throughput and reduce the load on the FortiGuard servers the FortiGate does not determine a category rating on scripts and css files.

Configuring FortiGuard Web Filter usage quotas

In addition to using category and classification blocks and overrides to limit user access to URLs, you can set a daily timed access quota by category, category group, or classification.

Quotas allow access for a specified length of time, calculated separately for each user. Quotas are reset every day at midnight.

Users must authenticate with the FortiGate unit. The quota is applied to each user individually so the FortiGate must be able to identify each user. One way to do this is to configure a security policy using the identity based policy feature. Apply the web filter profile in which you have configured FortiGuard Web Filter and FortiGuard Web Filter quotas to such a security policy.

The use of FortiGuard Web Filter quotas requires that users authenticate to gain web access. The quotas are ignored if applied to a security policy in which user authentication is not required.

When a user first attempts to access a URL, they're prompted to authenticate with the FortiGate unit. When they provide their user name and password, the FortiGate unit recognizes them, determines their quota allowances, and monitors their web use. The category and classification of each page they visit is checked and FortiGate unit adjusts the user's remaining available quota for the category or classification.

Editing the web filter profile resets the quota timers for all users.

1. Select the *Monitor* action.
2. Enable *Enforce Quota* to activate the quota for the selected categories and category groups.
3. Select *Hours*, *Minutes*, or *Seconds* and enter the number of hours, minutes, or seconds. This is the daily quota allowance for each user.
4. Select *OK*.
5. Select *Apply*.

Apply the web filter profile to an identity-based security policy. All the users subject to that policy are restricted by the quotas.

Quota hierarchy

You can apply quotas to categories and category groups. Only one quota per user can be active at any one time. The one used depends on how you configure the FortiGuard Web Filter.

When a user visits a URL, the FortiGate unit queries the FortiGuard servers for the category of the URL. From highest to lowest, the relative priority of the quotas are:

1. Category
2. Category group

Overriding FortiGuard website categorization

In most things there is an exception to the rule. When it comes to the rules about who is allowed to go to which websites in spite of the rules or in this case, policies, it seems that there are more exceptions than to most rules. There are numerous valid reasons and scenarios for exceptions so it follows that there needs to be a way to accommodate this exceptions.

The different methods of override

There are actually two different ways to override web filtering behavior based on FortiGuard categorization of a websites. The second method has 2 variations in implementation and each of the three has a different level of granularity.

1. Using Alternate Categories

Rating Override

This method manually assigns a specific website to a different Fortinet category or a locally created category.

2. Using Alternate Profiles

Administrative Override or Allow Blocked Override

In this method all of the traffic going through the FortiGate unit, using identity based policies and a Web Filtering profile has the option where configured users or IP addresses can use an alternative Web Filter profile when attempting to access blocked websites.

Using Alternate Categories

Rating Overrides

There are two approached to overriding the FortiGuard Web Filtering. The first is an identity based method that can be configured using a combination of identity based policies and specifically designed webfilter profiles. This has been addressed in the Firewall Handbook.

The second method is the system wide approach that locally (on the FortiGate Firewall) reassigns a URL to a different FortiGuard Category and even subcategory. This is where you can set assign a specific URL to the FortiGuard Category that you want to you can also set the URL to one of the Local Categories that you have created

The Rating Overrides option is available because different people will have different criteria for how they categorize websites. Even is the criteria is the same an organization may have reason to block the bulk of a category but need to be able to access specific URLs that are assigned to that category.

A hypothetical example could be that a website, example.com is categorized as being in the Sub-Category Pornography. The law offices of Barrister, Solicitor and Lawyer do not want their employees looking at pornography at work so they have used the FortiGuard Webfilter to block access to sites that have been assigned to the Category "Pornography". However, the owners of example.com are clients of the law office and they are aware that example.com is for artists that specialize in nudes and erotic images. In this case to approaches can be taken. The first is that the Rating Override function can be used to assign example.com to Nudity and Risque instead of Pornography for the purposes of matching the criteria that the law office goes by or the site can be assigned to a Custom Category that is not blocked because the site belongs to one of their clients and they always want to be able to access the site.

Another hypothetical example from the other side of the coin. A private school has decided that a company that specializes in the online selling of books that could be considered inappropriate for children because of their violent subject matter, should not be accessible to anyone in the school. The categorization by Fortinet of the site example2.com is General Interest - Business with the subcategory of Shopping and Auction, which is a category that is allowed at the school. In this case they school could reassign the site to the Category Adult Material which is a blocked category.

Local Categories

User-defined categories can be created to allow users to block groups of URLs on a per-profile basis. The categories defined here appear in the global URL category list when configuring a web filter profile. Users can rate URLs based on the local categories.

Users can create user-defined categories then specify the URLs that belong to the category. This allows users to block groups of web sites on a per profile basis. The ratings are included in the global URL list with associated categories and compared in the same way the URL block list is processed.

The local assignment of a category overrides the FortiGuard server ratings and appear in reports as "Local" Categories or "Custom" Categories depending on the context.

Local categories are configured in the CLI.

To create a Local Category:

```
config webfilter ftgd-local-cat
  edit local_category_1
    set id 140
    end
```

There is also a way to create a new category in the Web Based Manager.

Go to *Security Profiles* > *Web Filter* > *Rating Overrides*. In the process creating a new override, you can choose to override that URL to the *Category* "Custom Categories". When you are choosing the *Sub-Category* from the drop down menu, the last item in the list will be *Create New*. If you select this item you will be given the option to fill out a field called *Category Name*. Just type in the name of the new category from this point on it will be listed in the *Custom Categories* of the *Rating Overrides* and as one of the Local Categories in the FortiGuard Webfilter.

Configuring Rating Overrides

1. Go to *Security Profiles* > *Web Filter* > *Rating Overrides.*
2. Select Create New
3. Type in the URL field the URL of the Website that you wish to recategorize.
4. Select the Lookup Rating button to verify the current categorization that is being assigned to the URL.
5. Change the Category field to one of the more applicable options from the drop down menu.
6. Change the Sub-Category field to a more narrowly defined option within the main category.
7. Select OK.

It is usually recommended that you choose a category that you know will be addressed in existing Webfilter profiles so that you will not need to engage in further configuration.

Using Alternate Profiles

Allow Blocked Overrides or Web Overrides

Depending on which patch level you are on of Firmware Version 5 you will have one of 2 options for the implementation of the alternate profile approach to overriding a Web Filter profile that blocks access to a URL.

Before FortiOS 5 Patch 4

The only feature available was referred to as "Allow Blocked Override". The configuration settings for this feature where found in the individual Web Filter profiles by going to the Advanced section near the bottom of the configuration window.

Starting with FortiOS Patch 4

The Administrative Override feature for Web Filtering was added and is found by going to *Security Profiles > Web Filter > Web Overrides*. This opening window will display a listing of all of the overrides of this type. The editing window referred to the configuration as an Administrative Override.

The Concept

When a Web filter profile is overridden it does not necessarily remove all control and restrictions that were previously imposed by the Web Filter. The idea is to replace a restrictive filter with a different one. In practice, it makes sense that this will likely be a profile that is less restrictive the the original one but there is nothing that forces this. The degree to which that the alternate profile is less restrictive is open. It can be as much as letting the user access everything on the Internet or as little as allowing only one addition website. The usual practice though is to have as few alternate profiles as are needed to allow approved people to access what they need during periods when an exception to the normal rules is needed but still having enough control that the organizations web usage policies are not compromised.

You are not restricted to having only one alternative profile as an option to the existing profile. The new profile depends on the credentials or IP address making the connection. For example, John connecting through the "Standard" profile could get the "Allow_Streaming_Video" profile while George would get the "Allow_Social_Networking_Sites" profile.

The other thing to take into account is the time factor on these overrides. They are not indefinite. The longest that an override can be enabled is for 1 year less a minute. Often these overrides are set up for short periods of time for specific reasons such as a project. Having the time limitation means that the System Administrator does not have to remember to go back and turn the feature off after the project is finished.

Identity or Address

In either case what these override features do is, for specified users, user groups or IP addresses, allow sites blocked by Web Filtering profiles to be overridden for a specified length of time. The drawback of this method of override is that it takes more planning and preparation than the rating override method. The advantage is that once this has been set up, this method requires very little in the way of administrative overhead to maintain.

When planning to use the alternative profile approach keep in mind the following: In Boolean terms, one of the following "AND" conditions has to be met before overriding the Web Filter is possible

Based on the IP address:

- The Web Filter profile must be specified as allowing overrides
- AND the user's computer is one of the IP addresses specified
- AND the time is within the expiration time frame.

While the conditions are fewer for this situation there is less control over who has the ability to bypass the filtering configured for the site. All someone has to do is get on a computers that is allowed to override the Web Filter and they have access.

Based on user or user group:

• The Web Filter profile must be specified as allowing overrides

• AND the policy the traffic is going through must be identity based

• AND the user's credentials matches the identity credentials specified

• AND the time is within the expiration time frame.

This method is the one most likely to be used as it gives more control in that the user has to have the correct credential and more versatile because the user can use the feature from any computer that uses the correct policy to get out on the Internet.

Settings

When using an alternate profile approach to Web Filter overrides the following settings are used to determine authentication and outcome. Not every setting is used in both methods but enough of them are common to describe them collectively.

Apply to Group(s)

This is found in the Allow Blocked Overrides configuration. Individual users can not be selected. You can select one or more of the User Groups that are recognized my the FortiGate unit, whether they are local to the system or from a third part authentication device such as a AD server through FSSO.

Original Profile

This is found in the Administrative Override configuration. In the Allow Blocked Overrides setting the configuration is right inside the profile so there was no need to specify which profile was the original one, but the Administrative Override setup is done separately from the profiles themselves.

Assign to Profile or New Profile

Despite the difference in the name of the field, this is the same thing in both variations of the feature. You select from the drop down menu the alternate Web Filter Profile that you wish to set up for this override.

Scope or Scope Range

When setting up the override in the "Allow Blocked Overrides" variation you are given a drop down menu next to the field name Scope while in the Administrative Override configuration you are asked to select a radio button next to the same options. In both cases this is just a way of selecting which form of credentials will be required to approve the overriding of the existing Web Filter profile.

When the Web Filter Block Override message page appears it will display a field named "Scope:" and depending on the selection, it will show the type of credentials used to determine whether or not the override is allowed. The available options are:

User

This means that the authentication for permission to override will be based on whether or not the user is using a specific user account.

User Group

This means that the authentication for permission to override will be based on whether on not the user account supplied as a credential is a member of the specified User Group.

IP

This means that the authentication for permission to override will be based on the IP address of the computer that was used to authenticate. This would be used with computers that have multiple users. Example: If Paul logs on to the computer, engages the override using

his credentials and then logs off, if the scope was based on the IP address of the computer, anybody logging in with any account on that computer would now be using the alternate override Web Filter profile.

When entering an IP address in the Administrative Override version, only individual IP addresses are allowed.

Differences between IP and Identity based scope

- Using the IP scope does not require the use of an Identity based policy.
- When using the Administrative Override variation and IP scope, you may not see a warning message when you change from using the original Web Filter profile to using the alternate profile. There is no requirement for credentials from the user so, if allowed, the page will just come up in the browser.

Ask

This option is available only in the "Allowed Blocked Overrides" variation and when used configures the message page to ask which scope the user wished to use. Normally, when the page appears the scope options are greyed out an not editable, but by using the ask option the option is dark and the user can choose from the choice of:

- User
- User Group
- IP Address

Duration Mode

This option is available only in the "Allowed Blocked Overrides" variation. The Administrative Override sets a specified time frame that is always used for that override. The available options from the drop down menu are:

Constant

Using this setting will mean that what ever is set as the duration will be the length of time that the override will be in effect. If the Duration variable is set to 15 minutes the length of the override will always be 15 minutes. The option will be visible in the Override message page but the setting will be greyed out.

Ask

Uusing this setting will give the person the option of setting the duration to the override when it is engaged. The duration time which is greyed out if the Constant setting is used will be dark and editable. The user can set the duration in terms of Day, Hours and or Minutes.

Duration

Duration is on of the areas where the two variations takes a different approach, on two aspects of the setting. As already indicated the "Administrative Override" only uses a static time frame there is no option for the user to select on the fly how long it will last. The other way in which the two variation differ is that the "Allow Blocked Overrides" starts the clock when the user logs in with his credentials. For example, if the duration is 1 hour and John initiates an override at 2:00 p.m. on January 1, at the end of that hour he will revert back to using the original profile but he can go back and re-authenticate and and start the process over again. The Administrative override variation starts the clock from when the override was configured, which is why is shows an expiration date and time when your are configuring it.

This option, which is available when the Duration Mode is set to Constant is the time in minutes that the override will last when engaged by the user.

When setting up a constant duration in the Web Based Interface, minutes is the only option for units of time. To set a longer time frame or to use the units of hours or days you can use the CLI.

```
config webfilter profile
   edit <name of webfilter profile>
      config override
         set ovrd-dur <###d##h##m>
         end
```

When configuring the duration you don't have to set a value for a unit you are not using. If you are not using days or hours you can use

```
         set ovrd-dur 30m
```

instead of

```
         set ovrd-dur 0d0h30m
```

However, each of the units of time variable has their own maximum level

```
###d cannot be more than 364
##h cannot be more than 23
##m cannot be more than 59
```

So the maximum length that the override duration can be set to is 364 days, 23 hours, and 59 minutes(a minute shy of 1 year)

SafeSearch

SafeSearch is a feature of popular search sites that prevents explicit web sites and images from appearing in search results. Although SafeSearch is a useful tool, especially in educational environments, the resourceful user may be able to simply turn it off. Enabling SafeSearch for the supported search sites enforces its use by rewriting the search URL to include the code to indicate the use of the SafeSearch feature. For example, on a Google search it would mean adding the string "&safe=active" to the URL in the search.

The search sites supported are:

- Google
- Yahoo
- Bing
- Yandex

Enabling SafeSearch — CLI

```
config webfilter profile
   edit default
      config web
         set safe-search url
      end
   end
```

This enforces the use of SafeSearch in traffic controlled by the firewall policies using the web filter you configure.

YouTube Education Filter

YouTube for Schools is a way to access educational videos from inside a school network. This YouTube feature gives schools the ability to access a broad set of educational videos on YouTube EDU and to select the specific videos that are accessible from within the school network.

Before this feature can be used an account has to be set up for the school with YouTube. Once the account is set up a unique ID will be provided. This ID becomes part of the filter that is used to all access to the educational content of YouTube for use in schools even if YouTube is blocked by the policy.

More details can be found by going to http://www.youtube.com/schools.

Enabling YouTube Education Filter in CLI

```
config webfilter profile
  edit default
    config web
      set safe-search url header
      set youtube-edu-filter-id ABCD1234567890abcdef
      end
  end
```

Deep Scanning Restrictions

This section doesn't have a label such as "Deep Scanning Restrictions" but there are 2 settings in the profile that relate to the topic. In the profile, they appear as:

• *Enable HTTPS URL Scan Only*

• *Categories Exempt from Deep Scanning*...

Enable HTTPS URL Scan Only

When Deep Scanning is turned on traffic that is encrypted using SSL is scanned for issues just as unencrypted traffic is. However, scanning encrypted traffic puts a larger load on the resources of the FortiGate unit.

Even if the scanning of the contents of the traffic is not a requirement many administrator prefer to scan the URLs being sent over HTTPS so that users cannot bypass the blocking of access to a site by putting "https://" as a prefix to a URL. The setting restricts the deep scanning of the traffic to the URL destination which is in the header. This way the resources tied up in decrypting the traffic a are minimized, yet the administrator can still enforce policy regarding access to prohibited websites

Categories Exempt from Deep Scanning

For the purposes of personal privacy, there are 3 categories that can be exempted from deep scanning by the FortiGate unit. They are Banking, Health Care and Personal Privacy.

When HTTPS URL Scan Only is enabled you will notice that the option to exclude these categories from deep scanning is removed. This is because if only the URL is being scanned then the contents of the traffic is not being scanned anyway so there is no need to exclude it.

Web Site Filter

You can allow or block access to specific URLs by adding them to the Web Site Filter list. You add the URLs by using patterns containing text and regular expressions. The FortiGate unit allows or blocks web pages matching any specified URLs or patterns and displays a replacement message instead.

URL blocking does not block access to other services that users can access with a web browser. For example, URL blocking does not block access to ftp:// ftp.example.com. Instead, use firewall policies to deny ftp connections.

When adding a URL to the URL filter list, follow these rules:

- Type a top-level URL or IP address to control access to all pages on a web site. For example, `www.example.com` or `192.168.144.155` controls access to all pages at this web site.
- Enter a top-level URL followed by the path and file name to control access to a single page on a web site. For example, `www.example.com/news.html` or `192.168.144.155/news.html` controls access to the news page on this web site.
- To control access to all pages with a URL that ends with example.com, add `example.com` to the filter list. For example, adding example.com controls access to www.example.com, mail.example.com, www.finance.example.com, and so on.
- Control access to all URLs that match patterns using text and regular expressions (or wildcard characters). For example, `example.*` matches example.com, example.org, example.net and so on.

URLs with an action set to exempt or monitor are not scanned for viruses. If users on the network download files through the FortiGate unit from a trusted web site, add the URL of this web site to the URL filter list with an action to pass it so the FortiGate unit does not virus scan files downloaded from this URL.

URL formats

When adding a URL to the URL filter list, follow these rules:

How URL formats are detected when using HTTPS

If your unit does not support SSL content scanning and inspection or if you have selected the *URL filtering* option in web content profile for *HTTPS content filtering mode* under *Protocol Recognition*, filter HTTPS traffic by entering a top level domain name, for example, `www.example.com`. HTTPS URL filtering of encrypted sessions works by extracting the CN from the server certificate during the SSL negotiation. Since the CN only contains the domain name of the site being accessed, web filtering of encrypted HTTPS sessions can only filter by domain names.

If your unit supports SSL content scanning and inspection and if you have selected Deep Scan, you can filter HTTPS traffic in the same way as HTTP traffic.

How URL formats are detected when using HTTP

URLs with an action set to exempt are not scanned for viruses. If users on the network download files through the unit from trusted web site, add the URL of this web site to the URL

filter list with an action set to exempt so the unit does not virus scan files downloaded from this URL.

- Type a top-level URL or IP address to control access to all pages on a web site. For example, `www.example.com` or `192.168.144.155` controls access to all pages at this web site.

- Enter a top-level URL followed by the path and filename to control access to a single page on a web site. For example, `www.example.com/news.html` or `192.168.144.155/news.html` controls the news page on this web site.

- To control access to all pages with a URL that ends with `example.com`, add `example.com` to the filter list. For example, adding `example.com` controls access to `www.example.com`, `mail.example.com`, `www.finance.example.com`, and so on.

- Control access to all URLs that match patterns created using text and regular expressions (or wildcard characters). For example, `example.*` matches `example.com`, `example.org`, `example.net` and so on.

- Fortinet URL filtering supports standard regular expressions.

> If virtual domains are enabled on the unit, web filtering features are configured globally. To access these features, select *Global Configuration* on the main menu.

Web Site Filter actions

You can select one of four actions for how traffic will be treated as it attempts to reach a site in the list.

Block

Attempts to access any URLs matching the URL pattern are denied. The user will be presented with a replacement message.

Allow

Any attempt to access a URL that matches a URL pattern with an allow action is permitted. The traffic is passed to the remaining antivirus proxy operations, including FortiGuard Web Filter, web content filter, web script filters, and antivirus scanning.

Allow is the default action. If a URL does not appear in the URL list, it is permitted.

Monitor

Traffic to, and reply traffic from, sites matching a URL pattern with a monitor be allowed through in the same way as the "Allow" action. The difference with the Monitor action being that a log message will be generated each time a matching traffic session is established. The requests will also be subject to all other Security Profiles inspections that would normally be applied to the traffic.

Exempt

Exempt allows trusted traffic to bypass the antivirus proxy operations, but it functions slightly differently. In general, if you're not certain that you need to use the *Exempt* action, use *Monitor*.

HTTP 1.1 connections are persistent unless declared otherwise. This means the connections will remain in place until closed or the connection times out. When a client loads a web page, the client opens a connection to the web server. If the client follows a link to another page on

the same site before the connection times out, the same connection is used to request and receive the page data.

When you add a URL pattern to a URL filter list and apply the *Exempt* action, traffic sent to and replies traffic from sites matching the URL pattern will bypass all antivirus proxy operations. The connection itself inherits the exemption. This means that all subsequent reuse of the existing connection will also bypass all antivirus proxy operations. When the connection times out, the exemption is cancelled.

For example, consider a URL filter list that includes `example.com/files` configured with the Exempt action. A user opens a web browser and downloads a file from the URL `example.com/sample.zip`. This URL does not match the URL pattern so it is scanned for viruses. The user then downloads `example.com/files/beautiful.exe` and since this URL does match the pattern, the connection itself inherits the exempt action. The user then downloads `example.com/virus.zip`. Although this URL does not match the exempt URL pattern, a previously visited URL did, and since the connection inherited the exempt action and was re-used to download a file, the file is not scanned.

If the user next goes to an entirely different server, like `example.org/photos`, the connection to the current server cannot be reused. A new connection to example.org is established. This connection is not exempt. Unless the user goes back to example.com before the connection to that server times out, the server will close the connection. If the user returns after the connection is closed, a new connection to example.com is created and it is not exempt until the user visits a URL that matches the URL pattern.

Web servers typically have short time-out periods. A browser will download multiple components of a web page as quickly as possible by opening multiple connections. A web page that includes three photos will load more quickly if the browser opens four connections to the server and downloads the page and the three photos at the same time. A short time-out period on the connections will close the connections faster, allowing the server to avoid unnecessarily allocating resources for a long period. The HTTP session time-out is set by the server and will vary with the server software, version, and configuration.

Using the *Exempt* action can have unintended consequences in certain circumstances. You have a web site at example.com and since you control the site, you trust the contents and configure `example.com` as exempt. But example.com is hosted on a shared server with a dozen other different sites, each with a unique domain name. Because of the shared hosting, they also share the same IP address. If you visit example.com, your connection your site becomes exempt from any antivirus proxy operations. Visits to any of the 12 other sites on the same server will reuse the same connection and the data you receive is exempt from scanned.

Use of the *Exempt* action is not suitable for configuration in which connections through the FortiGate unit use an external proxy. For example, you use proxy.example.net for all outgoing web access. Also, as in the first example, URL filter list that includes a URL pattern of `example.com/files` configured with the *Exempt* action. Users are protected by the antivirus protection of the FortiGate unit until a user visits a URL that matches the of `example.com/files` URL pattern. The pattern is configured with the *Exempt* action so the connection to the server inherits the exemption. With a proxy however, the connection is from the user to the proxy. Therefore, the user is entirely unprotected until the connection times out, no matter what site he visits.

Ensure you are aware of the network topology involving any URLs to which you apply the Exempt action.

Status

The Web Site Filter has the option to either enable or disable individual web sites in the list. This allows for the temporary removal of the actions against a site so that it can be later reengaged without having to rewrite the configuration.

Configuring a Web Site Filter

To create a URL Filter list

1. Go to *Security Profiles > Web Filter > Profiles.*

2. Select the Web Filter Profile that you which to add the Web Site Filter to.

3. About half way down the Edit Web Filter Profile page check the box next to Enable Web Site Filter.

4. Select *Create New.*

5. Enter a URL for the website.

6. Enter optional comments to describe it.

7. Select *OK.*

Configuring a URL filter list

Each URL filter list can have up to 5000 entries. For this example, the URL www.example*.com will be used. You configure the list by adding one or more URLs to it.

To add a URL to a URL filter list

1. Go to *Security Profiles > Web Filter > URL Filter.*

2. Select an existing list and choose *Edit*.

3. Select *Create New*.

4. Enter the URL, without the "http", for example: `www.example*.com`.

5. Select a *Type*: *Simple*, *Wildcard* or *Regular Expression*.

6. In this example, select Wildcard.

7. Select the *Action* to take against matching URLs: Exempt, Block, Allow, or Monitor.

8. Select *Enable*.

9. Select *OK.*

Web content filter

You can control web content by blocking access to web pages containing specific words or patterns. This helps to prevent access to pages with questionable material. You can also add words, phrases, patterns, wild cards and Perl regular expressions to match content on web pages. You can add multiple web content filter lists and then select the best web content filter list for each web filter profile.

Enabling web content filtering involves three separate parts of the FortiGate configuration.

• The security policy allows certain network traffic based on the sender, receiver, interface, traffic type, and time of day.

• The web filter profile specifies what sort of web filtering is applied.

• The web content filter list contains blocked and exempt patterns.

The web content filter feature scans the content of every web page that is accepted by a security policy. The system administrator can specify banned words and phrases and attach a numerical value, or score, to the importance of those words and phrases. When the web content filter scan detects banned content, it adds the scores of banned words and phrases in the page. If the sum is higher than a threshold set in the web filter profile, the FortiGate unit blocks the page.

General configuration steps

Follow the configuration procedures in the order given. Also, note that if you perform any additional actions between procedures, your configuration may have different results.

1. Create a web content filter list.
2. Add patterns of words, phrases, wildcards, and regular expressions that match the content to be blocked or exempted.
3. You can add the patterns in any order to the list. You need to add at least one pattern that blocks content.
4. In a web filter profile, enable the web content filter and select a web content filter list from the options list.

To complete the configuration, you need to select a security policy or create a new one. Then, in the security policy, enable *Webfilter* and select the appropriate web filter profile from the list.

Creating a web filter content list

You can create multiple content lists and then select the best one for each web filter profile. Creating your own web content lists can be accomplished only using the CLI.

This example shows how to create a web content list called inappropriate language, with two entries, offensive and rude.

To create a web filter content list

```
config webfilter content
   edit 3
      set name "inappropriate language"
      config entries
        edit offensive
           set action block
           set lang western
           set pattern-type wildcard
           set score 15
           set status enable
        next
        edit rude
           set action block
           set lang western
           set pattern-type wildcard
           set score 5
           set status enable
        end
      end
   end
```

How content is evaluated

Every time the web content filter detects banned content on a web page, it adds the score for that content to the sum of scores for that web page. You set this score when you create a new pattern to block the content. The score can be any number from zero to 99999. Higher scores indicate more offensive content. When the sum of scores equals or exceeds the threshold score, the web page is blocked. The default score for web content filter is 10 and the default

threshold is 10. This means that by default a web page is blocked by a single match. Blocked pages are replaced with a message indicating that the page is not accessible according to the Internet usage policy.

Banned words or phrases are evaluated according to the following rules:

- The score for each word or phrase is counted only once, even if that word or phrase appears many times in the web page.
- The score for any word in a phrase without quotation marks is counted.
- The score for a phrase in quotation marks is counted only if it appears exactly as written.

The following table describes how these rules are applied to the contents of a web page. Consider the following, a web page that contains only this sentence: "The score for each word or phrase is counted only once, even if that word or phrase appears many times in the web page."

Table 7: Banned Pattern Rules

	Assigned score	Score added to the sum for the entire page	Threshold score	Comment
word	20	20	20	Appears twice but only counted once. Web page is blocked.
word phrase	20	40	20	Each word appears twice but only counted once giving a total score of 40. Web page is blocked
word sentence	20	20	20	"word" appears twice, "sentence" does not appear, but since any word in a phrase without quotation marks is counted, the score for this pattern is 20. Web page is blocked.
"word sentence"	20	0	20	"This phrase does not appear exactly as written. Web page is allowed.
"word or phrase"	20	20	20	This phrase appears twice but is counted only once. Web page is blocked.

Enabling the web content filter and setting the content threshold

When you enable the web content filter, the web filter will block any web pages when the sum of scores for banned content on that page exceeds the content block threshold. The threshold will be disregarded for any exemptions within the web filter list.

To enable the web content filter and set the content block threshold

1. Go to *Security Profiles > Web Filter > Profiles.*

2. Select the *Create New* icon on the Edit Web Filter Profile window title bar.

3. In the *Name* field, enter the name of the new web filter profile.

4. Optionally, you may also enter a comment. The comment can remind you of the details of the sensor.

5. Select the *Inspection Method*.

   Proxy-based detection involves buffering the file and examining it as a whole. Advantages of proxy-based detection include a more thorough examination of attachments, especially archive formats and nesting.

   Flow-based detection examines the file as it passes through the FortiGate unit without any buffering. Advantages of flow-based detection include speed and no interruption of detection during conserve mode.

6. Expand the *Advanced Filter* heading.

7. Enable *Web Content Filter*.

8. Select the required web filter content list from the *Web Content Filter* drop-down list.

9. Select *Apply*.

The web filter profile configured with web content filtering is ready to be added to a firewall profile.

Advanced web filter configurations

Allow websites when a rating error occurs

Enable to allow access to web pages that return a rating error from the FortiGuard Web Filter service.

If your FortiGate unit cannot contact the FortiGuard service temporarily, this setting determines what access the FortiGate unit allows until contact is re-established. If enabled, users will have full unfiltered access to all web sites. If disabled, users will not be allowed access to any web sites.

ActiveX filter

Enable to filter ActiveX scripts from web traffic. Web sites using ActiveX may not function properly with this filter enabled.

Block HTTP redirects by rating

Enable to block HTTP redirects.

Many web sites use HTTP redirects legitimately but in some cases, redirects may be designed specifically to circumvent web filtering, as the initial web page could have a different rating than the destination web page of the redirect.

This option is not supported for HTTPS.

Block Invalid URLs

Select to block web sites when their SSL certificate CN field does not contain a valid domain name.

FortiGate units always validate the CN field, regardless of whether this option is enabled. However, if this option is not selected, the following behavior occurs:

- If the request is made directly to the web server, rather than a web server proxy, the FortiGate unit queries for FortiGuard Web Filtering category or class ratings using the IP address only, not the domain name.
- If the request is to a web server proxy, the real IP address of the web server is not known. Therefore, rating queries by either or both the IP address and the domain name is not reliable. In this case, the FortiGate unit does not perform FortiGuard Web Filtering.

Cookie filter

Enable to filter cookies from web traffic. Web sites using cookies may not function properly with this enabled.

Provide Details for Blocked HTTP 4xx and 5xx Errors

Enable to have the FortiGate unit display its own replacement message for 400 and 500-series HTTP errors. If the server error is allowed through, malicious or objectionable sites can use these common error pages to circumvent web filtering.

HTTP POST action

Select the action to take with HTTP POST traffic. HTTP POST is the command used by your browser when you send information, such as a form you have filled-out or a file you are uploading, to a web server.

The available actions include:

Normal	Allow use of the HTTP POST command as normal.
Comfort	Use client comforting to slowly send data to the web server as the FortiGate unit scans the file. Use this option to prevent a server time-out when scanning or other filtering is enabled for outgoing traffic.
	The client comforting settings used are those defined in the Proxy Options profile selected in the security policy. For more information, see "Configuring client comforting" on page 34.
Block	Block the HTTP POST command. This will limit users from sending information and files to web sites.
	When the post request is blocked, the FortiGate unit sends the http-post-block replacement message to the web browser attempting to use the command.

Java applet filter

Enable to filter java applets from web traffic. Web sites using java applets may not function properly with this filter enabled.

Rate Images by URL

Enable to have the FortiGate retrieve ratings for individual images in addition to web sites. Images in a blocked category are not displayed even if they are part of a site in an allowed category.

Blocked images are replaced on the originating web pages with blank place-holders. Rated image file types include GIF, JPEG, PNG, BMP, and TIFF.

Rate URLs by Domain and IP Address

Enable to have the FortiGate unit request the rating of the site by URL and IP address separately, providing additional security against attempts to bypass the FortiGuard Web Filter.

If the rating determined by the domain name and the rating determined by the IP address defer the Action that is enforce will be determined by a weighting assigned to the different categories. The higher weighted category will take precedence in determining the action. This will have the side effect that sometimes the Action will be determined by the classification based on the domain name and other times it will be determined by the classification that is based on the IP address.

FortiGuard Web Filter ratings for IP addresses are not updated as quickly as ratings for URLs. This can sometimes cause the FortiGate unit to allow access to sites that should be blocked, or to block sites that should be allowed.

An example of how this would work would be if a URL's rating based on the domain name indicated that it belonged in the category Lingerie and Swimsuit, which is allowed but the category assigned to the IP address was Pornography which has an action of Block, because the Pornography category has a higher weight the effective action is Block.

Web resume download block

Enable to prevent the resumption of a file download where it was previously interrupted. With this filter enabled, any attempt to restart an aborted download will download the file from the beginning rather than resuming from where it left off.

This prevents the unintentional download of viruses hidden in fragmented files.

Note that some types of files, such as PDF, fragment files to increase download speed and enabling this option can cause download interruptions. Enabling this option may also break certain applications that use the Range Header in the HTTP protocol, such as YUM, a Linux update manager.

Working with the Interface

In order to find out the status of your configuration it helps to understand the interface. In the Web filtering section there are a number of pages that you will need to be able to read.

Profile page

Lists each web filter profile that you created. On this page, you can edit, delete or create a new web filter profile. You are redirected to this page when you select *View List* on the Edit Web Filter Profile page.

Note: Web filtering overrides are profile-based, allowing a rule to be created that changes the web filter profile that applies to a user. An override link appears in all related blocked pages. This is available only in the CLI.

New Web Filter Profile page

Provides settings for configuring a web filter profile. Advanced features, such as web content filtering and FortiGuard web filtering, is configured in the CLI.

This page appears when you select *Create New* on the Edit Web Filter Profile page. If you are on the Profile page, and you select *Create New*, you will be redirected to the New Web Filter Profile page.

Note: Logging is enabled in the CLI.

The following explains the web filtering options in the Web Filtering menu. If your unit supports SSL content scanning and inspection you can also configure web filtering for HTTPS traffic.

If you want to configure advanced settings, such as web content filter, you must configure them within the CLI. Advanced settings also includes overrides.

This topic includes the following:

- Profile
- Browser cookie-based FortiGuard Web Filtering overrides
- URL Filter
- Rating Overrides

Profile

The Profile menu allows you to configure a web filter profile to apply to a firewall policy. A profile is specific information that defines how the traffic within a policy is examined and what action may be taken based on the examination.

Web profile configuration settings

The following are web filter profile configuration settings in *Security Profiles > Web Filter > Profiles*. If you want to configure advanced settings, such as FortiGuard web filtering overrides, you must configure these settings within the CLI.

Profile page

Lists each web filter profile that you created. On this page, you can edit, delete or create a new web filter profile. You are redirected to this page when you select *View List* on the Edit Web Filter Profile page.

Note: Web filtering overrides are profile-based, allowing a rule to be created that changes the web filter profile that applies to a user. An override link appears in all related blocked pages. This is available only in the CLI.

Create New	Creates a new web filter profile. When you select *Create New*, you are automatically redirected to the New Web Filter Profile page.
Edit	Modifies settings within a web filter profile. When you select *Edit*, you are automatically redirected to the Edit Web Filter Profile page.

Delete	Removes a web filter profile from within the list on the Profile page.
	To remove multiple web filter profiles from within the list, on the Profile page, in each of the rows of the file filter lists you want removed, select the check box and then select *Delete*.
	To remove all web filter profiles from the list, on the Profile page, select the check box in the check box column and then select *Delete*.
Name	The name of the web filter profile.
Comments	A description given to the web filter profile. This is an optional setting.
Ref.	Displays the number of times the object is referenced to other objects. For example, av_1 profile is applied to a firewall policy; on the Profile page (*Security Profiles > Antivirus > Profile*), 1 appears in *Ref.*.
	To view the location of the referenced object, select the number in *Ref.*, and the Object Usage window appears displaying the various locations of the referenced object.
	To view more information about how the object is being used, use one of the following icons that is aviable within the Object Usage window:
	• **View the list page for these objects** – automatically redirects you to the list page where the object is referenced at.
	• **Edit this object** – modifies settings within that particular setting that the object is referenced with. For example, av_1 profile is referenced with a firewall policy and so, when this icon is selected, the user is redirected to the Edit Policy page.
	• **View the details for this object** – table, similar to the log viewer table, contains information about what settings are configured within that particular setting that the object is referenced with. For example, av_1 profile is referenced with a firewall policy, and that firewall policy's settings appear within the table.

New Web Filter Profile page

Provides settings for configuring a web filter profile. Advanced features, such as web content filtering and FortiGuard web filtering, is configured in the CLI.

This page appears when you select *Create New* on the Edit Web Filter Profile page. If you are on the Profiile page, and you select *Create New*, you will be redirected to the New Web Filter Profile page.

Note: Logging is enabled in the CLI.

Name	Enter a name for the web filter profile.
	If you want to edit the name at any time, select the profile and enter a new name in the *Name* field. Select *Apply* to save the change.

Comments	Enter a description for the web filter profile. This is optional.
	If you want to edit the description at any time, select the profile and enter the new description in the *Comments* field. Select *Apply* to save the change.
Inspection mode	Select to enable either flow-based web filtering or proxy-based.
	Flow-based web filtering is a non-proxy solution, which provides high concurrent session, high session rate, and low-latency web filtering service.
FortiGuard Categories	A list of FortiGuard category groups and categories that are used to rate web sites. Selecting a category group will automatically select all of the categories within the group. For example, if you select Security Risk, you can see that all of the categories within are selected if you expand the group. You can however, select or deselect categories within groups as required.
Show	Select an action to view all of the categories that are currently configured with the selected action.
Change Action for Selected Categories to	Select an action, and all of the selected categories will have the selected action applied. Selected category groups will have the action applied to all categories within the group.
Quota on Categories	Users can have their web browsing time limited by category through the use of quotas. Quotas can be applied only to categories that are configured with the Monitor action.
	If you create a quota for a single category, every authenticated user subject to the security policy in which the web filter profile is applied is limited in browsing web sites in the category to the duration you specify. If you create a single quota that includes multiple categories, the quota will apply to the categories as a whole.
	Quotas are ignored for unauthenticated users. To enforce quotas, configure the security policy to require authentication.
Enable Safe Search (Support Search Engines: Google, Yahoo and Bing)	When enabled, the supported search engines exclude offensive material from search results.
HTTPS Scanning	Available only on models that support HTTPS.
	Select to have all of the web filtering specified in the web filter profile to HTTPS traffic as well as HTTP traffic.
Advanced Filter	Expand this heading for advanced web filtering options.
Web URL Filter	Enable to block access to URLs listed in the selected URL list.

Web Resume Download Block	Enable to prevent the resumption of a file download where it was previously interrupted. With this filter enabled, any attempt to restart an aborted download will download the file from the beginning rather than resuming from where it left off.
	This prevents the unintentional download of viruses hidden in fragmented files.
	Note that some types of files, such as PDF, fragment files to increase download speed and enabling this option can cause download interruptions. Enabling this option may also break certain applications that use the Range Header in the HTTP protocol, such as YUM, a Linux update manager.
Block Invalid URLs	Select to block web sites when their SSL certificate CN field does not contain a valid domain name.
	FortiGate units always validate the CN field, regardless of whether this option is enabled. However, if this option is not selected, the following behavior occurs:
	• If the request is made directly to the web server, rather than a web server proxy, the FortiGate unit queries for FortiGuard Web Filtering category or class ratings using the IP address only, not the domain name.
	• If the request is to a web server proxy, the real IP address of the web server is not known. Therefore, rating queries by either or both the IP address and the domain name is not reliable. In this case, the FortiGate unit does not perform FortiGuard Web Filtering.
HTTP POST Action	Select the action to take with HTTP POST traffic. HTTP POST is the command used by your browser when you send information, such as a form you have filled-out or a file you are uploading, to a web server.
	The available actions include:
	• Normal: Allow use of the HTTP POST command as normal.
	• Comfort: Use client comforting to slowly send data to the web server as the FortiGate unit scans the file. Use this option to prevent a server time-out when scanning or other filtering is enabled for outgoing traffic. The client comforting settings used are those defined in the Proxy Options profile selected in the security policy.
	• Block: Block the HTTP POST command. This will limit users from sending information and files to web sites. When the post request is blocked, the FortiGate unit sends the http-post-block replacement message to the web browser attempting to use the command.
Remove Java Applet Filter	Enable to filter java applets from web traffic. Web sites using java applets may not function properly with this filter enabled.
Remove ActiveX Filter	Enable to filter ActiveX scripts from web traffic. Web sites using ActiveX may not function properly with this filter enabled.

Remove Cookie Filter	Enable to filter cookies from web traffic. Web sites using cookies may not function properly with this enabled.
Search Engine Keyword Filter	Enter the keywords that you want to monitor when users enter those same or similar keywords during a search within the supported search engines.
Web Content Filter	Enable to block access to web pages that include the words included in the selected web content filter list.
Provide Details for Blocked HTTP 4xx and 5xx Errors	Enable to have the FortiGate unit display its own replacement message for 400 and 500-series HTTP errors. If the server error is allowed through, malicious or objectionable sites can use these common error pages to circumvent web filtering.
Rate Images by URL (Blocked images will be replaced with blanks)	Enable to have the FortiGate retrieve ratings for individual images in addition to web sites. Images in a blocked category are not displayed even if they are part of a site in an allowed category.
	Blocked images are replaced on the originating web pages with blank place-holders. Rated image file types include GIF, JPEG, PNG, BMP, and TIFF.
Allow Websites When a Rating Error Occurs	Enable to allow access to web pages that return a rating error from the FortiGuard Web Filter service.
	If your FortiGate unit cannot contact the FortiGuard service temporarily, this setting determines what access the FortiGate unit allows until contact is re-established. If enabled, users will have full unfiltered access to all web sites. If disabled, users will not be allowed access to any web sites.
Strict Blocking	This setting determines when the FortiGate unit blocks a site. Enable strict blocking to deny access to a site if any category or classification assigned to the site is set to Block. Disable strict blocking to deny access to a site only if all categories and classifications assigned to the site are set to Block.
	All rated URLs are assigned one or more categories. URLs may also be assigned a classification. If Rate URLs by domain and IP address is enabled, the site URL and IP address each carry separately assigned categories and classifications. Depending on the FortiGuard rating and the FortiGate configuration, a site could be assigned to at least two categories and up to two classifications.

Rate URLs by Domain and IP Address	Enable to have the FortiGate unit request the rating of the site by URL and IP address separately, providing additional security against attempts to bypass the FortiGuard Web Filter.
	FortiGuard Web Filter ratings for IP addresses are not updated as quickly as ratings for URLs. This can sometimes cause the FortiGate unit to allow access to sites that should be blocked, or to block sites that should be allowed.
Block HTTP Redirects by Rating	Enable to block HTTP redirects.
	Many web sites use HTTP redirects legitimately but in some cases, redirects may be designed specifically to circumvent web filtering, as the initial web page could have a different rating than the destination web page of the redirect.
	This option is not supported for HTTPS.

Browser cookie-based FortiGuard Web Filtering overrides

By using browser cookie-based FortiGuard Web Filtering overrides, you can identify users according to their web browser cookie instead of their IP address and then to use this identification to apply FortiGuard Web Filtering overrides to individual users.

This feature uses the dynamic profile feature to assign a web filter profile that includes FortiGuard Web Filtering to a communication session. Just like normal FortiGuard Web Filtering overrides, when FortiGuard Web Filtering blocks access to a web page, the user can authenticate to override FortiGuard Web Filtering. However, with Browser cookie-based overrides enabled, the browser cookie is used to identify the user instead of the user's IP address.

You can also go to *Security Profiles > Web Filter > Configuration* and configure the following browser cookie-based override settings.

Additional browser cookie-based configuration settings are available from the CLI using the `config webfilter cookie-ovrd` command.

Cookie (Browser Based) Override Configuration page

Provides settings for configuring the browser cookie-based override.

Override Validation Hostname Enter the override validation hostname in the field.

Override Validation Port Enter the port number in the field.

How browser cookie-based FortiGuard Web Filtering overrides work

The following steps occur when a user's session that can use browser cookie-based FortiGuard Web Filtering overrides is received:

1. The Dynamic Profile applies a profile to the user session in the normal way.
2. The user issues a request to a remote site blocked by FortiGuard Web Filtering.
3. For example, http://www.example.com.
4. FortiGuard Web Filtering blocks the page and provides an override link.

5. The user selects the override option and successfully authenticates.

6. The unit sends a cookie to the remote site that seems to come from the *Override Validation Hostname*.

7. The unit creates a second cookie to the user's browser for the domain of the remote site.

   For example, the domain could be example.com.

The rest of the communication between the user and the remote site is authorized with the unit by these cookies

URL Filter

Allow or block access to specific URLs by adding them to the URL filter list. Add patterns using text and regular expressions (or wildcard characters) to allow or block URLs. The unit allows or blocks web pages matching any specified URLs or patterns and displays a replacement message.

You can add multiple URL filter lists and then select the best URL filter list for each profile.

You can add the following to block or exempt URLs:

- complete URLs
- IP addresses
- partial URLs to allow or block all sub-domains

Each URL filter list can have up to 5000 entries.

URL filter configuration settings

The following are URL filter configuration settings in *Security Profiles > Web Filter > URL Filter*.

URL blocking does not block access to other services that users can access with a web browser. For example, URL blocking does not block access to `ftp://ftp.example.com`. Instead, use firewall policies to deny FTP connections.

URL Filter page

Lists each URL filter that you created. On this page, you can edit, delete or create a new URL filter.

Create New	Creates a new URL filter list. When you select *Create New*, you are automatically redirected to the New List page. This page provides a name field and comment field. You must enter a name to go to the URL Filter Settings page.
Edit	Modifies settings within a URL filter list. When you select *Edit*, you are automatically redirected to the URL Filter Settings page.

Delete	Removes the URL filter list from the list on the URL Filter page. The *Delete* icon is only available if the URL filter list is not selected in any profiles.
	To remove multiple URL filter list from within the list, on the URL Filter page, in each of the rows of the file filter lists you want removed, select the check box and then select *Delete*.
	To remove all URL filter list from the list, on the URL Filter page, select the check box in the check box column and then select *Delete*.
Name	The available URL filter lists.
# Entries	The number of URL patterns in each URL filter list.
MMS Profiles (FortiOS Carrier only)	The name of the MMS profile
Comments	Optional description of each URL filter list.
Ref.	Displays the number of times the object is referenced to other objects. For example, av_1 profile is applied to a firewall policy; on the Profile page (*Security Profiles > AntiVirus > Profiles*), 1 appears in *Ref.*.
	To view the location of the referenced object, select the number in *Ref.*, and the Object Usage window appears displaying the various locations of the referenced object.
	To view more information about how the object is being used, use one of the following icons that is avialable within the Object Usage window:
	• **View the list page for these objects** – automatically redirects you to the list page where the object is referenced at.
	• **Edit this object** – modifies settings within that particular setting that the object is referenced with. For example, av_1 profile is referenced with a firewall policy and so, when this icon is selected, the user is redirected to the Edit Policy page.
	• **View the details for this object** – table, similar to the log viewer table, contains information about what settings are configured within that particular setting that the object is referenced with. For example, av_1 profile is referenced with a firewall policy, and that firewall policy's settings appear within the table.

URL Filter Settings page

Provides settings for configuring URLs that make up the URL filter, and also lists the URLs that you created. You are automatically redirected to this page from the New List Page. If you are editing a URL filter, you are automatically redirected to this page.

Name	If you are editing an existing URL filter setting and want to change the name, enter a new name in this field. You must select *OK* to save the change.
Comments	If you are editing an existing URL filter setting and want to change or add a description, enter the new text in this field. You must select *OK* to save these changes.
Create New	Adds a URL address and filter settings to the list. When you select *Create New*, you are automatically redirected to the New URL Filter list.

Edit	Modifies the settings within a URL filter.
Delete	Removes an entry from the list.
	To remove multiple URL filters from within the list, on the URL Filter Settings page, in each of the rows of the filters you want removed, select the check box and then select *Delete*.
	To remove all URL filters from the list, on the URL Filter Settings page, select the check box in the check box column and then select *Delete*.
Enable	Enables a filter in the list.
Disable	Disables a filter in the list.
Move To	Moves the URL to any position in the list. When you select *Move To*, the Move URL Filter window appears.
	To move a URL, select the new position *Before* or *After*, which will place the current URL entry before or after the entry you enter in the *(URL)* field. For example, 1example.com is being moved after 3example.com, so 3example.com is entered in the *(URL)* field.
Remove All Entries	Removes all filter entries within the list on the URL Filter Settings page.
Enable	Indicates whether the URL is enable or disabled. A green check mark indicates that the URL is enabled; a gray check mark indicates that the URL is disabled.
URL	The URL address.
Action	The type of action the unit will take when there is a match.
Type	The type of URL. For example, the type of URL is *Regex*.

New URL Filter page

Provides settings for configuring a URL to add to the filter list.

URL	Enter the URL.
Type	Select a type from the drop-down list: *Simple*, *Regex* (regular expression), or *Wildcard*.

Action	Select an action the unit will take.
	• *Allow* – any attempt to access a URL that matches a URL pattern with an allow action is permitted.
	• *Exempt* – similar to *Pass* in that it allows trusted traffic to bypass the antivirus proxy operations, but it functions slightly differently; ensure you are aware of the network topology involving URLs that you applied the Exemption action. Additional information about the Exempt action is found in the Security Profiles chapter of the FortiOS Handbook.
	• *Block* – attempts to access any URLs matching the URL pattern are denied; user is presented with a replacement message.
	• *Pass* – traffic to, and replay traffic from sites that match a URL pattern with a pass action will bypass all antivirus proxy operations, including FortiGuard Web Filter, web content filter, web script filters, and antivirus scanning. Make sure you trust the content of any site you pass, otherwise there may be a security risk.
Enable	Select to enable the URL. By default, the URL is enabled.

Web filtering example

Web filtering is particularly important for protecting school-aged children. There are legal issues associated with improper web filtering as well as a moral responsibility not to allow children to view inappropriate material. The key is to design a web filtering system in such a way that students and staff do not fall under the same web filter profile in the FortiGate configuration. This is important because the staff may need to access websites that are off-limits to the students.

School district

The background for this scenario is a school district with more than 2300 students and 500 faculty and staff in a preschool, three elementary schools, a middle school, a high school, and a continuing education center. Each elementary school has a computer lab and the high school has three computer labs with connections to the Internet. Such easy access to the Internet ensures that every student touches a computer every day.

With such a diverse group of Internet users, it was not possible for the school district to set different Internet access levels. This meant that faculty and staff were unable to view websites that the school district had blocked. Another issue was the students' use of proxy sites to circumvent the previous web filtering system. A proxy server acts as a go-between for users seeking to view web pages from another server. If the proxy server has not been blocked by the school district, the students can access the blocked website.

When determining what websites are appropriate for each school, the district examined a number of factors, such as community standards and different needs of each school based on the age of the students.

The district decided to configure the FortiGate web filtering options to block content of an inappropriate nature and to allow each individual school to modify the options to suit the age of the students. This way, each individual school was able to add or remove blocked sites almost immediately and have greater control over their students' Internet usage.

In this simplified example of the scenario, the district wants to block any websites with the word *example* on them, as well as the website www.example.com. The first task is to create web content filter lists for the students and the teachers.

To create a web content filter list for the students

```
config webfilter content
   edit 5
      set name "Student Web Content List"
      config entries
         edit example
            set action block
            set status enable
      end
   end
```

It might be more efficient if the Teacher Web Content List included the same blocked content as the student list. From time to time a teacher might have to view a blocked page. It would then be a matter of changing the *Action* from *Block* to *Allow* as the situation required.

To create a web content filter list for the teachers

```
config webfilter content
   edit 5
      set name "Teacher Web Content List"
      config entries
         edit example
            set action exempt
            set status enable
      end
   end
```

URL filter lists with filters to block unwanted web sites must be created for the students and teachers. For this example the URL *www.example.com* will be used.

To create a URL filter for the students

1. Go to *Security Profiles > Web Filter > URL Filter*.

2. Select *Create New*.

3. Enter `Student URL List` as the URL filter *Name*.

4. Enter optional comments to describe the contents of the list.

5. Select *OK*.

   The URL filter for the students has been created. Now it must be configured.

6. Select *Create New*.

7. Enter `example.com` in the URL field.

8. Select *Simple* from the *Type* list.

9. Select *Block* from the *Action* list.

10. Select *Enable.*

11. Select *OK*.

12. Select *OK*.

The teachers should be able to view the students' blocked content, however, so an addition URL filter is needed.

To create a URL filter for the teachers

1. Go to *Security Profiles > Web Filter > URL Filter*.
2. Select *Create New*.
3. Enter `Teacher URL List` as the URL filter *Name*.
4. Enter optional comments to describe the list.
5. Select *OK*.

   The URL filter for the students has been created. Now it must be configured.
6. Select *Create New*.
7. Enter `www.example.com` in the *URL* field.
8. Select *Simple* from the *Type* list.
9. Select *Exempt* from the *Action* list.
10. Select *Enable*.
11. Select *OK*.
12. Select *OK*.

A web filter profile must be created for the students and the teachers.

To create a web filter profile for the students

1. Go to *Security Profiles > Web Filter > Profiles*.
2. Select the *Create New* icon in the Edit Web Filter window title bar.
3. Enter `Students` as the *Profile Name*.
4. Enter optional comments to identify the profile.
5. Expand the *Advanced Filter* heading.
6. Enable *Web Content Filter*.
7. Select *Student Web Content List* from the *Web Content Filter* drop-down list.
8. Enable *Web URL Filter*.
9. Select *Student URL List* from the *Web URL Filter* drop-down list.
10. Enable *Web Resume Download Block*.

    Selecting this setting will block downloading parts of a file that have already been downloaded and prevent the unintentional download of virus files hidden in fragmented files. Note that some types of files, such as PDFs, are fragmented to increase download speed, and that selecting this option can cause download interruptions with these types.
11. Select *OK*.

To create a security policy for the students

1. Go to *Policy > Policy > Policy.*
2. Select *Create New*.
3. Enable *Web Filter*.
4. Select *Students* from the web filter drop-down list.
5. Enter optional comments.
6. Select *OK*.

To create a web filter profile for the teachers

1. Go to *Security Profiles > Web Filter > Profiles*.
2. Select the *Create New* icon in the Edit Web Filter window title bar.
3. Enter `Teachers` as the *Profile Name*.
4. Enter optional comments to identify the profile.
5. Expand the *Advanced Filter* heading.
6. Enable *Web Content Filter*.
7. Select *Teacher Web Content List* from the *Web Content Filter* drop-down list.
8. Enable *Web URL Filter*.
9. Select *Teacher URL List* from the *Web URL Filter* drop-down list.
10. Enable *Web Resume Download Block*.
11. Select *OK*.

To create a security policy for Teachers

1. Go to *Policy > Policy > Policy.*
2. Select *Create New*.
3. Enable *Web Filter*.
4. Select *Teachers* from the web filter drop-down list.
5. Enter optional comments.

Select *OK*.

Data leak prevention

The FortiGate data leak prevention (DLP) system allows you to prevent sensitive data from leaving your network. When you define sensitive data patterns, data matching these patterns will be blocked, or logged and allowed, when passing through the FortiGate unit. You configure the DLP system by creating individual filters based on file type, file size, a regular expression, an advanced rule, or a compound rule, in a DLP sensor and assign the sensor to a security policy.

Although the primary use of the DLP feature is to stop sensitive data from leaving your network, it can also be used to prevent unwanted data from entering your network and to archive some or all of the content passing through the FortiGate unit.

This section describes how to configure the DLP settings.

The following topics are included:

- Data leak prevention concepts
- Enable data leak prevention
- Fingerprint
- File filter
- DLP archiving
- DLP examples

Data leak prevention concepts

Data leak prevention examines network traffic for data patterns you specify. You define whatever patterns you want the FortiGate unit to look for in network traffic. The DLP feature is broken down into a number of parts.

DLP sensor

A DLP sensor is a package of filters. To use DLP, you must enable it in a security policy and select the DLP sensor to use. The traffic controlled by the security policy will be searched for the patterns defined in the filters contained in the DLP sensor. Matching traffic will be passed or blocked according to how you configured the filters.

DLP filter

Each DLP sensor has one or more filters configured within it. Filters can examine traffic for known files using DLP fingerprints, for files of a particular type or name, for files larger than a specified size, for data matching a specified regular expression, or for traffic matching an advanced rule or compound rule.

You can configure the action taken when a match is detected. The actions include:

- None
- Log Only,
- Block
- Quarantine User,
- Quarantine IP address
- Quarantine Interface

Log Only is enabled by default.

Fingerprint

Fingerprint scanning allows you to create a library of files for the FortiGate unit to examine. It will create checksum fingerprints so each file can be easily identified. Then, when files appear in network traffic, the FortiGate will generate a checksum fingerprint and compare it to those in the fingerprint database. A match triggers the configured action.

File filter

File filters use file filter lists to examine network traffic for files that match either file names or file types. For example, you can create a file filter list that will find files called secret.* and also all JPEG graphic files. You can create multiple file filter lists and use them in filters in multiple DLP sensors as required.

File size

This filter-type checks for files exceeding a configured size. All files larger than the specified size are subject to the configured action.

Regular expression

The FortiGate unit checks network traffic for the regular expression specified in a regular expression filter. The regular expression library used by Fortinet is a variation of a library called PCRE (Perl Compatible Regular Expressions). A number of these filters can be added to a sensor making a sort of 'dictionary' subset within the sensor.

Some other, more limited DLP implementations, use a list of words in a text file to define what words are searched for. While the format used here is slightly different than what some people are used to, the resulting effect is similar. Each Regular Expression filter can be thought of as a more versatile word to be searched against. In this dictionary (or sensor), the list of words is not limited to just predefined words. It can include expressions that can accommodate complex variations on those words and even target phrases. Another advantage of the individual filter model of this dictionary over the list is that each word can be assigned its own action, making this implementation much more granular.

Watermark

Watermarking is essentially marking files with a digital pattern to mark the file as being proprietary to a specific company. Fortinet has a utility that will apply a digital watermark to files. The utility adds a small (approx. 100 byte) pattern to the file that is recognised by the DLP Watermark filter. the pattern is invisible to the end user.

When watermarking a file it should be verified that the pattern matches up to a category found on the FortiGate firewall. For example, if you are going to watermark a file with the sensitivity

level of "Secret" you should verify that "Secret" is a sensitivity level that has been assigned in the FortiGate unit.

Software Versions

Before planning on using watermarking software it is always best to verify that the software will work with your OS. Currently the utility was only available for the Linux and Windows operating systems.

The Linux version can be found in one of 3 command line executable programs.

- watermark_linux_amd64
- watermark_linux_arm
- watermark_linux_x86

The Windows version is part of the FortiExplorer software.

File types

The Watermark tool does not work with every file type. The following file types are supported by the watermark tool:

- .txt
- .pdf
- .doc
- .xls
- .ppt
- .docx
- .pptx
- .xlsx

Currently the DLP only works with Fortinet's watermarking software.

Using the FortiExplorer Watermark tool

The FortiExplorer software can be downloaded from the Fortinet Support Site.

1. Choose whether to "Apply Watermark To:"
   - Select File
   - Entire Directory
2. Fill in the fields:
   a. *Select File*

   This Field has a browse icon next to it which will allow the user to browse to and select a single file or directory to apply the water mark to.

   b. *Sensitivity Level*

   This field is a drop down menu that lists the available sensitivity levels that the FortiGate can scan for

   c. *Identifier*

   This is a unique identifier string of characters to identify the company that the document belongs to.

   d. *Output Directory*

   This Field has a browse icon next to it which will allow the user to browse to a directory where the altered file will be placed. If the output directory is the same as the source

directory the original file will be overwritten. If the output directory is different than the source directory then the watermarked version of the file will be place there and the unaltered original will be left in the source directory.

3. Select *Apply Watermark* to start the process.

You should get output in the window similar to this:

```
> fortinet-watermark-win.exe -v -f "C:\Users\TestUser\Documents\test
    document.txt" -i "123456ABC" -l "Private" -o
    "C:\Users\TestUser\Watermarked Documents"
  Creating watermark. Pattern:
  =-=-=-=-=-=-=-=-=-=-=-=-=-=-=-=-=-=identifier=123456ABC
   sensitivity=Private=-=-=-=-=-=-=-=-=-=-=-=-=-=-=-=-=-=
  --> 'C:\Users\TestUser\Documents\test document.txt'
  Inserted watermark size 231
  ------------------------------------------------------
  1 file(s) processed. (success = 1, failure = 0)
```

Installation of the watermark utility on Linux

Add the watermark file to a location on the system that is in the $PATH.

To see what the path is use the command

`echo $PATH`

Example results:

```
/usr/local/sbin:/usr/local/bin:/usr/sbin:/usr/bin:/sbin:/bin:/usr/ga
    mes
```

for example you could move or copy the file to the :/bin directory.

Permissions on the watermark file:

Check the existing permissions:

The command in Linux for listing file along with the permissions is:

`ls -l`

Run the check to see if the permission status. The results may be something along these lines:

`-rw-r--r-- 1 root root 2053868 Jan 10 11:44 watermark`

You will see that in this case it has no executable permissions

To change the permissions on the watermark file:

It will be assume for this command that the utility is in the bin directory and that you have ownership level access.

`chmod o+x /bin/watermark`

To verify the change:

`ls -l wa*`

`-rw-r--r-x 1 root root 2053868 Jan 10 11:44 watermark`

You can see how the x for executable has been added to the permissions for the *others* group.

Syntax of the Watermark utility

The tool is executed in a Linux environment by passing in files or directories of files to insert a watermark.

USAGE:

```
watermark <options> -f <file name> -i <identifier> -l <sensitivity
     level>
watermark <options> -d <directory> -i <identifier> -l <sensitivity
     level>
```

Options:

```
-h print help
-v verbose information
-I inplace watermarking (don't copy file)
-o output directory
-e encode <to non-readable>
-a add additional watermark (by default replaces watermarks existing
     watermarks)
-D delete all watermarks
```

Using the watermark utility

Now if you are in your home directory and you want to watermark a file in the Documents directory you could plan out the command like this:

watermark [because that is the executable to be used]

-v [so that you can get as much feedback as possible]

-I [because you don't want a new file you just want to watermark the existing one]

-f [because you only want to change the one file not the entire directory]

filename.pdf [the name of the file]

-i 123456 [to set the identifier to 123456 - this is a required setting

-l Private [to set the sensitivity level to "Private"]

Now at the command prompt enter all of these components in order:

```
watermark -v -I -f filename.pdf -i 12345 -l Private
Creating watermark.  Pattern:
=-=-=-=-=-=-=-=-=-=-=-=-=-=-=-=identifier=12345
     sensitivity=Private=-=-=-=-=-=-=-=-=-=-=-=-=-=-=-=
Watermarking file: 'filename.pdf'
Inserted watermark size 148
```

Enable data leak prevention

DLP examines your network traffic for data patterns you specify. The FortiGate unit then performs an action based on the which pattern is found and a configuration set for each filter trigger.

General configuration steps

Follow the configuration procedures in the order given. Also, note that if you perform any additional actions between procedures, your configuration may have different results.

1. Create a DLP sensor.

New DLP sensors are empty. You must create one or more filters in a sensor before it can examine network traffic.

2. Add one or more filters to the DLP sensor.

Each filter searches for a specific data pattern. When a pattern in the active DLP sensor appears in the traffic, the FortiGate unit takes the action configured in the matching filter. Because the order of filters within a sensor cannot be changed, you must configure DLP in sequence.

3. Add the DLP sensor to one or more firewall policies that control the traffic to be examined.

Creating a DLP sensor

DLP sensors are collections of filters. You must also specify an action for the filter when you create it in a sensor. Once a DLP sensor is configured, you can select it a security policy profile. Any traffic handled by the security policy will be examined according to the DLP sensor configuration.

To create a DLP sensor

1. Go to *Security Profiles > Data Leak Prevention > Sensor*.

2. Select the *Create New* icon on the Edit DLP Sensor window title bar.

3. In the *Name* field, enter the name of the new DLP sensor.

4. Optionally, you may also enter a comment. The comment appears in the DLP sensor list and can remind you of the details of the sensor.

5. Select *OK*.

The DLP sensor is created and the sensor configuration window appears.

6. Select OK.

A newly created sensor is empty, containing no filters. Without filters, the DLP sensor will do nothing.

Adding filters to a DLP sensor

Once you have created a DLP sensor, you need to add filters.

1. To add filters to a DLP sensor

2. Go to *Security Profiles > Data Leak Prevention> Sensor*.

3. Select the Sensor in the Edit DLP Sensor window title bar drop-down list.

4. Select *Create New*.

5. Enter a filter name.

6. Select the type of filter. You can choose either Messages or Files. Depending on which of these two are chosen different options will be available

Message filter will have these configuration options:

- [radio button] Containing: [drop down menu including: Credit Card # or SSN]
- [radio button] Regular Expression [input field]
- [radio button] Encrypted

Examine the following Services:

Web Access
- HTTP-POST

Email
- [check box] SMTP
- [check box] POP3

- [check box] IMAP
- [check box] MAPI

Others

- [check box] NNTP

Action [from drop down menu]

- None
- Log Only,
- Block
- Quarantine User,
- Quarantine IP address
- Quarantine Interface

Files filter will have these options:

- [radio button] Containing: drop down menu including: Credit Card # or SSN
- [radio button] File Size >= [ ]kb
- [radio button] File Type included in [drop down menu of File Filters]
- [radio button] File Finger Print : [drop down menu]
- [radio button] Watermark Sensitivity: [drop down menu] and Corporate Identifier [id field]
- [radio button] Regular Expression [input field]
- [radio button] Encrypted

Examine the following Services:

Web Access

- [check box] HTTP-POST
- [check box] HTTP-GET

Email

- [check box] SMTP
- [check box] POP3
- [check box] IMAP
- [check box] MAPI

Others

- [check box] FTP
- [check box] NNTP

Action [from drop down menu]

- None
- Log Only,
- Block
- Quarantine User,
- Quarantine IP address
- Quarantine Interface

Table 8: Option explanations

	• Credit Card numbers - The number formats used by American Express, Visa, and Mastercard credit cards are detected. • Social Security Numbers.

Table 9: Action Options

Action	Description
None	No action is taken if filter even if filter is triggered
Log Only	The FortiGate unit will take no action on network traffic matching a rule with this action. The filter match is logged, however. Other matching filters in the same sensor may still operate on matching traffic.
Block	Traffic matching a filter with the block action will not be delivered. The matching message or download is replaced with the data leak prevention replacement message.

Table 9: Action Options

Action	Description
Quarantine User	If the user is authenticated, this action blocks all traffic to or from the user using the protocol that triggered the rule and adds the user to the Banned User list. If the user is not authenticated, this action blocks all traffic of the protocol that triggered the rule from the user's IP address.
	If the banned user is using HTTP, FTP, or NNTP (or HTTPS if the FortiGate unit supports SSL content scanning and inspection) the FortiGate unit displays the "Banned by data leak prevention" replacement message. If the user is using IM, the IM and P2P "Banned by data leak prevention" message replaces the banned IM message and this message is forwarded to the recipient. If the user is using IMAP, POP3, or SMTP (or IMAPS, POP3S, SMTPS if your FortiGate unit supports SSL content scanning and inspection) the Mail "Banned by data leak prevention" message replaces the banned email message and this message is forwarded to the recipient. These replacement messages also replace all subsequent communication attempts until the user is removed from the banned user list.
	If this action is chosen the additional field for [ ] minutes will appear so that a time limit can be set for the duration of the quarantine. This field cannot be left blank.
Quarantine IP Address	This action blocks access for any IP address that sends traffic matching a filter with this action. The IP address is added to the Banned User list. The FortiGate unit displays the "NAC Quarantine DLP Message" replacement message for all connection attempts from this IP address until the IP address is removed from the banned user list.
	If this action is chosen the additional field for [ ] minutes will appear so that a time limit can be set for the duration of the quarantine. This field cannot be left blank.
Quarantine Interface	This action blocks access to the network for all users connecting to the interface that received traffic matching a filter with this action. The FortiGate unit displays the "NAC Quarantine DLP Message" replacement message for all connection attempts to the interface until the interface is removed from the banned user list.
	If this action is chosen the additional field for [ ] minutes will appear so that a time limit can be set for the duration of the quarantine. This field cannot be left blank.

Quarantine User, *Quarantine IP*, and *Quarantine Interface* provide functionality similar to NAC quarantine. However, these DLP actions block users and IP addresses at the application layer while NAC quarantine blocks IP addresses and interfaces at the network layer.

7. Select OK.

8. Repeat Steps 6 and 7 for each filter.

9. Select Apply to confirm the settings of the sensor.

If you have configured DLP to block IP addresses and if the FortiGate unit receives sessions that have passed through a NAT device, all traffic from that NAT device — not just traffic from individual users — could be blocked. You can avoid this problem by implementing authentication.

> To view or modify the replacement message text, go to *System > Config > Replacement Message*.

DLP document fingerprinting

One of the DLP techniques to detect sensitive data is fingerprinting (also called document fingerprinting). Most DLP techniques rely on you providing a characteristic of the file you want to detect, whether it's the file type, the file name, or part of the file contents. Fingerprinting is different in that you provide the file itself. The FortiGate unit then generates a checksum fingerprint and stores it. The FortiGate unit generates a fingerprint for all files detected in network traffic, and it is compared to all of the fingerprints stored in its fingerprint database. If a match is found, the configured action is taken.

The document fingerprint feature requires a FortiGate unit with internal storage. The document fingerprinting menu item does not appear on models without internal storage.

Any type of file can be detected by DLP fingerprinting and fingerprints can be saved for each revision of your files as they are updated.

To use fingerprinting you select the documents to be fingerprinted and then add fingerprinting filters to DLP sensors and add the sensors to firewall policies that accept the traffic to which to apply fingerprinting.

Fingerprinted Documents

The FortiGate unit must have access to the documents for which it generates fingerprints. One method is to manually upload documents to be fingerprinted directly to the FortiGate unit. The other is to allow the FortiGate unit to access a network share that contains the documents to be fingerprinted.

If only a few documents are to be fingerprinted, a manual upload may be the easiest solution. If many documents require fingerprinting, or if the fingerprinted documents are frequently revised, using a network share makes user access easier to manage.

To configure manual document fingerprints

1. Go to *Security Profiles > Data Leak Prevention > Document Fingerprinting*.

2. In the Manual Document Fingerprints section, select *Create New*.

3. Select the file to be fingerprinted.

4. Choose a Sensitivity level. The default choices are *Critical*, *Private* and *Warning*, but more can be added in the CLI.

5. If the file is an archive containing other files, select *Process files inside archive* if you also want the individual files inside the archive to have fingerprints generated in addition to the archive itself.

6. Select *OK*.

   The file is uploaded and a fingerprint generated.

To configure a fingerprint document source

1. Go to *Security Profiles > Data Leak Prevention > Document Fingerprinting*.

2. In the Document Sources section, select *Create New*.

3. Configure the settings:

Name	Enter a descriptive name for the document source.
Server Type	This refers to the type of server share that is being accessed. The default is Windows Share but this will also work on Samba shares.
Server Address	Enter the IP address of the server.
User Name Password	Enter the user name and password of the account the FortiGate unit uses to access the server network share.
Path	Enter the path to the document folder.
Filename Pattern	You may enter a filename pattern to restrict fingerprinting to only those files that match the pattern. To fingerprint all files, enter an asterisk ("*").
Sensitivity Level	Select a sensitivity level. The sensitivity is a tag for your reference that is included in the log files. It does not change how fingerprinting works.
Scan Periodically	To have the files on the document source scanned on a regular basis, select this option. This is useful if files are added or changed regularly. Once selected, you can choose Daily, Weekly, or Monthly update options, and enter the time of day the files are fingerprinted.
Advanced	Expand the Advanced heading for additional options.
Fingerprint files in subdirectories	By default, only the files in the specified path are fingerprinted. Files in subdirectories are ignored. Select this option to fingerprint files in subdirectories of the specified path.
Remove fingerprints for deleted files	Select this option to retain the fingerprints of files deleted from the document source. If this option is disabled, fingerprints for deleted files will be removed when the document source is rescanned.
Keep previous fingerprints for modified files	Select this option to retain the fingerprints of previous revisions of updated files. If this option is disabled, fingerprints for previous version of files will be deleted when a new fingerprint is generated.

4. Select *OK*.

File filter

File filter is a DLP option that allows you to block files based on their file name or their type.

- **File patterns** are a means of filtering based purely on the names of files. They may include wildcards (*). For example, blocking *.scr will stop all files with an scr file extension, which is

commonly used for Windows screen saver files. Files trying to pass themselves off as Windows screen saver files by adopting the file-naming convention will also be stopped.

- Files can specify the full or partial file name, the full or partial file extension, or any combination. File pattern entries are not case sensitive. For example, adding `*.exe` to the file pattern list also blocks any files ending with `.EXE`.

- Files are compared to the enabled file patterns from top to bottom, in list order.

- In addition to the built-in patterns, you can specify more file patterns to block. For details, see .

- **File types** are a means of filtering based on an examination of the file contents, regardless of the file name. If you block the file type *Archive (zip)*, all zip archives are blocked even if they are renamed with a different file extension. The FortiGate examines the file contents to determine what type of file it is and then acts accordingly.

The FortiGate unit can take either of the following actions toward the files that match a configured file pattern or type:

- **Block**: the file is blocked and a replacement message is sent to the user. If both file pattern filtering and virus scan are enabled, the FortiGate unit blocks files that match the enabled file filter and does not scan these files for viruses.

- **Allow**: the file is allowed to pass.

The FortiGate unit also writes a message to the Security log and sends an alert email message if configured to do so.

> File filter does not detect files within archives. You can use file filter to block or allow the archives themselves, but not the contents of the archives.

General configuration steps

The following steps provide an overview of file filter configuration. For best results, follow the procedures in the order given. Also, note that if you perform any additional actions between procedures, your configuration may have different results.

1. Create a file filter list.
2. Create one or more file patterns or file types to populate the file filter list.
3. Enable the file filter list by adding it to a filter in a DLP sensor.
4. Select the DLP sensor in a security policy.

Creating a file filter list

Before your FortiGate unit can filter files by pattern or type, you must create a file filter list. The action triggered by the contents of a file filter list will be decided in the sensor so it is important to make sure that when you are building a list that you intend for the same action to be applied to all of the entries in the same list.

To create a file filter list

1. Go to *Security Profiles > Data Leak Prevention > File Filter*.
2. Select *Create New*.
3. Enter a *Name* for the new file filter list.
4. Select *OK*.

The new list is created and the edit file filter list window appears. The new list is empty. You need to populate it with one or more file patterns or file types.

Creating a file pattern

A file pattern allows you to block or allow files based on the file name. File patterns are created within file filter lists.

To create a file pattern

1. Go to *Security Profiles > Data Leak Prevention > File Filter*.
2. Select a file filter list.
3. Select the *Edit* icon.
4. Select *Create New*.
5. Select *File Name Pattern* as the *Filter Type*.
6. Enter the pattern in the *Pattern* field. The file pattern can be an exact file name or can include wildcards (*). The file pattern is limited to a maximum of 80 characters.
7. Select *OK*.

Creating a file type

A file type allows you to block or allow files based on the kind of file. File types are created within file filter lists.

To create a file type

1. Go to *Security Profiles > Data Leak Prevention > File Filter*.
2. Select the *Edit* icon of the file filter list to which you will add the file type.
3. Select *Create New*.
4. Select *File Type* as the *Filter Type*.
5. Select the kind of file from the *File Type* list.
6. Select *OK*.

DLP can detect the following file types:

- Archive (arj)
- Archive (bzip)
- Archive (bzip2)
- Archive (cab)
- Archive (gzip)
- Archive (Archive (lzh)
- Archive (rar)
- Archive (tar)
- Archive (zip)
- Audio (wav)
- Audio (wma)
- BMP (bmp)
- Batch File (bat)
- Common Console Document (msc)
- Encoded Data (base64)

- Encoded Data (binhex)
- Encoded Data (mime)
- Encoded Data (uue)
- Executable (elf)
- Executable (exe)
- GIF Image (gif)
- HTML Application (hta)
- HTML File (html)
- Ignored File Type (ignored)
- JPEG Image (jpeg)
- Java Application Descriptor (jad)
- Java Class File (class)
- Java Compiled Bytecode (cod)
- JavaScript File (javascript)
- Microsoft Office (msoffice)
- PDF (pdf)
- PNG Image (png)
- Packer (aspack)
- Packer (fsg)
- Packer (petite)
- Packer (upx)
- PalmOS Application (prc)
- Real Media Streaming (rm)
- Symbian Installer System File (sis)
- TIFF Image (tiff)
- Torrent (torrent)
- Unknown File Type (unknown)
- Video (mov)
- Video (mpeg)
- Windows Help File (hlp)
- activemime (activemime)

> The "unknown" type is any file type that is not listed in the table. The "ignored" type is the traffic the unit typically does not scan. This includes primarily streaming audio and video.

Preconfigured sensors

A number of preconfigured sensors are provided with your FortiGate unit. These can be edited or added to more closely match your needs.

Some of the preconfigured sensors with filters ready to go are:

- Credit-Card - This sensor logs the traffic, both files and messages, that contain credit card numbers in the formates used by American Express, MasterCard and Visa.
- Large-File - This sensor logs the traffic consisting of files larger than 5120 kB or approximately 5 MB.
- SSN-Sensor - This sensor logs the traffic, both files and messages, that contain Social Security Numbers with the exception of those that are WebEx invitation emails.

These rules affect only unencrypted traffic types. If you are using a FortiGate unit that can decrypt and examine encrypted traffic, you can enable those traffic types in these rules to extend their functionality if required.

Before using the rules, examine them closely to ensure you understand how they will affect the traffic on your network.

DLP archiving

DLP is typically used to prevent sensitive information from getting out of your company network, but it can also be used to record network use. This is called DLP archiving. The DLP engine examines email, FTP, IM, NNTP, and web traffic. Enabling archiving for rules when you add them to sensors directs the FortiGate unit to record all occurrences of these traffic types when they are detected by the sensor.

Since the archive setting is configured for each rule in a sensor, you can have a single sensor that archives only the things you want.

You can archive Email, FTP, HTTP, IM, and session control content:

- Email content includes IMAP, POP3, and SMTP sessions. Email content can also include email messages tagged as spam by Email filtering. If your unit supports SSL content scanning and inspection, Email content can also include IMAPS, POP3S, and SMTPS sessions.
- HTTP content includes HTTP sessions. If your unit supports SSL content scanning and inspection HTTP content can also include HTTPS sessions.
- IM content includes AIM, ICQ, MSN, and Yahoo! sessions.

DLP archiving comes in two forms: *Summary Only*, and *Full*.

Summary archiving records information about the supported traffic types. For example, when an email message is detected, the sender, recipient, message subject, and total size are recorded. When a user accesses the Web, every URL the user visits recorded. The result is a summary of all activity the sensor detected.

For more detailed records, full archiving is necessary. When an email message is detected, the message itself, including any attachments, is archived. When a user accesses the Web, every page the user visits is archived. Far more detailed than a summary, full DLP archives require more storage space and processing.

Because both types of DLP archiving require additional resources, DLP archives are saved to a FortiAnalyzer unit or the FortiGuard Analysis and Management Service (subscription required).

You can use DLP archiving to collect and view historical logs that have been archived to a FortiAnalyzer unit or the FortiGuard Analysis and Management Service. DLP archiving is available for FortiAnalyzer when you add a FortiAnalyzer unit to the Fortinet configuration. The FortiGuard Analysis server becomes available when you subscribe to the FortiGuard Analysis and Management Service.

Two sample DLP sensors are provided with DLP archiving capabilities enabled. If you select the `Content_Summary` sensor in a security policy, it will save a summary DLP archive of all traffic the security policy handles. Similarly, the `Content_Archive` sensor will save a full DLP archive of all traffic handled the security policy you apply it to. These two sensors are configured to detect all traffic of the supported types and archive them.

DLP archiving is set in the CLI only.

To set the archive to Full

```
config dlp sensor
  edit <name of sensor>
    set full-archive-proto smtp pop3 imap http ftp nntp aim icq msn
        yahoo mapi
    end
```

To set the archive to Summary Only

```
config dlp sensor
  edit <name of sensor>
    set summary-proto smtp pop3 imap http ftp nntp aim icq msn yahoo
        mapi
    end
```

DLP examples

Blocking content with credit card numbers

When the objective is to block credit card numbers one of the important things to remember is that 2 filters will need to be used in the sensor.

In the default Credit-Card sensor, you will notice a few things.

- The Action is set to Log Only
- In the Files filter not all of the services are being examined.

If you wish to block as much content as possible with credit card numbers in it instead of just logging most the traffic that has it, the existing sensor will have to be edited.

Security Profile > Data Leak Prevention > Sensor.

Use the drop down menu to select Credit-Card.

1. Edit the first filter.
   a. Change the Action to Block
   b. Make sure all of the services are being examined.
   c. Select OK
2. Repeat for the second filter
3. Select Apply
4. Edit the appropriate policies so that under *Security Profiles*, *DLP* is turned on and the *Credit-Card* sensor is selected.

Blocking emails larger than 15 MB and logging emails from 5 MB to 15 MB

Because the designated size is over 10 MB the proxy options cannot be used to block the file size. Multiple filters will have to be used in this case and the order that they are used is important. Because there is no mechanism to move the filters within the sensor the order that they are added to the sensor is important.

Security Profile > Data Leak Prevention > Sensor.

Create a new sensor

Use the following values

Name	large_emails
Comment	<optional>

Once the Sensor has been created a new filter will need to be added.

Create New

Use the following values

Filter:

* Choice between *Meassages* and *Files*: choose *Files.*
* Choose radio button to the left of *File Size*
* In the field for the file size type 15360
    1MB = 1024kB, 15 MB = 15 x 1024kB = 15360kB

Examine the following Services

SMTP	enabled
POP3	enabled
IMAP	enabled
HTTP	not enabled
FTP	not enabled
AIM	not enabled
ICQ	not enabled
MSN	not enabled
Yahoo!	not enabled
NNTP	not enabled
MAPI	not enabled

Action

* From the drop down menu choose Block

Select OK

A second filter will need to be added.

Create New

Use the following values

Filter:

- Choice between Meassages and Files: choose Files
- Choose radio button to the left of *File Size*
- In the field for the file size type 1024

Examine the following Services

SMTP	enabled
POP3	enabled
IMAP	enabled
HTTP	not enabled
FTP	not enabled
AIM	not enabled
ICQ	not enabled
MSN	not enabled
Yahoo!	not enabled
NNTP	not enabled
MAPI	not enabled

Action

- From the drop down menu choose *Log Only*.

Select OK

Select Apply

Add the sensor to the appropriate policy.

The reason that the block filter is placed first is because the filters are applied in sequence and once the traffic triggers a filter the action is applied and then the traffic is passed on to the next test. If the Log Only filter which checks for anything over 1MB is triggered this would include traffic over 15MB, so a 16 MB file would only be logged. In the described order, the 16 MB file will be blocked and the 3 MB file will be logged.

Selective blocking based on a finger print

The following is a fairly complex example but shows what can be done by combining various components in the correct configuration.

The company has a number of copyrighted documents that it does not want "escaping" to the Internet but it does want to be able to send those documents to the printers for turning into hardcopy.

The policies and procedures regarding this issue state that:

- Only members of the group *Senior_Editors* can send copyrighted material to the printers.
- Every member of the company by default is included in the group *employees*.
- Even permitted transmission of copyrighted material should be recorded.
- All of the printers IP addresses are in a group called *approved_printers*.
- There is a file share called copyrighted where any file that is copyrighted is required to have a copy stored.
- It doesn't happen often but for legal reasons sometimes these files can be changed, but all versions of a file in this directory need to be secured.
- All network connections to the Internet must have Antivirus enabled using at least the default profile.
- The SSL/SSH Inspection profile used will be *default*.

It is assumed for the purposes of this example that:

- Any addresses or address groups have been created.
- User accounts and groups have been created.
- The account used by the FortiGate is fgtaccess.
- The Copyrighted sensitivity level needs to be created.
- The copyrighted material is stored at \\192.168.27.50\books\copyrighted\

Sensitivity Level Addition

```
config dlp fp-sensitivity
   edit copyrighted
   end
```

Finger print configuration

Security Profile > Data Leak Prevention > Document Fingerprinting.

In the *Document Sources* section *select Create New*

Use the following field values

Name	copyrighted_material
Server Type	Windows Share
Server Address	192.168.27.50
User Name	fgtaccess
Password	******
Path	books/copyrighted/
Filename Pattern	*.pdf
Sensitivity	copyrighted
Scan Periodically	enabled
<Frequency>	Daily, Hour: 2, Min: 0
Advanced	

Fingerprint files in subdirectories	enabled
Remove fingerprints for deleted files	not enabled
Keep previous fingerprints for modified files	enabled

Create DLP Sensors

Security Profile > Data Leak Prevention > Sensor

Create a new sensor. This can be done one of two ways.

- In the menu bar at the top on the right hand, use the Create New icon (circle with + symbol inside).
- In the menu bar at the top on the right hand, use the View List icon to go to the list window and use the Create New icon on the top left of that page.

Two Sensors need to be created. One for blocking the transmission of copyrighted material and a second for allowing the passing of copyrighted material under specific circumstances.

Configuration for the first sensor that blocks transmission.

Use the following field values:

Name	block_copyrighted
Comment	<optional>

Once the Sensor has been created a new filter will need to be added.

Create New

Use the following values

Filter:

- Choice between Meassages and Files: choose Files
- Choose radio button to the left of File Finger Print
- From the drop down for File Finger Print choose "copyrighted"

Examine the following Services

SMTP	enabled
POP3	enabled
IMAP	enabled
HTTP	enabled
FTP	enabled
AIM	enabled
ICQ	enabled
MSN	enabled

Yahoo!	enabled
NNTP	enabled
MAPI	enabled

Action

- From the drop down menu choose Block

Configuration for the second sensor that allows transmission.

Use the following field values:

Name	allow_copyrighted
Comment	<optional>

Once the Sensor has been created a new filter will need to be added.

This will be identical to the filter in the block_copyrighted sensor except that the action will be *Log Only*.

Create policies and attach DLP sensors

Policy to allow transmission of copyrighted material

Policy > Policy > Policy

Create New

Use the following values in the Policy:

Policy Type	Firewall
Policy Subtype	User Identity
Incoming Interface	LAN
Source Address	all
Outgoing Interface	wan1
Enable NAT	enabled -- Use Destination Interface Address
Enable Web cache	<optional>
Enable WAN Optimization	<optional>
Skip this policy for unauthenticated user	do not enable
Disclaimer	<optional>
Customize authentication Messages	<optional>

Configure Authentication Rules:

Destination Address	approved_printers
Group(s)	Senior_Editors
User(s)	<optional>
Schedule	always
Service	ALL
Action	ACCEPT
Log Allowed Traffic	<optional>
Security Profiles	
Antivirus	<ON> default
Webfilter	<optional>
Application Control	<optional>
IPS	<optional>
Email Filter	<optional>
DLP Sensor	<ON> Copyrighted
VoIP	<optional>
ICAP	<optional>
Proxy Options	
SSL/SSH Inspection	<ON>
Traffic Shaping	<optional>

This policy should be place as close to the beginning of the list of policies so the it is among the first tested against.

Policy to block transmission of copyrighted material

This will in effect be the default template for all following policies in that they will have to use the DLP profile that blocks the transmission of the copyrighted material.

Policy > Policy > Policy

Create New or *Edit* the existing policies.

The fields should include what ever values you need to accomplish your requirements are but each policy should include the DLP sensor block_copyrighted or if a different DLP configuration is required it should include a filter that blocks *copyrighted* fingerprinted file.

If you need to create a policy that is identity based make sure that there is an Authentication rule for the group *employees* that uses the DLP sensor that blocks copyrighted material.

Application control

Using the application control Security Profile feature, your FortiGate unit can detect and take action against network traffic depending on the application generating the traffic. Based on FortiGate Intrusion Protection protocol decoders, application control is a user-friendly and powerful way to use Intrusion Protection features to log and manage the behavior of application traffic passing through the FortiGate unit. Application control uses IPS protocol decoders that can analyze network traffic to detect application traffic even if the traffic uses non-standard ports or protocols.

The FortiGate unit can recognize the network traffic generated by a large number of applications. You can create application control sensors that specify the action to take with the traffic of the applications you need to manage and the network on which they are active, and then add application control sensors to the firewall policies that control the network traffic you need to monitor.

Fortinet is constantly increasing the list of applications that application control can detect by adding applications to the FortiGuard Application Control Database. Because intrusion protection protocol decoders are used for application control, the application control database is part of the FortiGuard Intrusion Protection System Database and both of these databases have the same version number.

You can find the version of the application control database that is installed on your unit, by going to the *License Information* dashboard widget and find IPS Definitions version.

You can go to the FortiGuard Application Control List to see the complete list of applications supported by FortiGuard. This web page lists all of the supported applications. You can select any application name to see details about the application.

If you enable virtual domains (VDOMs) on the Fortinet unit, you need to configure application control separately for each virtual domain.

The following topics are included in this section:

- Application control concepts
- Application considerations
- Application traffic shaping
- Application control monitor
- Enable application control
- Application control examples

Application control concepts

You can control network traffic generally by the source or destination address, or by the port, the quantity or similar attributes of the traffic itself in the security policy. If you want to control the flow of traffic from a specific application, these methods may not be sufficient to precisely define the traffic. To address this problem, the application control feature examines the traffic itself for signatures unique to the application generating it. Application control does not require knowledge of any server addresses or ports. The FortiGate unit includes signatures for over 1000 applications, services, and protocols.

Updated and new application signatures are delivered to your FortiGate unit as part of your FortiGuard Application Control Service subscription. Fortinet is constantly increasing the

number of applications that application control can detect by adding applications to the FortiGuard Application Control Database. Because intrusion protection protocol decoders are used for application control, the application control database is part of the FortiGuard Intrusion Protection System Database and both of these databases have the same version number.

To view the version of the application control database installed on your FortiGate unit, go to the *License Information* dashboard widget and find the *IPS Definitions* version.

To see the complete list of applications supported by FortiGuard Application Control go to the FortiGuard Application Control List. This web page lists all of the supported applications. You can select any application name to see details about the application.

Application considerations

Some applications behave differently from most others. You should be aware of these differences before using application control to regulate their use.

Automatically allowing basic applications

A common practice is to block applications by category, because the alternative is to list each specific traffic on an individual basis. While listing the applications individually gives a great deal of granularity it does tend to allow for missing some of them. On the other hand, blocking by category has the drawback of blocking some traffic that was not intended to be blocked.

There are a number of basic applications that you may want to be allowed on a default basis. For example, DNS. If you were to block the category Network Services you would end up blocking your web browsing, unless your users are members of a very limited group that do their web browsing by using IP addresses instead of URLs. Without DNS the systems will not be able to resolve URLs into IP addresses.

Using a set of options in the CLI the FortiGate unit can be configured to automatically allow the following types of traffic, regardless of whether or not their category is blocked:

• DNS
• ICMP
• Generic HTTP Web browsing
• Generic SSL communications

Syntax

```
config application list
  edit appcontrol
    set options allow-dns allow-icmp allow-http allow-ssl
  end
```

As the example indicates, DNS is vitally important to multiple other types of traffic so by default it is set to be allowed, however the other settings must be specifically enabled.

IM applications

The Application Control function for a number of IM application is not in the Web Based Manager, in the CLI of the FortiGate unit. These applications are:

- AIM
- ICQ
- MSN
- Yahoo

These applications are controlled by either permitting or denying the users from logging in to the service. Individual IM accounts are configured as to whether or not they are permitted and then there is a global policy for how to action unknown users, by the application, and whether to add the user to the black list or the white list.

The configuration details for these settings can be found in the CLI Reference guide under the heading of imp2p.

Skype

Based on the NAT firewall type, Skype takes advantage of several NAT firewall traversal methods, such as STUN (Simple Traversal of UDP through NAT), ICE (Interactive Connectivity Establishment) and TURN (Traversal Using Relay NAT), to make the connection.

The Skype client may try to log in with either UDP or TCP, on different ports, especially well-known service ports, such as HTTP (80) and HTTPS (443), because these ports are normally allowed in firewall settings. A client who has previously logged in successfully could start with the known good approach, then fall back on another approach if the known one fails.

The Skype client could also employ Connection Relay. This means if a reachable host is already connected to the Skype network, other clients can connect through this host. This makes any connected host not only a client but also a relay server.

Application traffic shaping

You can apply traffic shaping for application list entries you configure to pass. Traffic shaping enables you to limit or guarantee the bandwidth available to the application or applications specified in an application list entry. You can also prioritize traffic by using traffic shaping.

You can create or edit traffic shapers by going to *Firewall Objects > Traffic Shaper > Shared*. Per-IP traffic shapers are not available for use in application traffic shaping.

Direction of traffic shaping

When Traffic Shaping is enabled the direction that traffic shaping will be applied must also be chosen.

Forward direction traffic shaping refers to the direction of the initial connection. This would be the direction described by the policy that the Application Control Sensor is assigned to. If the policy has an Incoming Interface of LAN and an Outgoing Interface of wan1 then any Forward Direction Traffic Shaping profile will apply to network traffic heading in that direction only. If the connection used by that policy involved a response that included a download of Gigabytes of traffic the shaper would not be applied to that traffic.

Reverse Direction Traffic Shaping is applied to traffic that is flowing in the opposite direction indicated by the direction of the policy. If the policy has an Incoming Interface of LAN and an

Outgoing Interface of wan1 then the shaper would only be applied to the traffic that was coming from the wan1 interface to the LAN interface.

For example, if you find that your network bandwidth is being overwhelmed by streaming HTTP video, one solution is to limit the bandwidth by applying a traffic shaper to an application control entry that allows the HTTP.Video application. Your users access the Web using a security policy that allows HTTP traffic from the internal interface to the external interface. Firewall policies are required to initiate communication so even though web sites respond to requests, a policy to allow traffic from the external interface to the internal interface is not required for your users to access the Web. The internal to external policy allows them to open communication sessions to web servers, and the external servers can reply using the existing session.

If you enable *Traffic Shaping* and select the Forward Direction shaper in an application sensor specified in the security policy, the problem will continue. The reason is the shaper you select for *Traffic Shaping* is applied only to the application traffic moving in the direction stated in the security policy. In this case, that is from the internal interface to the external interface. The security policy allows the user to visit the web site and start the video, but the video itself is streamed from the server to the user, or from the external interface to the internal interface. This is the reverse of the direction specified in the security policy. To solve the problem, you must enable *Reverse Direction Traffic Shaping* and select the appropriate shaper.

Shaper re-use

Shapers are created independently of firewall policies and application sensors so you are free to reuse the same shapers in multiple list entries and policies. Shared shapers can be configured to apply separately to each security policy or across all policies. This means that if a shaper is configured to guaranteed 1000 KB/s bandwidth, each security policy using the shaper will have its own 1000 KB/s reserved, or all of the policies using the shaper will share a pool if 1000 KB/s, depending on how it is configured.

The same thing happens when a shaper is used in application sensors. If an application sensor using a shaper is applied to two separate policies, how the bandwidth is limited or guaranteed depends on whether the shaper is set to apply separately to each policy or across all policies. In fact, if a shaper is applied directly to one security policy, and it is also included in an application sensor that is applied to another security policy, the same issue occurs. How the bandwidth is limited or guaranteed depends on the shaper configuration.

If a shaper is used more than once within a single application sensor, all of the applications using the shaper are restricted to the maximum bandwidth or share the same guaranteed bandwidth.

For example, you want to limit the bandwidth used by Skype and Facebook chat to no more than 100 KB/s. Create a shaper, enable *Maximum Bandwidth*, and enter 100. Then create an application sensor with an entry for Skype and another entry for Facebook chat. Apply the shaper to each entry and select the application sensor in the security policy that allows your users to access both services.

This configuration uses the same shaper for each entry, so Skype **and** Facebook chat traffic are limited to no more than 100 KB/s in total. That is, traffic from both applications is added and the total is limited to 100 KB/s. If you want to limit Skype traffic to 100 KB/s and Facebook chat traffic to 100 KB/s, you must use separate shapers for each application control entry.

Application control monitor

The application monitor enables you to gain an insight into the applications generating traffic on your network. When monitor is enabled in an application sensor entry and the list is selected in a security policy, all the detected traffic required to populate the selected charts is logged to the

SQL database on the FortiGate unit hard drive. The charts are available for display in the executive summary section of the log and report menu.

Because the application monitor relies on a SQL database, the feature is available only on FortiGate units with an internal hard drive.

While the monitor charts are similar to the top application usage dashboard widget, it offers several advantages. The widget data is stored in memory so when you restart the FortiGate unit, the data is cleared. Application monitor data is stored on the hard drive and restarting the system does not affect old monitor data.

Application monitor allows you to choose to compile data for any or all of three charts: top ten applications by bandwidth use, top ten media users by bandwidth, and top ten P2P users by bandwidth. Further, there is a chart of each type for the traffic handled by each security policy with application monitor enabled. The top application usage dashboard widget shows only the bandwidth used by the top applications since the last system restart.

Application Control monitor

Once you have configured application control and associated the sensors with firewall policies, you can monitor the results. The applications that will be reported on the ones that are included in sensors that are assigned to firewall policies.

Security Profile > Monitor > Application Monitor.

Here you will find some widgets that include charts:

- Top Applications by Bandwidth
- Top Applications by Session Count
- Top IP/User for <application>

The number of "Top" can be set to the value of 5, 10 or 15 on any of these widgets.

Enable application control

Application control examines your network traffic for traffic generated by the applications you want it to control.

General configuration steps

Follow the configuration procedures in the order given. Also, note that if you perform any additional actions between procedures, your configuration may have different results.

1. Create an application sensor.
2. Configure the sensor to include the signatures for the application traffic you want the FortiGate unit to detect. Configure each entry to allow or pass the traffic.
3. Enable application control in a security policy and select the application sensor.

Creating an application sensor

You need to create an application sensor before you can enable application control.

To create an application sensor

1. Go to *Security Profiles > Application Control > Application Sensor*.

2. Select the *Create New* icon in the title bar of the *Edit Application Sensor* window.

3. In the *Name* field, enter the name of the new application sensor.

4. Optionally, you may also enter a comment.

5. Select *OK*.

The application sensor is created and the sensor configuration window appears. A newly created application sensor is empty. Without applications, the application sensor will have no effect.

Adding applications to an application sensor

Once you have created an application sensor, you need to need to define the applications that you want to control.

You can add applications using application entries and application filters. Entries allow you to choose individual applications. Filters allow you to choose application attributes and all the applications with matching attributes are included in the filter.

The sequence of the entries in the table is significant. The entries are checked against the traffic in sequence, from top to bottom. If a match is found and the action is *Block* or *Reset*, the action is performed and further checking is stopped. If the action is *Monitor* the traffic is checked against all of the signatures in the sensor and the best match to the signature is the one that is logged.

To add an application entry to an application sensor

1. Go to *Security Profiles > Application Control > Application Sensors*.

2. Select an application sensor from the drop-down list in the Edit Application Sensor window title bar.

3. Select the *Create New* icon in the sensor area and this will bring up a new window entitled *New Application Filter*.

4. Choose the format of the filter. There are two types of entries that can be added to a sensor. The type of entry is detemined by the selection of the sensor type. The choices are either *Filter Based* or *Specify Applications*.

   • Filter Based

   This option is for choosing groups of similar applications based on the filters of Category, Popularity, Technology and Risk. Once the parameters of the 4 filter types have been chosen every application that falls in to that filtered list will be included in the list that the Application Control engine will use to filter the network traffic.

   • Specify Applications

   This option is good for a more granular approach to picking appliction to be filtered. It will allow for the use of the same filters that were used in the *Filter Based* option to develop a list of applications to be filtered but the Specify Applications option can be selective of which applications in that list are actually filtered. They are selected individually.

The difference in the Web-based Manager, when alternating between the Sensor Types, is that when the Filter Based option is chosen the Filter Options section will appear by default. If the Specify Applications sensor type is chosen you can click on the [Filter Options] link to make it appear and use it to narrow down the list of possible applications but it does not show up by default. The other difference is that with the Specify Applications option you are

given an additional field at the top of the Application List that allow you to type out the name of an application to search for it in that manner.

To use the search field, located above the application list, start typing any portion of the application name. The mail list of application will adjust accordingly.

5. Narrow down the list of applications to be filtered. This will depend a little on which *Sensor Type* was chosen. If the Filter Based option was chosen, by default, the top section of the window will show the properties by which the list of application filters can be filtered into a more managable list. These properties are broken into 4 sections representing the properties of Category, Popularity, Technology and Risk. Between the property filter section and the Action section of the window there is a listing of the individual application filters that have been configured into the appliance.

Each of these individual application filters is assigned values in each of the 4 properties. The values that can be assigned to these properties are listed in the 4 sections. By enabling the check boxes next to the propterties in the sections the list can be narrowed down until it only includes the subset of the individual application filters that you wish to make up the sensor entry or Application Filter.

When choosing a property, if the specific value is unknown do not disable the property section as this will cause the list of individual application filters to be empty.

The properties have been broken down into the following sections:

   a. Category

   These are the types of application that are available to filter by:

Table 10: Property Values listed in Category section along with ID#

Category Name	Category ID#
Botnet	19
eMail	21
File.Sharing	24
Game	8
General.Interest	12
IM	1
Media	5
Network.Service	15
P2P	2
Proxy	6
Remote.Access	7
Social.Networking	23
Storage.Backup	22
Update	17
VoIP	3
Web.Surfing	25

there is also a category designation reserved for future use.

These categories should cover the bulk of application based network traffic. If you wanted to disallow the use of Peer to Peer (P2P) applications because you didn't want your users tieing up your bandwidth with torrent downloads you would select the P2P category and set the Action to Block

b. Popularity

Popularity is broken down into 5 levels of popularity represented by stars. 5 stars representing the most popular applications and 1 star representing applications that are the least popular. The Popularity property works well when trying to narrow down the list of one of the categories. Using the previous category example of P2P traffic but you wanted to monitor the activity of the most popular applications, which numbers about 30 as opposed to over 100, you would choose P2P from Category and the 5 star popularity.

c. Technology

Technology is broken down into 3 technology models as well as the more basic Network-Protocol which would can be used as a catch all for anything not covered by the more narrowly defined technologies of:

- Browser-Based
- Client-Server
- Peer -to-Peer

d. Risk

The Risk property does not indicate the level of risk but the type of impact that is likely to occur by allowing the traffic from that application to occur. The Risk list is broken down into the following

- Botnet
- Excessive-Bandwidth
- None

6. Pick the individual applications if using the Specify Applications Sensor type.

From the list of possible applications highlight the application by selecting the application.If you choose an application in error you can unhighlight or desellect the application by clicking on it again.

If the Filter Based sensor type is being used this will not be an option.

7. Select the *Action* the FortiGate unit will take when it detects network traffic from the application:

- *Monitor* allows the application traffic to flow normally and log all occurrences.

  If you set the action to *Monitor*, you have the option of enabling traffic shaping for the application or applications specified in this application list entry. For more information about application control traffic shaping, see

- *Block* will stop all traffic from the application and log all occurrences.

- *Reset* will reset the network connection on the session that the specified application traffic was detected on.

- *Traffic Shaping* will allow a Traffic Shaping profile to be applied to the applicatin traffic that triggered the sensor.

  Choosing the Traffic shaping action will cause to appear the secondary options of:

  - Forward Direction Traffic Shaping with a checkbox
  - Reverse Direction Traffic Shaping with a checkbox

  If the checkbox is enable for these options a dropdown menu will appear next to that option that will allow you to choose one of the existing Traffic Shaping profiles. If you are

going to want to use Traffic Shaping as an action in Application Control it is best to set up any of the Traffic Shaping profiles that you will want in advance.

Viewing and searching the application list

Go to *Security Profiles > Application Control > Application List* to view the list of applications the FortiGate unit recognizes. You may find applications by paging manually through the list, apply filters, or by using the search field.

Searching manually

Applications are displayed in a paged list, with 50 applications per page. The bottom of the screen shows the current page and the total number of pages. You can enter a page number and press enter, to skip directly to that page. Previous Page and Next Page buttons move you through the list, one page at a time. The First Page and Last Page button take you to the beginning or end of the list.

Applying application list filters

You can enter criteria for one or more columns, and only the applications matching all the conditions you specify will be listed.

To apply filters

1. Go to *Security Profiles > Application Control > Application List*.
2. Goto the column you intend to filter by and select the filter icon in the heading to the left of the column name.
3. A small window will appear which will have a field for the value to intend to filter by and a checkbox for *NOT* so that you can choose to view all of the values except the one you enter into the field. You can also input multiple values if you separate them using commas.
4. Select *Apply*.
5. Continue to add more filters to narrow your search, if required.
6. Select *OK*.

Creating a New Custom Application Signature

If you have to deal with an application that is not already in the *Application List* you have the option to create a new one.

Creating a new Application Filter

1. Go to *Security Profiles > Application Control > Application List*.
2. Select *Create New.*
3. The *New Custom Application Signature* window will appear.
4. In the *Name* field give the new signature a unique name.
5. In the *Comments* field give a brief discription of the application or what you intend to filter by.
6. In the *Signature* field include the signature that you intend to base your application filtering on. For more details on how to design a signature see "Creating a custom IPS signature" on page 74
7. As an optional step you can select the Submit Signature link to submit your newly created signature to Fortiguard for possible inclusion in future predefined application lists.
8.

Enabling application traffic shaping

Enabling traffic shaping in an application sensor involves selecting the required shaper. You can create or edit shapers in *Firewall Objects > Traffic Shaper > Shared*.

To enable traffic shaping

1. Go to *Security Profiles > Application Control > Application Sensors*.

2. Select an application sensor from the drop-down list in the Edit Application Sensor window title bar.

3. Select the application control list entry and choose *Edit*.

4. Select *Traffic Shaping* and choose the required traffic shaper from the list.

   If the action is set to *Block*, the traffic shaping option is not available. Only allowed traffic can be shaped.

5. Select *Reverse Direction Traffic Shaping* and choose the required traffic shaper from the list if traffic flowing in the opposite direction also requires shaping.

6. Select *OK*.

Any security policy with this application sensor selected will shape application traffic according to the applications specified in the list entry and the shaper configuration.

Application control examples

Blocking all instant messaging

Instant messaging use is not permitted at the Example Corporation. Application control helps enforce this policy.

First you will create an application sensor with a single entry that includes all instant messaging applications. You will set the list action to block.

To create the application sensor

1. Go to *Security Profiles > Application Control > Application Sensors*.

2. Select the *Create New* icon in the title bar of the *Edit Application Sensor* window.

3. In the *Name* field, enter `no_IM` for the application sensor name.

4. Select *OK*.

5. Select the *Create New* icon in the sensor.

6. For the Sensor Type select *Filter Based*.

7. For *Category*, select only *IM*.

8. For *Popularity*, *Technology* and *Risk*, make sure that all of the options are selected.

9. For *Action*, select *Block*.

10. Select *OK* to save the new filter.

11. Select *Apply* to save the sensor.

Next you will assign the sensor to a policy.

To enable application control and select the application sensor

1. Go to *Policy > Policy > Policy*.

2. Select the security policy that allows the network users to access the Internet and choose *Edit*.

3. Under the heading *Security Profiles* toggle the button next to *Application Control* to turn it on.

4. In the drop down menu field next to the *Application Control* select the *no_IM* application sensor.

5. Select *OK*.

No IM use will be allowed by the security policy. If other firewall policies handle traffic that users could use for IM, enable application control with the *no IM* application sensor for those as well.

Allowing only software updates

Some departments at Example Corporation do not require access to the Internet to perform their duties. Management therefore decided to block their Internet access. Software updates quickly became an issue because automatic updates will not function without Internet access and manual application of updates is time-consuming.

The solution is configuring application control to allow only automatic software updates to access the Internet.

To create an application sensor — web-based manager

1. Go to *Security Profiles > Application Control > Application Sensors*.

2. Select the *Create New* icon in the title bar of the *Edit Application Sensor* window.

3. In the *Name* field, enter `Updates_Only` as the application sensor name.

4. Select *OK*.

5. Select the *Create New* icon in the sensor.

6. For the Sensor Type select *Filter Based*.

7. Enable only *Update* in the *Category* list.

8. Select *Monitor* from the *Action* list.

9. Select *OK* to save the filter to the sensor.

The filter just finished filter will allow all software update application traffic.

10. Select the application filter *All Other Known Applications*.

11. Select *Edit*.

12. Select *Block* from the *Action* list.

13. Select *OK*.

The filter just finished filter will block all traffic from recognized applications that are not specified in this application sensor.

14. Select the *All Other Unknown Applications* entry.

15. Select *Edit*.

16. Select *Block* from the *Action* list.

17. Select *OK*.

The filter just finished filter will block all traffic from applications that are not recognized by the application control feature.

18. Select *Apply* to save the application sensor.

To create an application sensor — CLI

```
config application list
  edit Updates_Only
    config entries
      edit 1
        set category 17
        set action pass
      end
    set other-application-action block
    set unknown-application-action block
  end
```

You will notice that there are some differences in the naming convention between the Web Based Interface and the CLI. For instance the *Action* in the CLI is "`pass`" and the *Action* in the Web Based Manager is "*Monitor".*

Selecting the application sensor in a security policy

An application sensor directs the FortiGate unit to scan network traffic only when it is selected in a security policy. When an application sensor is selected in a security policy, its settings are applied to all the traffic the security policy handles.

To select the application sensor in a security policy — web-based manager

1. Go to *Policy > Policy > Policy*.
2. Select a policy.
3. Select the *Edit* icon.
4. Under the heading *Security Profiles* toggle the button next to *Application Control* to turn it on.
5. In the drop down menu field next to the *Application Control* select the `Updates_only` list.
6. Select *OK*.

To select the application sensor in a security policy — CLI

```
config firewall policy
  edit 1
    set utm-status enable
    set profile-protocol-options default
    set application-list Updates_Only
  end
```

Traffic handled by the security policy you modified will be scanned for application traffic. Software updates are permitted and all other application traffic is blocked.

ICAP

ICAP is the acronym for Internet Content Adaptation Protocol The purpose of the feature is to off load work that would normally take place on the firewall to a separate server specifically set up for the specialized processing of the incoming traffic. This takes some of the resource strain off of the FortiGate firewall leaving it to concentrate its resources on things that only it can do.

Off-loading value-added services from Web servers to ICAP servers allows those same web servers to be scaled according to raw HTTP throughput versus having to handle these extra tasks.

ICAP servers are focused on a specific function, for example:

- Ad insertion
- Virus scanning
- Content translation
- HTTP header or URL manipulation
- Language translation
- Content filtering

> ICAP does not appear by default in the web-based manager. You must enable it in *System > Admin > Settings* to display ICAP in the web-based manager.

The following topics are included in this section:

- The Protocol
- Offloading using ICAP
- Configuration Settings
- Example ICAP sequence
- Example Scenerio

The Protocol

The protocol is a lightweight member of the TCP/IP suite of protocols. It is an Application layer protocol and its specifications are set out in RFC 3507. The default TCP that is assigned to it is 1344. Its purpose is to support HTTP content adaptation by providing simple object-based content vectoring for HTTP services. ICAP is usually used to implement virus scanning and content filters in transparent HTTP proxy caches. Content Adaptation refers to performing the particular value added service, or content manipulation, for an associated client request/response.

Essentially it allows an ICAP client, in this case the FortiGate firewall, to pass HTTP messages to an ICAP server like a remote procedure call for the purposes of some sort of transformation or other processing adaptation. Once the ICAP server has finished processing the the content, the modified content is sent back to the client.

The messages going back and forth between the client and server are typically HTTP requests or HTTP responses. While ICAP is a request/response protocol similar in semantics and usage

to HTTP/1.1 it is not HTTP nor does it run over HTTP, as such it cannot be treated as if it were HTTP. For instance ICAP messages can not be forwarded by HTTP surrogates.

Offloading using ICAP

If you enable ICAP in a security policy, HTTP traffic intercepted by the policy is transferred to an ICAP server in the ICAP profile added to the policy. Responses from the ICAP server are returned to the FortiGate unit which forwards them to an HTTP client or server.

You can offload HTTP responses or HTTP requests (or both) to the same or different ICAP servers.

If the FortiGate unit supports HTTPS inspection, HTTPS traffic intercepted by a policy that includes an ICAP profile is also offloaded to the ICAP server in the same way as HTTP traffic.

When configuring ICAP on the FortiGate unit, you must configure an ICAP profile that contains the ICAP server information; this profile is then applied to a security policy.

Configuration Settings

There are 2 sections where ICAP is configured:

Servers

The available settings to be configured regarding the profile are

IP Type (in the GUI) or IP address version (in the CLI)

The options for this field in the GUI are 2 radio buttons labelled "IPv4" and "IPv4". In the CLI the approach is slightly different. There is a field "ip-version" that can be set to "4" or "6".

IP address

depending on whether you've set the IP version to 4 or 6 will determine the format that the content of this field will be set into. In the GUI it looks like the same field with a different format but in the CLI it is actually 2 different fields named "ip-address" and ip6-address.

Maximum Connections

This value refers to the maximum number of concurrent connections that can be made to the ICAP server. The default setting is 100. This setting can only be configured in the CLI.

The syntax is:

```
config icap server
  edit <icap_server_name>
    set max-connections <integer>
    end
```

Port

this is the TCP port used for the ICAP traffic. The range can be from 1 to 65535. The default value is 1344.

Enable Request Processing

Enabling this setting allows the ICAP server to process request messages.

If enabled this setting will also require:

* *Server* - This is the name of the ICAP server. It is chosen from the drop down menu in the field. The servers are configure in the Security Profiles > ICAP > Server section.
* *Path* - This is the path on the server to the processing compent. For instance if the Windows share name was "Processes" and the directory within the share was "Content-Filter" the path would be "/Processes/Content-Filter/"
* *On Failure* - There are 2 options. You can choose by the use of radio buttons either *Error* or *Bypass*.

Enable Response Processing

Enabling this setting allows the ICAP server to process response messages.

If enabled this setting will also require:

* *Server* - This is the name of the ICAP server. It is chosen from the drop down menu in the field. The servers are configure in the Security Profiles > ICAP > Server section.
* *Path* - This is the path on the server to the processing compent. For instance if the Windows share name was "Processes" and the directory within the share was "Content-Filter" the path would be "/Processes/Content-Filter/"

On Failure - There are 2 options. You can choose by the use of radio buttons either *Error* or *Bypass*.

Enable Streaming Media Bypass

Enabling this setting allows streaming media to ignore offloading to the ICAP server.

Example ICAP sequence

This example is for an ICAP server performing web URL filtering on HTTP requests

1. A user opens a web browser and sends an HTTP request to connect to a web server.
2. The FortiGate unit intercepts the HTTP request and forwards it to an ICAP server.
3. The ICAP server receives the request and determines if the request is for URL that should be blocked or allowed.
   * If the URL should be blocked the ICAP server sends a response to the FortiGate unit. The FortiGate unit returns this response to the user's web browser. This response could be a message informing the user that their request was blocked.
   * If the URL should be allowed the ICAP server sends a request to the FortiGate unit. The FortiGate unit forwards the request to the web server that the user originally attempted to connect to.
   * When configuring ICAP on the FortiGate unit, you must configure an ICAP profile that contains the ICAP server information; this profile is then applied to a security policy.

Example Scenerio

Information relavent to the following example:

- The ICAP server is designed to do proprietary content filtering specific to the organization so it will have to receive the messages and sent back appropriate responses.
- The content filter is a required security precaution so it if the message cannot be processed it is not allowed through.
- Resources on both the Fortigate and the ICAP server are considerable so the maximum connections setting will set at a double the default value to analyse the impact on performance.
- The ICAP server's IP address is 172.16.100. 55.
- The path to the processing component is "/proprietary_code/content-filter/".
- Streaming media is not something that the filter considers, but is allowed through the policy so processing it would be a waste of resources.
- The ICAP profile is to be added to an existing firewall policy.
- It is assumed that the display of the policies has already been configured to show the column "ID".

1. Enter the following to configure the ICAP server:

Go to *Security Profiles* > *ICAP* > *Server*.

Use the following values:

Name	content-filtration-server4
IP Type	4
IP Address	172.16.100.55
Port	1344

Use the CLI to set the max-connections value.

```
config icap server
  edit content-filtration-server4
    set max-connections 200
    end
```

2. Enter the following to configure the ICAP profile to then apply to a security policy:

Use the following values:

Name	Prop-Content-Filtration
Enable Request Processing	enable
Server	content-filtration-server4
Path	/proprietary_code/content-filter/
On Failure	Error
Enable Response Processing	enable
Server	content-filtration-server4

Path	/proprietary_code/content-filter/
On Failure	Error
Enable Streaming Media Bypass	enable

3. Apply the ICAP profile to policy:

The purposes of this particular ICAP profile is to filter the content of the traffic coming through the firewall via policy ID#17

- **a.** Go to *Policy > Policy > Policy.*
- **b.** Open the existing policy ID# 17 for editing.
- **c.** Go to the section *Security Profiles.*
- **d.** Select the button next to *ICAP* so that it indicates that it's status is *ON*.
- **e.** Select the field with the profile name and use the drop down menu to select *Prop-Content-Filtration.*
- **f.** Select *OK*.

Other Security Profiles considerations

The following topics are included in this section:

- Profile Groups
- Security Profiles and Virtual domains (VDOMs)
- Conserve mode
- SSL content scanning and inspection
- Monitoring Security Profiles activity
- Using wildcards and Perl regular expressions
- Monitor interface reference

Profile Groups

One of the options when adding Security profiles to policies is the use of the Profile Groups feature. This works much the same way as an address group or a service group. You assign a selection of Security profiles to the Group and assign the group to a policy. This can be very convenient in an environment that has a large number of policies because instead of deciding each time you make a policy which Security profiles are going to be used you can have a small selection of Profile groups and every policy is assigned one of those groups. If changes need to be make, rather than going into each policy to make individual changes you only have to make changes to the group and the changes automatically propagate through to all of the policies that are using the Profile Group. It makes Security Profiles administration much simpler to implement, simpler to administrate and simpler to remember what Security Profiles features are being assigned to policies.

To refine the application of Security Profiles even further you can use the Profile Group in combination with Identity based policies and User Groups so that depending upon which User group a person belongs to that can be assigned a common set of Security profiles. A good example of this would a school environment. Staff and students are going to have significantly different permissions and restrictions associated with them. Staff will be allow access to websites that children are not (Web Filter). Staff will be allowed to transmit certain data under certain circumstances while students cannot transmit that type of data at all (DLP). Staff might have access to applications to communicate with colleagues in real time while students might be denied social networking access to get them from being distracted from their studies (Application Control). There are a number of permutations and possibilities made simpler and easier to administrate using these features together.

Creating a new group

Security profiles that can be grouped

When setting up a Profile Group you can assign to a group, or not as you want, the following Profile types:

- AntiVirus
- Web Filter
- Application Control
- IPS
- Email Filter
- DLP Sensor
- VoIP
- ICAP

Because the Security profiles need to use one, if you are assigning a Security profile to a policy you must assign a Proxy Option profile.

Using the Web-based Manager

To keep the interface simpler and less cluttered, by default, some versions of the firmware only display a default profile for each of the profile types and a default Profile Group. By going into the Admin Settings section and enabling the display of Multiple Security Profiles the option to have multiple Profile Groups in the Web Based Manager is also enabled.

1. Go to Security Profiles --> Profile Group --> Profile Group
2. Select Create New
3. Give the New Profile group a name.
4. Select the Security Profiles.
   a. Use the check-boxes to determine whether or not a particular Security profile will be assigned.
   b. Use the drop-down menu to determine which Security profile will be used.
   c. Select a Proxy Option profile.

   The Default Proxy Option Profile will be added by default if another profile is not selected.
5. Select OK.

Using the CLI

In the CLI enter the commands:

```
config firewall profile-group
   edit <profile_group_name>
      set profile-protocol-options <protocol_options_name>
      set av-profile <name_of_av-profile>
      set webfilter-profile <name_of_webfilter-profile>
      set spamfilter-profile <name_of_spamfilter-profile>
      set dlp-sensor <name_of_dlp-sensor>
      set ips-sensor <name_of_ips-sensor>
      set application-list <name_of_application-list>
      set voip-profile <name_of_voip-profile>
      set icap-profile <name_of_icap-profile>
      set deep-inspection-options <name_of_deep-inspection-options>
   next
end
```

Adding a Profile Group to a policy

Using the CLI

1. Go to the Firewall policy that you wish to associate the Profile Group

   a. For an Address Firewall policy:
   ```
   config firewall policy
      edit <policyID>
   ```

   b. For an Identity based policy
   ```
   config firewall policy
      edit <policyID>
         config identity-based-policy
            edit <policy_id>
   ```

2. To assign a Profile Group to a security policy the following additional settings need to be added to the policy configuration.
   ```
   set utm-status enable
   set profile-type group
   set profile-group <name of the profile group>
   end
   ```

When adding a Profile Group to a policy there are 2 potential points of confusion:

1. Depending on your interpretation, there may be some confusion on the profile-type setting.
   - `group` indicates the use of a profile group.
   - `single` indicates the use of individual Security profiles.

2. In the CLI, the context, or placement in the "syntax tree" of configuration settings, can make some options available or unavailable depending on other settings.

   In an Address Policy you only have to go down 2 "levels" to have the options for configuring the Profile Groups available.

   When an Identity policy is being used the Profile Group options are not available at the same level. You have to go down a further 2 levels, to inside the Authentication rule that is nested within the overall umbrella of the Firewall Policy. This is where the Profile Group settings will be available to you.

Security Profiles and Virtual domains (VDOMs)

If you enable virtual domains (VDOMs) on your FortiGate unit, all Security Profiles configuration is limited to the VDOM in which you configure it.

While configuration is not shared, the various databases used by Security Profiles features are shared. The FortiGuard antivirus and IPS databases and database updates are shared. The FortiGuard web filter and spam filter features contact the FortiGuard distribution network and access the same information when checking email for spam and web site categories and classification.

Conserve mode

FortiGate units perform all Security Profiles processing in physical RAM. Since each model has a limited amount of memory, conserve mode is activated when the remaining free memory is nearly exhausted or the AV proxy has reached the maximum number of sessions it can service. While conserve mode is active, the AV proxy does not accept new sessions.

The AV proxy

Most content inspection the FortiGate unit performs requires that the files, email messages, URLs, and web pages be buffered and examined as a whole. The AV proxy performs this function, and because it may be buffering many files at the same time, it uses a significant amount of memory. Conserve mode is designed to prevent all the component features of the FortiGate unit from trying to use more memory than it has. Because the AV proxy uses so much memory, conserve mode effectively disables it in most circumstances. As a result, the content inspection features that use the AV proxy are also disabled in conserve mode.

All of the Security Profiles features use the AV proxy with the exception of IPS, application control, DoS as well as flow-based antivirus, DLP, and web filter scanning. These features continue to operate normally when the FortiGate unit enters conserve mode.

Entering and exiting conserve mode

A FortiGate unit will enter conserve mode because it is nearly out of physical memory, or because the AV proxy has reached the maximum number of sessions it can service. The memory threshold that triggers conserve mode varies by model, but it is about 20% free memory. When memory use rises to the point where less than 20% of the physical memory is free, the FortiGate unit enters conserve mode.

The FortiGate unit will leave conserve mode only when the available physical memory exceeds about 30%. When exiting conserve mode, all new sessions configured to be scanned with features requiring the AV proxy will be scanned as normal, with the exception of a unit configured with the one-shot option.

Conserve mode effects

What happens when the FortiGate unit enters conserve mode depends on how you have `av-failopen` configured. There are four options:

off

The off setting forces the FortiGate unit to stop all traffic that is configured for content inspection by Security Profiles features that use the AV proxy. New sessions are not allowed but

current sessions continue to be processed normally unless they request more memory. Sessions requesting more memory are terminated.

For example, if a security policy is configured to use antivirus scanning, the traffic it permits is blocked while in conserve mode. A policy with IPS scanning enabled continues as normal. A policy with both IPS and antivirus scanning is blocked because antivirus scanning requires the AV proxy.

Use the off setting when security is more important than a loss of access while the problem is rectified.

pass

The pass setting allows traffic to bypass the AV proxy and continue to its destination. Since the traffic is bypassing the proxy, no Security Profiles scanning that requires the AV proxy is performed. Security Profiles scanning that does not require the AV proxy continues normally.

Use the pass setting when access is more important than security while the problem is rectified.

Pass is the default setting.

one-shot

The one-shot setting is similar to pass in that traffic is allowed when conserve mode is active. The difference is that a system configured for one-shot will force new sessions to bypass the AV proxy even after it leaves conserve mode. The FortiGate unit resumes use of the AV proxy only when the `av-failopen` setting is changed or the unit is restarted.

idledrop

The idledrop setting will recover memory and session space by terminating all the sessions associated with the host that has the most sessions open. The FortiGate may force this session termination a number of times, until enough memory is available to allow it to leave conserve mode.

The idledrop setting is primarily designed for situations in which malware may continue to open sessions until the AV proxy cannot accept more new sessions, triggering conserve mode. If your FortiGate unit is operating near capacity, this setting could cause the termination of valid sessions. Use this option with caution.

Configuring the av-failopen command

You can configure the av-failopen command using the CLI.

```
config system global
   set av-failopen {off | pass | one-shot | idledrop}
end
```

The default setting is pass.

SSL content scanning and inspection

If your FortiGate model supports SSL content scanning and inspection, you can apply antivirus scanning, web filtering, FortiGuard Web Filtering, and email filtering to encrypted traffic. You can

also apply DLP and DLP archiving to HTTPS, IMAPS, POP3S, and SMTPS traffic. To perform SSL content scanning and inspection, the FortiGate unit does the following:

- intercepts and decrypts HTTPS, IMAPS, POP3S, SMTPS, and FTPS sessions between clients and servers (FortiGate SSL acceleration speeds up decryption)

- applies content inspection to decrypted content, including:

- HTTPS, IMAPS, POP3S, and SMTPS Antivirus, DLP, and DLP archiving

- HTTPS web filtering and FortiGuard web filtering

- IMAPS, POP3S, and SMTPS email filtering

- encrypts the sessions and forwards them to their destinations.

Figure 13: FortiGate SSL content scanning and inspection packet flow
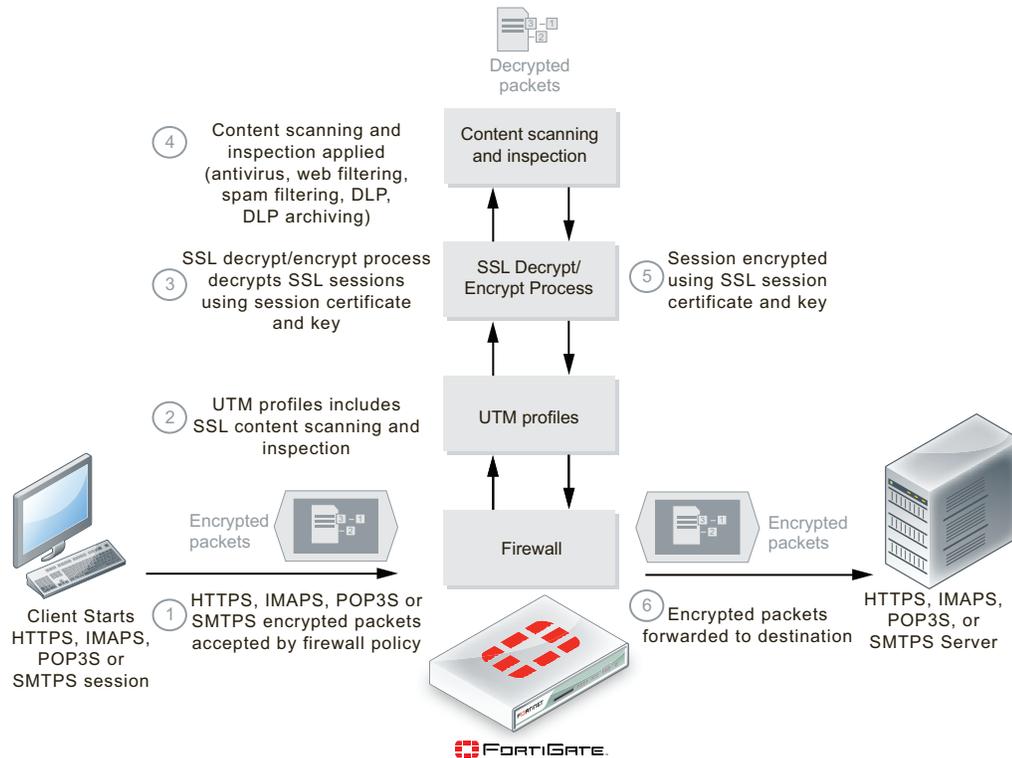
Setting up certificates to avoid client warnings

To use SSL content scanning and inspection, you need to set up and use a certificate that supports it. FortiGate SSL content scanning and inspection intercepts the SSL keys that are passed between clients and servers during SSL session handshakes and then substitutes spoofed keys. Two encrypted SSL sessions are set up, one between the client and the FortiGate unit, and a second one between the FortiGate unit and the server. Inside the FortiGate unit the packets are decrypted.

While the SSL sessions are being set up, the client and server communicate in clear text to exchange SSL session keys. The session keys are based on the client and server certificates. The FortiGate SSL decrypt/encrypt process intercepts these keys and uses a built-in signing CA certificate named Fortinet_CA_SSLProxy to create keys to send to the client and the server. This signing CA certificate is used only by the SSL decrypt/encrypt process. The SSL decrypt/encrypt process then sets up encrypted SSL sessions with the client and server and uses these keys to decrypt the SSL traffic to apply content scanning and inspection.

Some client programs (for example, web browsers) can detect this key replacement and will display a security warning message. The traffic is still encrypted and secure, but the security warning indicates that a key substitution has occurred.

You can stop these security warnings by importing the signing CA certificate used by the server into the FortiGate unit SSL content scanning and inspection configuration. Then the FortiGate unit creates keys that appear to come from the server and not the FortiGate unit.

You can add one signing CA certificate for SSL content scanning and inspection. The CA certificate key size must be 1024 or 2048 bits. 4096-bit keys are not supported for SSL content scanning and encryption.

You can replace the default signing CA certificate, Fortinet_CA_SSLProxy, with another signing CA certificate. To do this, you need the signing CA certificate file, the CA certificate key file, and the CA certificate password.

To add a signing CA certificate for SSL content scanning and inspection

1. Obtain a copy of the signing CA certificate file, the CA certificate key file, and the password for the CA certificate.

2. Go to *System > Certificates > Local Certificates* and select *Import*.

3. Set *Type* to *Certificate*.

4. For *Certificate file,* use the *Browse* button to select the signing CA certificate file.

5. For *Key file,* use the *Browse* button to select the CA certificate key file.

6. Enter the CA certificate *Password*.

7. Select *OK*.

   The CA certificate is added to the *Local Certificates* list. In this example the signing CA certificate name is Example_CA. This name comes from the certificate file and key file name. If you want the certificate to have a different name, change these file names.

8. Add the imported signing CA certificate to the SSL content scanning and inspection configuration. Use the following CLI command if the certificate name is Example_CA.

   ```
   config firewall ssl setting
      set caname Example_CA
   end
   ```

   The Example_CA signing CA certificate will now be used by SSL content scanning and inspection for establishing encrypted SSL sessions.

SSL content scanning and inspection settings

If SSL content scanning and inspection is available on your FortiGate unit, you can configure SSL settings. The following table provides an overview of the options available and where to find further instruction:

Table 11:SSL content scanning and inspection settings

Setting	Description
Predefined firewall services	The IMAPS, POP3S and SMTPS predefined services. You can select these services in a security policy and a DoS policy.
Protocol recognition	The TCP port numbers that the FortiGate unit inspects for HTTPS, IMAPS, POP3S, and SMTPS. Go to *Policy > Policy > Proxy Options*. Add or edit a Proxy Options profile, configure HTTPS, IMAPS, POP3S, SMTPS, and FTPS. Using *Proxy Options*, you can also configure the FortiGate unit to perform URL filtering of HTTPS or to use SSL content scanning and inspection to decrypt HTTPS so that the FortiGate unit can also apply antivirus and DLP content inspection and DLP archiving to HTTPS. Using SSL content scanning and inspection to decrypt HTTPS also allows you to apply more web filtering and FortiGuard Web Filtering options to HTTPS. To enable full SSL content scanning of web filtering, select *Enable Deep Scanning* under HTTPS in the Proxy Options profile.
Antivirus	Antivirus options including virus scanning and file filtering for HTTPS, IMAPS, POP3S, and SMTPS. Go to *AntiVirus > Profile*. Add or edit a profile and configure *Virus Scan* for HTTPS, IMAPS, POP3S, and SMTPS.
Antivirus quarantine	Antivirus quarantine options to quarantine files in HTTPS, IMAPS, POP3S, SMTPS, and FTPS sessions. Go to *Security Profiles > AntiVirus > Quarantine*. You can quarantine infected files, suspicious files, and blocked files found in HTTPS, IMAPS, POP3S, SMTPS, and FTPS sessions.
Web filtering	Web filtering options for HTTPS: • Web Content Filter • Web URL Filter • ActiveX Filter • Cookie Filter • Java Applet Filter • Web Resume Download Block • Block invalid URLs Go to *Security Profiles > Web Filter > Profile*. Add or edit a web filter profile and configure web filtering for HTTPS.

Table 11:SSL content scanning and inspection settings (continued)

Setting	Description
FortiGuard Web Filtering	FortiGuard Web Filtering options for HTTPS: • Enable FortiGuard Web Filtering • Enable FortiGuard Web Filtering Overrides • Provide Details for Blocked HTTP 4xx and 5xx Errors • Rate Images by URL (Blocked images will be replaced with blanks) • Allow Websites When a Rating Error Occurs • Strict Blocking • Rate URLs by Domain and IP Address • Block HTTP Redirects by Rating Go to *Security Profiles > Web Filter > Profile*. Add or edit a profile and configure FortiGuard Web Filtering for HTTPS.
Email filtering	Email filtering options for IMAPS, POP3S, and SMTPS: • FortiGuard Email Filtering IP Address Check, URL check, E-mail Checksum Check, and Spam Submission • IP Address BWL Check • E-mail Address BWL Check • Return S-mail DNS Check • Banned Word Check • Spam Action • Tag Location • Tag Format Go to *Security Profiles > Email Filter > Profile*. Add or edit a profile and configure email filtering for IMAPS, POP3S, and SMTPS.
Data Leak Prevention	DLP for HTTPS, IMAPS, POP3S, and SMTPS. To apply DLP, follow the steps below: • Go to *Security Profiles > Data Leak Prevention > Sensor,* create a new DLP sensor or edit an existing one and then add any combination of the DLP advanced rules, DLP compound rules, file filters, a Regular Expressions, and file size limits to a DLP sensor. • Go to *Policy > Policy > Proxy Options*. Add or edit a profile and select *Enable Deep Scan* under HTTPS. • Go to *Policy > Policy > Policy*, edit the required policy, enable *DLP Sensor* and select the DLP sensor. • Go to *Policy > Policy > Policy*, edit the required policy, enable *Proxy Options* and select a profile that has *Enable Deep Scan* selected under HTTPS. **Note**: If no Proxy Options profile is selected, or if *Enable Deep Scan* is not selected within the Proxy Options profile, DLP rules ***cannot*** inspect HTTPS.

Table 11:SSL content scanning and inspection settings (continued)

Setting	Description
DLP archiving	DLP archiving for HTTPS, IMAPS, POP3S, and SMTPS. Add DLP Rules for the protocol to be archived.
Monitor DLP content information on the system dashboard	DLP archive information on the Log and Archive Statistics widget on the system dashboard for HTTPS, IMAPS, POP3S, and SMTPS. Go to *Policy > Policy > Proxy Options.* Add or edit a profile. For each protocol you want monitored on the dashboard, enable *Monitor Content Information for Dashboard*. These options display meta-information on the Statistics dashboard widget.

Exeptions

Periodically, you will come across situations were SSL and certificates will interfer with the smooth operation of an application or website. For instance, there is a popular application called Dropbox that does not work when deep SSL inspection is enabled. The reason for this is that the trusted certificate authority that is recognised by Dropbox is imbedded in the software and Dropbox cannot be reconfigured to recognise the FortiGate certificates that are used when deep SSL inspection is implimented.

One way to by-pass the deep inspection for Dropbox is to add dropbox.com to a local category in webfiltering and add that local category to the `ftgd-wf-ssl-exempt` list in the webfilter profile. This way any connections with dropbox.com will be exempt from deep SSL inspection.

Whenever an exception is found, the reason that it causes an issue will have to be determined in order to figure out a way to accommodate that application or website.

Monitoring Security Profiles activity

The first two steps in monitoring activity covered by Security profiles is make sure that logging is enabled on the FortiGate and that the policies are configured to collect those logs as traffic goes through them.

Check the Logging and Reporting handbook for configuration of such details as to whether the logs are stored locally on a disk or in memory, or use a remote service of some kind such as a FortiAnalyzer or SNMP server. The important thing is that the storing of logs is taking place somewhere. This is configured by going to Log & Report > Log Config > Log Setting. If you are going to log locally you will also have to enable logging locally in the CLI.

The next step is to get the firewall policies to collect traffic logs. In the configuration of policies there are 3 logging options:

• No log

• Log Security events

• Log all Sessions

Make sure that either Log Security events or Log all sessions is selected.

There are two ways to view the Security Profiles activity based on the collected logs. The first gives you an overview based on a sampling of logs over time. This is good for spotting trends and giving you an idea of the overall impact of a type of Security Profiles threat. For instance you can see if you are a lot of your users are trying to get to sites that you have blocked or which email protocol is receiving the most blocked email.

Go to *Security Profiles > Monitor*. From here you can choose information from the differnt types of Security profile that you have running.

From the AV Monitor, you can see information relating to the Antivirus Profile.

- What are the Top Viruses coming through the FortiGate unit, listing:
  - Virus name,
  - Last time it was detected
  - A count of how many times it was detected.

From the Web Monitor you can find information relating to Web filtering. You can choose:

- Report by FortiGuard Webfilter Category
  - Top blocked Categories (pie chart and graph)
  - Total blocked requests
- Report by Webfilter Technicue
  - Pie chart of requests (allowed, etc.)
  - Blocked Requests (Bar chart)
  - Spam
  - Banned Word
  - Virus Archive
  - FortiGuard
  - URL Filter
  - Fragmented
  - DLP

From the Application Monitor you can get an idea of which applications are being used over your network and who is using them by looking at the charts:

- Top Application by Bandwidth
- Top Applications by Session Count
- Top IP/User for…

From the Intrusion Monitor you can determine what are the Top Attacks against your network. The report will list:

- Attack Name
- Last time the attack was detected
- A count of how many times the attack was used

From the Email Monitor

- Total Emails (pie chart)
- Blocked Emails, broken down by
  - Protocol used
  - Reason/technique used to block

From the Archive & Data Leak Monitor you can see what is the:

- Top DLP usage by policy
- Total Dropped Archives

From the FortiGuard Quota you can monitor the status of quotas by seeing which ones are in effect listing:

- User name
- Webfilter Profile
- Used Quota

The second way to look at the logs of the Security Profiles activity is to look at the individual logs. This is useful for trouble shooting and verification of what is being tracked and how because individual log display more information about what happened to the traffic in question.

To look at the logs go to Log and Report > Traffic Log > Forward Traffic and search for individual logged events. In order to see just the Security Profiles based events you may have to first display a column that relates to Security Profiles such as Security Action, Security Event or Security Sub type. Once the appropriate column is displayed in the log window you can then filter based on the criteria that you are searching on. For instance if you were looking for examples of where your DLP profile stopped some traffic from getting out you could go to the Security event column and then filter for the event "dlp". The log page will now only display dlp events. You could not further refine your filter until you were only looking at the logs that relate to the events they you are trying to track.

Configuring packet logging options

You can use a number of CLI commands to further configure packet logging.

Limiting memory use

When logging to memory, you can define the maximum amount of memory used to store logged packets.

```
config ips settings
   set packet-log-memory 256
end
```

The acceptable range is from 64 to 8192 kilobytes. This command affects only logging to memory.

Limiting disk use

When logging to the FortiGate unit internal hard disk, you can define the maximum amount of space used to store logged packets.

```
config ips settings
   set ips-packet-quota 256
end
```

The acceptable range is from 0 to 4294967295 megabytes. This command affects only logging to disk.

Configuring how many packets are captured

Since the packet containing the signature is sometimes not sufficient to troubleshoot a problem, you can specify how many packets are captured before and after the packet containing the IPS signature match.

```
config ips settings
   packet-log-history
   packet-log-post-attack
end
```

The `packet-log-history` command specifies how many packets are captured before and including the one in which the IPS signature is detected. If the value is more than `1`, the packet containing the signature is saved in the packet log, as well as those preceding it, with the total number of logged packets equalling the `packet-log-history` setting. For example, if `packet-log-history` is set to `7`, the FortiGate unit will save the packet containing the IPS signature match and the six before it.

The acceptable range for `packet-log-history` is from 1 to 255. The default is 1.

Setting `packet-log-history` to a value larger than `1` can affect the performance of the FortiGate unit because network traffic must be buffered. The performance penalty depends on the model, the setting, and the traffic load.

The `packet-log-post-attack` command specifies how many packets are logged after the one in which the IPS signature is detected. For example, if `packet-log-post-attack` is set to `10`, the FortiGate unit will save the ten packets following the one containing the IPS signature match.

The acceptable range for `packet-log-post-attack` is from 0 to 255. The default is 0.

Using wildcards and Perl regular expressions

Many Security Profiles feature list entries can include wildcards or Perl regular expressions.

For more information about using Perl regular expressions, see http://perldoc.perl.org/perlretut.html.

Regular expression vs. wildcard match pattern

A wildcard character is a special character that represents one or more other characters. The most commonly used wildcard characters are the asterisk (*), which typically represents zero or more characters in a string of characters, and the question mark (?), which typically represents any one character.

In Perl regular expressions, the '.' character refers to any single character. It is similar to the '?' character in wildcard match pattern. As a result:

- example.com not only matches example.com but also examplea.com, exampleb.com, examplec.com, and so on.

To add a question mark (`?`) character to a regular expression from the FortiGate CLI, enter Ctrl+V followed by ?. To add a single backslash character (`\`) to a regular expression from the CLI you must add precede it with another backslash character. For example, `example\\.com`.

To match a special character such as '.' and '*' use the escape character '\'. For example:

- To match example.com, the regular expression should be: example\.com

In Perl regular expressions, '*' means match 0 or more times of the character before it, not 0 or more times of any character. For example:

- exam*.com matches exammmm.com but does not match example.com

To match any character 0 or more times, use '.*' where '.' means any character and the '*' means 0 or more times. For example, the wildcard match pattern exam*.com should therefore be exam.*\.com.

Word boundary

In Perl regular expressions, the pattern does not have an implicit word boundary. For example, the regular expression "test" not only matches the word "test" but also any word that contains "test" such as "atest", "mytest", "testimony", "atestb". The notation "\b" specifies the word boundary. To match exactly the word "test", the expression should be \btest\b.

Case sensitivity

Regular expression pattern matching is case sensitive in the web and Email Filter filters. To make a word or phrase case insensitive, use the regular expression `/i`. For example, `/bad language/i` will block all instances of "bad language", regardless of case.

Perl regular expression formats

Table 12 lists and describes some example Perl regular expressions.

Table 12: Perl regular expression formats

Expression	Matches
abc	"abc" (the exact character sequence, but anywhere in the string)
^abc	"abc" at the beginning of the string
abc$	"abc" at the end of the string
a|b	Either "a" or "b"
^abc|abc$	The string "abc" at the beginning or at the end of the string
ab{2,4}c	"a" followed by two, three or four "b"s followed by a "c"
ab{2,}c	"a" followed by at least two "b"s followed by a "c"
ab*c	"a" followed by any number (zero or more) of "b"s followed by a "c"
ab+c	"a" followed by one or more b's followed by a c
ab?c	"a" followed by an optional "b" followed by a" c"; that is, either "abc" or "ac"
a.c	"a" followed by any single character (not newline) followed by a" c "
a\.c	"a.c" exactly
[abc]	Any one of "a", "b" and "c"
[Aa]bc	Either of "Abc" and "abc"
[abc]+	Any (nonempty) string of "a"s, "b"s and "c"s (such as "a", "abba", "acbabcacaa")
[^abc]+	Any (nonempty) string which does not contain any of "a", "b", and "c" (such as "defg")
\d\d	Any two decimal digits, such as 42; same as \d{2}
/i	Makes the pattern case insensitive. For example, `/bad language/i` blocks any instance of `bad language` regardless of case.

Table 12: Perl regular expression formats (continued)

`\w+`	A "word": A nonempty sequence of alphanumeric characters and low lines (underscores), such as foo and 12bar8 and foo_1
100\s*mk	The strings "100" and "mk" optionally separated by any amount of white space (spaces, tabs, newlines)
abc\b	"abc" when followed by a word boundary (for example, in "abc!" but not in "abcd")
perl\B	"perl" when not followed by a word boundary (for example, in "perlert" but not in "perl stuff")
\x	Tells the regular expression parser to ignore white space that is neither preceded by a backslash character nor within a character class. Use this to break up a regular expression into (slightly) more readable parts.
/x	Used to add regular expressions within other text. If the first character in a pattern is forward slash '/', the '/' is treated as the delimiter. The pattern must contain a second '/'. The pattern between '/' will be taken as a regular expressions, and anything after the second '/' will be parsed as a list of regular expression options ('i', 'x', etc). An error occurs if the second '/' is missing. In regular expressions, the leading and trailing space is treated as part of the regular expression.

Examples of regular expressions

Block any word in a phrase

/block|any|word/

Block purposely misspelled words

Spammers often insert other characters between the letters of a word to fool spam blocking software.

/^.*v.*i.*a.*g.*r.*o.*$/i

/cr[eéèêë][\+\-*=<>\.\,;!\?%&§@\^°\$£€\{\}()\[\]\|_01]dit/i

Block common spam phrases

The following phrases are some examples of common phrases found in spam messages.

/try it for free/i

/student loans/i

/you're already approved/i

/special[\+\-*=<>\.\,;!\?%&~#§@\^°\$£€\{\}()\[\]\|_1]offer/i

Monitor interface reference

The Monitor submenus allow you to view the Security Profiles activity occurring on your network. You must have Security Profiles and sensors applied to firewall policies, as well as logging enabled for the profiles and sensors, for the monitors to display any information regarding this activity.

This topic contains the following:

- AV Monitor
- Intrusion Monitor
- Web Monitor
- Email Monitor
- Archive & Data Leak Monitor
- Application Monitor

AV Monitor

The AV Monitor submenu allows you to view statistical information regarding viruses that were detected on your unit from *Security Profiles > Monitor > AV Monitor*. The information displays in a bar chart as well as in a table below the bar chart. The table contains detailed information.

You must have antivirus logging enabled for this within the profile itself, as well as within log settings and an antivirus profile is applied to a firewall policy.

AV Monitor page

Displays monitored information about viruses that were detected by the unit.

Tip: To view information about a specific virus, select a bar within the chart; the virus FortiGuard definition displays.

Refresh	Select to refresh the information on the page.
Reset	Select to reset the information to clear the current information from the page. New information is included on the page.
Top Viruses (all policies) since <yyyy-mm-dd hh:mm:ss>	The top viruses detected by the unit using all firewall policies.
#	The order that the viruses are listed in the table.
Virus Name	The name of the virus.
Last Detected	The last time that the virus was detected.
Count	The number of times the virus has been detected.

Intrusion Monitor

The Intrusion Monitor submenu allows you to view statistical information regarding attacks that were detected on your unit from *Security Profiles > Monitor > Intrusion Monitor*. The information

displays in a bar chart as well as in a table below the bar chart. The table contains detailed information.

Intrusion Monitor page

Displays monitored information about attacks that were detected by the unit.

Tip: To view information about a specific attack, select a bar within the chart; the attack FortiGuard definition displays.

Refresh	Select to refresh the information on the page.
Reset	Select to reset the information to clear the current information from the page. New information is included on the page.
Top Attacks (all policies) since <yyyy-mm-dd hh:mm:ss>	A bar chart displaying the top attacks detected by the unit.
#	The order that the attacks are listed in the table.
Attack Name	The name of the attack.
Last Detected	The last time that the attack was detected.
Count	The number of times the attack has been detected.

Web Monitor

The Web Monitor submenu allows you to view statistical information regarding the web activity from *Security Profiles > Monitor > Web Monitor*. The information displays in both a pie chart and a bar chart .

Web Monitor page	
Displays monitored information about web activity detected by the unit.	
Refresh	Select to refresh the information on the page.
Reset	Select to reset the information to clear the current information from the page. New information is included on the page.
Report By	Select whether to view the web filter monitored information by web filter technique or by FortiGuard web filter category. If you choose FortiGuard web filter category, you are viewing the information that was gathered from the category settings for FortiGuard web filter from the web filter profile.
Web Monitor since <yyyy-mm-dd hh:mm:ss>	
Total Requests (HTTP)	A pie chart representing the total requests detected.

Blocked Requests (HTTP)	A bar chart representing the total blocked requests detected. The information is broken down to spam, banned words, file filter, viruses, archives, FortiGuard, URL filter, and fragmented.
Total Web Requests (HTTP): \<number>	The total number of web requests over HTTP that occurred.

Email Monitor

The Email Monitor submenu allows you to view statistical information regarding email filtering from *Security Profiles > Monitor > Email Monitor*. The information displays in both a pie chart and bar chart.

Email Monitor page

Displays monitored information about email filter activity detected by the unit.

Refresh	Select to refresh the information on the page.
Reset	Select to reset the information to clear the current information from the page. New information is included on the page.
Total Emails	A pie chart representing the total number of emails scanned by the unit.
Blocked Emails	A bar chart representing the total number of blocked emails, broken down by protocol. The colors indicate the type of scanning that occurred.
Total Emails: \<number>	The total number of email messages detected by the unit.

Archive & Data Leak Monitor

The Archive & Data Leak Monitor submenu allows you to view statistical information regarding log archives, as well as DLP usage. This page displays the information in a bar chart in *Security Profiles > Monitor > Archive & Data Leak Monitor*.

Archive & Data Leak Monitor page

Displays monitored information about archive and DLP activity detected by the unit.

Refresh	Select to refresh the information on the page.
Reset	Select to reset the information to clear the current information from the page. New information is included on the page.
Report By:	Select what type of DLP information you want to view. You can view DLP usage by DLP sensor, firewall policy usage, or by protocol.
Top DLP Usage by DLP Sensor \<yyyy-mm-dd hh:mm:ss>	The bar chart that displays DLP usage monitored using DLP sensor information.

Top DLP Usage by Policy **<yyyy-mm-dd hh:mm:ss>**	The bar chart that displays DLP usage monitored using firewall policy traffic information.
Top DLP Usage by Protocol **<yyyy-mm-dd hh:mm:ss>**	The bar chart that displays DLP usage monitored using protocol information.
Total Dropped Archives: **<number>**	The total number of dropped DLP archives.

Application Monitor

The Application Monitor submenu allows you to view statistical information regarding application usage in *Security Profiles > Monitor > Application Monitor*.

Application Monitor page	
Displays monitored information about the application usage detected by the unit. **Tip:** To view top source IP addresses for a specific application, select a bar in the chart to view that application's source IP addresses.	
Refresh	Select to refresh the information on the page.
Reset	Select to reset the information to clear the current information from the page. New information is included on the page.
Top Application Usage by **<yyyy-mm-dd hh:mm:ss>**	The bar chart that displays the top applications being used detected by the unit.
Resolve Host Name	Appears after selecting a bar for a specific application, for example SSL. Select to resolve the host name. **Tip:** Hover your mouse over the bar to view the address and total MB (or KB) used for that application.
Report By:	Appears after selecting a bar for a specific application, for example, SSL. Select to view the detailed information by destination address, or source address.
Display User Name	Appears after selecting *Source Address* from the drop-down list beside *Report By*. Select to display user names.

FortiGuard Quota

The FortiGuard Quota submenu allows you to view statistical information regarding quota usage by users in *Security Profiles > Monitor > FortiGuard Quota*.

FortiGuard Quota page

Lists the users and the amount of quota that they have used.

Page Controls	Use to navigate through the list.
User Name	The user name of the user that has FortiGuard quota enabled for them.
Webfilter Profile	The web filter profile that was used for detecting users' FortiGuard quota usage.
Used Quota	The amount of used quota by a user.

Endpoint Monitor

You can view monitored endpoints in *Security Profiles > Monitor > Endpoint Monitor*. An endpoint is added to the list when it uses a security policy that has *Endpoint Security* enabled.

Endpoint Monitor page

Provides information about endpoints, such as endpoint traffic.

Note: The pie chart displays information in percent and indicates which is non-compliant and which is compliant.

Refresh	Updates the list, providing current endpoints that are being monitored.
Report By	Select to view endpoint information by traffic, status or application usage.
	When you select *Status*, a pie chart appears along with information about the total endpoints (*Total Endpoints*). When you select *Traffic* or *Application usage*, a bar chart appears; select a bar to view detailed information.

Index