



CHAPTER 6

Configuring Zone Filters

This chapter describes how to configure the Cisco Guard (Guard) network traffic filters.

This chapter refers to the Cisco Detector (Detector), the companion product of the Guard. The Detector is a Distributed Denial of Service (DDoS) attack detection device that analyzes a copy of the zone traffic. The Detector can activate the Guard attack mitigation services when the Detector determines that the zone is under attack. The Detector can also synchronize zone configurations with the Guard. For more information about the Detector, see the *Cisco Traffic Anomaly Detector Module Configuration Guide* and *Cisco Traffic Anomaly Detector Configuration Guide*.

This chapter contains the following sections:

- [Understanding Zone Filters](#)
- [Configuring Flex-Content Filters](#)
- [Configuring Bypass Filters](#)
- [Configuring User Filters](#)
- [Configuring Dynamic Filters](#)

Understanding Zone Filters

Zone filters define how the Guard handles a specific traffic flow. You can configure filters to customize the traffic flow and control the Distributed Denial of Service (DDoS) attack mitigation operation.

Zone filters enable the Guard to perform the following functions:

- Analyze zone traffic for anomalies
- Apply the basic or strong level of protection to separate legitimate traffic from malicious traffic
- Drop malicious packets
- Forward traffic directly to the zone, bypassing the Guard protection features

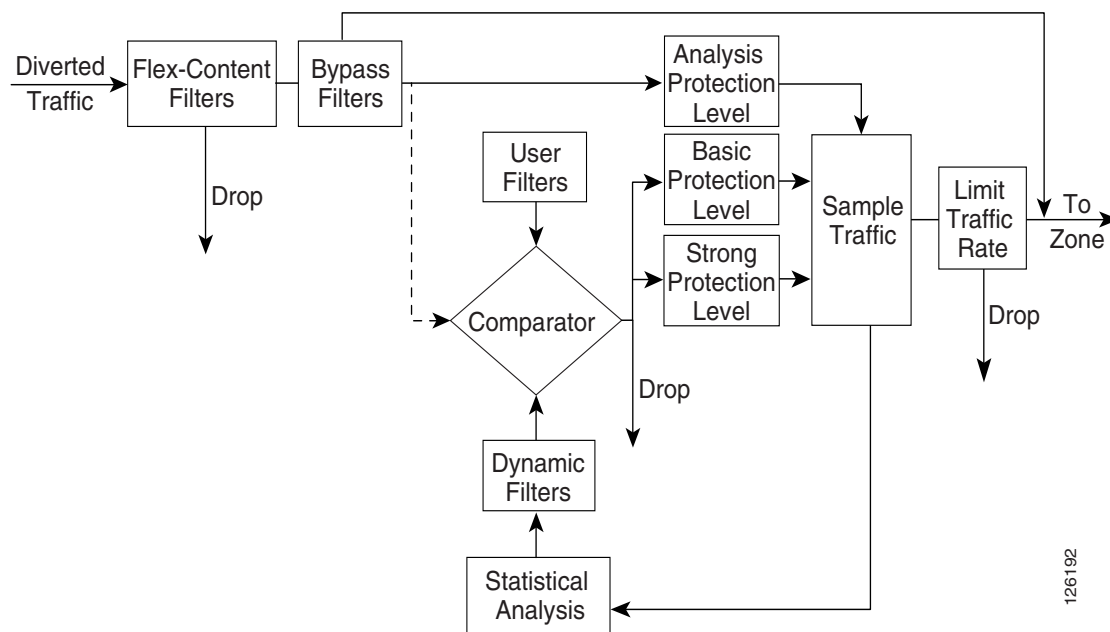
The Guard has the following types of filters:

- User filters—Apply the required protection level to the specified traffic flow. User filters define the first actions that the Guard takes when it identifies abnormal or malicious traffic. The zone configuration includes a default set of user filters configured for on-demand protection that can handle a wide range of attack types. You can modify user filters to customize the Guard protection capabilities and to set rules about how the Guard handles traffic flows when an attack is suspected. See the [“Configuring User Filters” section on page 6-13](#) for more information.

- Bypass filters—Prevent the Guard from analyzing specific traffic flows. You can direct trusted traffic away from the Guard protection features and forward it directly to the zone. See the “Configuring Bypass Filters” section on page 6-11 for more information.
- Flex-content filters—Count or drop a specific traffic flow. Flex-content filters provide extremely flexible filtering capabilities, such as filtering according to fields in the IP and TCP headers, filtering based on the payload content, and filtering based on complex Boolean expressions. See the “Configuring Flex-Content Filters” section on page 6-3 for more information.
- Dynamic filters—Apply the required protection level to the specified traffic flow. The Guard creates dynamic filters based on the analysis of traffic flow and continuously modifies this set of filters to zone traffic and the type of DDoS attack. The dynamic filters have a limited life span and are deleted by the Guard when the attack ends. See the “Configuring Dynamic Filters” section on page 6-18 for more information.

Figure 6-1 displays the Guard filter system.

Figure 6-1 Guard Filter System



When zone protection is enabled either by a user action or by a remote network-sensing DDoS element, such as the Detector, the Guard diverts the zone traffic to itself and analyzes the traffic.

The Guard monitors the rate of the traffic that flows to the zone. It drops traffic that exceeds the defined rate and forwards the legitimate traffic to the zone. The Guard performs statistical analysis of zone traffic and performs a closed-loop feedback cycle to adjust the protection measures to the dynamically changing zone traffic characteristics and the changing DDoS attack types.

To perform statistical analysis of traffic flow, the Guard uses the zone policies which are all configured to handle specific types of traffic. The zone policies constantly measure traffic flows and take action against a particular traffic flow if they identify that flow as malicious or abnormal, which occurs when the flow exceeds the policy threshold.

When the Guard identifies a traffic anomaly, it performs the following tasks:

1. Produces dynamic filters that are configured with actions to handle the attack. The Guard, by default, adds an initial dynamic filter that directs all traffic to user filters that provide the first line of defense against an evolving DDoS attack. Once the Guard has had enough time to analyze the attack, it begins producing dynamic filters to mitigate the attack.
2. Changes the flow of traffic within the Guard. Abnormal traffic flows into the Comparator, which is a component that receives input from the dynamic filters and the user filters. The Comparator compares the first user filter that matches the flow with the dynamic filters and chooses the most severe protection measure suggested. It applies the relevant protection level to authenticate the traffic.

By default, the Guard protects the zone until you deactivate zone protection.

Configuring Flex-Content Filters

Flex-content filters filter the zone traffic based on the fields in the packet header or the patterns in the packet payload. You can identify attacks that are based on the patterns that appear in the traffic. These patterns can identify known worms or flood attacks that have a constant pattern.

Use the flex-content filters to drop or count a desired packet flow and to identify a specific malicious source of traffic.

The flex-content filter applies the filtering criteria in the following order:

1. Filters packets based on the protocol and the port parameter values.
2. Filters packets based on the tcpdump-expression value.
3. Performs pattern matching with the pattern-expression value on the remaining packets.



Note

Flex-content filters consume a lot of CPU resources. We recommend that you limit the use of flex-content filters because they might affect the performance of the Guard. If you are using a flex-content filter to protect a specific attack that can be identified by a dynamic filter, such as TCP traffic to a specified port, we recommend that you filter the traffic using a dynamic filter.

This section contains the following topics:

- [Adding a Flex-Content Filter](#)
- [Displaying Flex-Content Filters](#)
- [Deleting Flex-Content Filters](#)
- [Changing the State of a Flex-Content Filter](#)

Adding a Flex-Content Filter

The Guard creates a list of flex-content filters that you create and activates the filters in an ascending order. When you add a new flex-content filter, make sure that you place it in the correct location in the filter list.

The Guard stops activating the flex-content filters when traffic matches a flex-content filter with a drop action.

To configure a flex-content filter, perform the following steps:

- Step 1** Display the list of flex-content filters and identify the location in the list in which you want to add the new filter (see the “[Displaying Flex-Content Filters](#)” section on page 6-9).
- Step 2** If the current row numbers are consecutive, renumber the flex-content filters in increments that allow you to insert the new flex-content filter by entering the following command in zone configuration mode:

```
flex-content-filter renumber [start [step]]
```

Table 6-1 provides the arguments for the **flex-content-filter renumber** command.

Table 6-1 Arguments for the flex-content-filter renumber Command

Parameter	Description
<i>start</i>	(Optional) Integer from 1 to 9999 that denotes the new starting number of the flex-content filter list. The default is 10.
<i>step</i>	(Optional) Integer from 1 to 999 that defines the increment between the flex-content filter row numbers. The default is 10.

- Step 3** (Optional) Filter a pattern expression of an ongoing attack or an attack that you have previously recorded. Activate the Guard to generate a signature of the attack by using the **show packet-dump signatures** command. See the “[Generating Attack Signatures from Packet-Dump Capture Files](#)” section on page 12-22 for more information.

- Step 4** Add a new flex-content filter by entering the following command:

```
flex-content-filter row-num {disabled | enabled} {drop | count} protocol port [start start-offset [end end-offset]] [ignore-case] expression tcpdump-expression pattern pattern-expression
```

Table 6-2 provides the arguments and keywords for the **flex-content-filter** command.

Table 6-2 Arguments and Keywords for the flex-content-filter Command

Parameter	Description
<i>row-num</i>	Unique number from 1 to 9999 that identifies the filter and defines the priority among the flex-content filters. The Guard operates the filters in ascending row-number order.
disabled	Sets the filter state to disabled. The filter does not monitor traffic.
enabled	Sets the filter state to enabled. The Guard monitors traffic and performs the action (drop or count) on the flow that matches the filter. This is the default state.
drop	Drops the flow that matches the filter.
count	Counts the flow that matches the filter.
<i>protocol</i>	Traffic from a specific protocol. Use an asterisk (*) to indicate any protocol. Enter an integer from 0 to 255. Review possible protocol numbers at the Internet Assigned Numbers Authority (IANA) website: http://www.iana.org/assignments/protocol-numbers

Table 6-2 Arguments and Keywords for the flex-content-filter Command (continued)

Parameter	Description
<i>port</i>	<p>Traffic destined to a specific destination port. Enter an integer from 0 to 65535. To define a specific port number, you must define a specific protocol number.</p> <p>Use an asterisk (*) to indicate any destination port. You can use an asterisk if you configure the protocol number to 6 (TCP) or 17 (UDP).</p> <p>Review possible port numbers at the Internet Assigned Numbers Authority (IANA) website: http://www.iana.org/assignments/port-numbers</p>
<i>start-offset</i>	<p>Offset, in bytes, from the beginning of the packet payload, where the pattern matching for the <i>pattern-expression</i> argument begins. The default is 0, which is the start of the payload. Enter an integer from 0 to 1800.</p> <p>If you copy the pattern from the show packet-dump signatures command output, copy this argument from the Start Offset field in the command output.</p>
<i>end-offset</i>	<p>Offset, in bytes, from the beginning of the packet payload, where the pattern matching for the <i>pattern-expression</i> argument ends. The default is the packet length, which is the end of the payload. Enter an integer from 0 to 1800.</p> <p>If you copy the pattern from the show packet-dump signatures command output, copy this argument from the End Offset field in the command output.</p>
ignore-case	<p>Defines the <i>pattern-expression</i> argument as case insensitive.</p> <p>By default, the <i>pattern-expression</i> argument is case sensitive.</p>
expression <i>tcpdump-expression</i>	<p>Specifies the expression that is matched with the packet. The expression is in Berkeley Packet filter format. See the “Configuring the tcpdump-expression Syntax” section on page 6-6 for more information and configuration examples.</p> <p>If you use spaces in the expression, enclose the expression in quotation marks (“ ”).</p> <p>To enter an empty expression, use double quotation marks (“ ”).</p> <p>To use a quotation mark in the expression, use the backslash escape character before the quotation mark (\”).</p> <p>Note Help is not available for the tcpdump-expression syntax.</p>
pattern <i>pattern-expression</i>	<p>Regular expression data pattern that is to be matched with the packet payload. See the “Configuring the pattern-expression Syntax” section on page 6-8 for more information.</p> <p>You can activate the Guard to generate the signature by using the show packet-dump signatures command. See the “Generating Attack Signatures from Packet-Dump Capture Files” section on page 12-22.</p> <p>If you use spaces in the expression, enclose the expression in quotation marks (“ ”).</p> <p>To enter an empty expression, use double quotation marks (“ ”).</p> <p>To use a quotation mark in the expression, use the backslash escape character before the quotation mark (\”).</p> <p>Note Help is not available for the pattern-expression syntax.</p>

You can change the filter state to enable or disable at any time. See the “[Changing the State of a Flex-Content Filter](#)” section on page 6-10 for more information.

You can delete a filter at any time (see the “Deleting Flex-Content Filters” section on page 6-10).

The following example shows how to configure the flex-content filter:

```
user@GUARD-conf-zone-scanner# flex-content-filter enabled count * * expression "ip[6:2] &
0x1fff=0" pattern
"/ HTTP/1\.\.1\ x0D\0AAccept: .*/.\.*\x0D\x0AAccept-Language: en*\x0D\x0AAccept-Encoding:
gzip, deflate\x0D\x0AUser-Agent: Mozilla/4\.\.0"
```

This section contains the following topics:

- [Configuring the tcpdump-expression Syntax](#)
- [Configuring the pattern-expression Syntax](#)

Configuring the tcpdump-expression Syntax

The tcpdump-expression is in the Berkeley Packet filter format and specifies the expression to be matched with the packet.



Note

You can use the tcpdump-expression to filter traffic based on the destination port and protocol, but the performance of the Guard may be affected. We recommend that you filter traffic based on these criteria using the flex-content filter *protocol* and *port* arguments.

The expression contains one or more elements which usually consist of an ID preceded by one or more qualifiers.

There are three types of qualifiers:

- Type qualifiers—Define the ID (name or number). Possible types are **host**, **net**, and **port**. The **host** type qualifier is the default.
- Direction qualifiers—Define the transfer direction. Possible directions are **src**, **dst**, **src or dst**, and **src and dst**. The direction qualifier **src or dst** is the default.
- Protocol qualifiers—Restrict the match to a particular protocol. Possible protocols are **ether**, **ip**, **arp**, **rarp**, **tcp**, and **udp**. If you do not specify a protocol qualifier, all protocols that apply to the type are matched. For example, port 53 means TCP or UDP port 53.

Table 6-3 describes the tcpdump-expression elements.

Table 6-3 tcpdump-expression Elements

Element	Description
dst host <i>host_ip_address</i>	Traffic to a destination host IP address.
src host <i>host_ip_address</i>	Traffic from a source host IP address.
host <i>host_ip_address</i>	Traffic to and from both source and destination host IP addresses.
net net mask <i>mask</i>	Traffic to a specific network.
net <i>net/len</i>	Traffic to a specific subnet.
dst port <i>destination_port_number</i>	TCP or UDP traffic to a destination port number.

Table 6-3 *tcpdump-expression Elements (continued)*

Element	Description
src port <i>source_port_number</i>	TCP or UDP traffic from a source port number.
port <i>port_number</i>	TCP or UDP traffic to and from both source and destination port numbers.
less <i>packet_length</i>	Packets with a length equal to or less than the specific length in bytes.
greater <i>packet_length</i>	Packets with a length equal to or greater than the specific length in bytes.
ip proto <i>protocol</i>	Packets with a protocol number of the following protocols: ICMP, UDP, and TCP.
ip broadcast	Broadcast IP packets.
ip multicast	Multicast packets.
ether proto <i>protocol</i>	Ether protocol packets of a specific protocol number or name such as IP, ARP, or RARP. The protocol names are also keywords. If you enter the protocol name, you must use a backslash (\) as an escape character before the name.
<i>expr relop expr</i>	Traffic that complies with the specific expression. Table 6-4 describes the tcpdump-expression rules.

[Table 6-4](#) describes the tcpdump-expression rules.

Table 6-4 *Flex-Content Filter Expression Rules*

Expression Rule	Description
<i>relop</i>	>, <, >=, <=, =, !=
<i>expr</i>	Arithmetic expression composed of integer constants (expressed in standard C syntax), the normal binary operators [+ , - , * , / , & ,], a length operator, and special packet data accesses. To access data inside the packet, use the following syntax: <i>proto [expr: size]</i>
<i>proto</i>	Protocol layer for the index operation. The possible values are ether, ip, tcp, udp, or icmp. The byte offset, relative to the indicated protocol layer, is given by the expr value. To access data inside the packet, use the following syntax: <i>proto [expr: size]</i> The <i>size</i> argument is optional and indicates the number of bytes in the field. The argument can be 1, 2, or 4. The default is 1.

You can combine expression elements using the following methods:

- A group of elements and operators in parentheses—The operators are the normal binary operators [+ , - , * , / , & , |] and a length operator.



Note To use a parenthesis in the expression, use the backslash escape character before the parenthesis (\ ().

- Negation—Use **!** or **not**.
- Concatenation—Use **&&** or **and**.
- Alternation—Use **||** or **or**.

Negation has the highest precedence. Alternation and concatenation have equal precedence and are associated from left to right. Explicit and tokens, not juxtaposition, are required for concatenation. If you specify an identifier without a keyword, the most recent keyword is used.

For a detailed explanation of the Berkeley Packet filter configuration options, go to this location:

<http://www.freesoft.org/CIE/Topics/56.htm>.

The following example shows how to count unfragmented datagrams and fragmented zeros of fragmented datagrams only. This filter is implicitly applied to the TCP and UDP index operations. For instance, `tcp[0]` always indicates the first byte of the TCP header and never indicates the first byte of an intervening fragment as shown in this example:

```
user@GUARD-conf-zone-scanner# flex-content-filter enabled count * * expression
ip[6:2]&0x1fff=0 pattern ""
```

The following example shows how to drop all TCP RST packets:

```
user@GUARD-conf-zone-scanner# flex-content-filter enabled drop * * expression tcp[13]&4!=0
pattern ""
```

The following example shows how to count all ICMP packets that are not echo requests/echo replies (ping):

```
user@GUARD-conf-zone-scanner# flex-content-filter enabled count * * expression "icmp
[0]!=8 and icmp[0] != 0" pattern ""
```

The following example shows how to count all TCP packets that are destined to port 80 and that did not originate from port 1000:

```
user@GUARD-conf-zone-scanner# flex-content-filter enabled count * * expression "tcp and
dst port 80 and not src port 1000" pattern ""
```

Configuring the pattern-expression Syntax

The pattern-expression syntax is a regular expression that describes a string of characters. The pattern-expression describes a set of strings without actually listing its elements. This expression consists of normal characters and special characters. Normal characters include all printable ASCII characters that are not considered to be special characters. Special characters have a special meaning and specify the type of matching that the Guard performs on the pattern-expression. The flex-content filter matches the pattern-expression with the content of the packet (the packet payload). For example, the three strings version 3.1, version 4.0, and version 5.2 are described by the following pattern: `version.*\.*`

[Table 6-5](#) describes the special characters that you can use.

Table 6-5 Special Characters Used in the pattern-expression

Special character	Description
.	Matches a string that may be present and can contain zero or more characters. For example, the pattern goo.*s matches the patterns goos, goods, good for dds, and so on.
\	Removes the special meaning of a special character. To use the special characters in this list as single-character patterns, remove the special meaning by preceding each character with a backslash (\). For example, two backslashes (\\) match one backslash (\), and one backslash and a period (\.) match one period (.). You must also precede an asterisk (*) with a backslash.
\xHH	Matches a hexadecimal value, where H is a hexadecimal digit and is not case sensitive. Hexadecimal values must be exactly two digits. For example, the pattern \x41 matches the hexadecimal value A.

By default, the pattern-expression is case sensitive. To define the pattern-expression as case insensitive, use the **flex-content-filter** command with the **ignore-case** keyword. See the “[Adding a Flex-Content Filter](#)” section on page 6-3 for more information.

The following example shows how to drop packets with a specific pattern in the packet payload. The pattern in the example was extracted from the Slammer worm. The *protocol*, *port*, and *tcpdump-expression* parameters are nonspecific.

```
user@GUARD-conf-zone-scanner# flex-content-filter enabled drop * * expression " " pattern
\x89\xe5qh\ .dllhel32hkernqhounthickChGetTf\xB911
Qh32\ .dhws2_f\xB9etQhsockf\xB9toQhsend\xBE\x18\x10\xAEB
```

Displaying Flex-Content Filters

To display the flex-content filters, use the following command in zone configuration mode:

```
show flex-content-filters
```

Table 6-6 describes the fields in the **show flex-content-filters** command output.

Table 6-6 Field Descriptions for the show flex-content-filters Command

Field	Description
Row	Flex-content filter priority.
State	Filter state (enabled or disabled).
Action	Action that the filter performs on the specific traffic type.
Protocol	Protocol number of the traffic that the filter processes.
Port	Destination port of the traffic that the filter processes.
Start	Offset, in bytes, from the beginning of the packet payload where the pattern matching begins. This offset applies to the <i>pattern</i> field.

Table 6-6 *Field Descriptions for the show flex-content-filters Command (continued)*

Field	Description
End	Offset, in bytes, from the beginning of the packet payload where the pattern matching ends. This offset applies to the <i>pattern</i> field.
Match-case	Whether the pattern-expression that the filter matches is case sensitive or not case sensitive. yes = case-sensitive, no = case-insensitive
TCPDump-expression	tcpdump-expression to be matched with the packet in Berkeley Packet filter format. See the “ Configuring the tcpdump-expression Syntax ” section on page 6-6 for the information on the tcpdump-expression syntax.
Pattern-filter	Regular expression data pattern to be matched with the packet payload. See the “ Configuring the pattern-expression Syntax ” section on page 6-8 for information on the pattern-expression syntax.
RxRate (pps)	Current traffic rate in packets per second that is measured for this filter.

Deleting Flex-Content Filters

You can delete a flex-content filter when you no longer need it to filter packets based on the filter expression.



Note

Do not delete a flex-content filter if you might need it at a later date. You can disable the flex-content filter and then enable it when needed (see the “[Changing the State of a Flex-Content Filter](#)” section on [page 6-10](#)).

To delete a flex-content filter, enter the following command in zone configuration mode:

```
no flex-content-filter row-num
```

The *row-num* argument specifies the flex-content filter row number to delete. To display the list of flex-content filters and identify the row number of the flex-content filter to delete, use the **show flex-content-filters** command (see the “[Displaying Flex-Content Filters](#)” section on [page 6-9](#)). To delete all flex-content filters, enter an asterisk (*) for the row number.

The following example shows how to delete a flex-content filter:

```
user@GUARD-conf-zone-scanner# no flex-content-filters 5
```

Changing the State of a Flex-Content Filter

You can disable a flex-content filter to prevent the Guard from filtering packets based on the filter expression and to prevent it from filtering specific types of traffic. When you disable the filter, it remains in the flex-content filter list, which allows you to enable the filter again if needed.

If you do not intend to use a flex-content filter again, you can delete it (see the “[Deleting Flex-Content Filters](#)” section on [page 6-10](#)).

To change the state of a flex-content filter, enter the following command in zone configuration mode:

```
flex-content-filter row-num {disabled | enabled}
```

Table 6-7 provides the arguments for the **flex-content-filter** command.

Table 6-7 Arguments and Keywords for the flex-content-filter Command

Parameter	Description
<i>row-num</i>	Specifies the flex-content filter row number.
disabled	Disables the specified flex-content filter.
enabled	Enables the specified flex-content filter.

To display the list of flex-content filters and identify the row number of the flex-content filter to enable or disable, enter the **show flex-content-filters** command (see the “[Displaying Flex-Content Filters](#)” section on page 6-9).

The following example shows how to disable a flex-content filter:

```
user@GUARD-conf-zone-scannet# flex-content-filters 5 disabled
```

Configuring Bypass Filters

The bypass filter allows you to specify traffic that you want the Guard to forward directly to the zone without applying any traffic protection functions, including the anti-spoofing and anti-zombie functions.



Note

The Guard injects traffic that passes through the bypass filters on to the zone without applying a limit on the traffic rate that was defined by using the **rate-limit** command.

This section contains the following topics:

- [Adding a Bypass Filter](#)
- [Displaying Bypass Filters](#)
- [Deleting Bypass Filters](#)

Adding a Bypass Filter

To add a bypass filter, use the following command in zone configuration mode:

```
bypass-filter row-num src-ip [ip-mask] protocol dest-port [fragments-type]
```

Table 6-8 provides the arguments for the **bypass-filter** command.

Table 6-8 Arguments for the bypass-filter Command

Parameter	Description
<i>row-num</i>	Unique number from 1 to 9999. The row-number identifies the filter and defines the priority among the bypass filters. The Guard operates the filters according to the ascending row-number order.
<i>src-ip</i>	Traffic from a specific IP address is processed. Use an asterisk (*) to indicate any IP address.

Table 6-8 Arguments for the bypass-filter Command (continued)

Parameter	Description
<i>ip-mask</i>	(Optional) Traffic from a specific subnet is processed. The subnet mask can contain only Class C values. The default subnet is 255.255.255.255.
<i>protocol</i>	Traffic from a specific protocol is processed. Use an asterisk (*) to indicate any protocol. Review possible protocol numbers at the Internet Assigned Numbers Authority (IANA) website: http://www.iana.org/assignments/protocol-numbers
<i>dest-port</i>	Traffic to a specific destination port is processed. Use an asterisk (*) to indicate any destination port. Review possible port numbers at the Internet Assigned Numbers Authority (IANA) website: http://www.iana.org/assignments/port-numbers
<i>fragments-type</i>	(Optional) Whether or not the filter processes fragmented traffic. The three fragmented types are as follows: <ul style="list-style-type: none"> • no-fragments—Nonfragmented traffic • fragments—Fragmented traffic • any-fragments—Fragmented and nonfragmented traffic The default is no-fragments .



Note You cannot specify both a fragments type and a destination port. To set the fragments type, enter an asterisk (*) for the destination port.

Displaying Bypass Filters

To display the list of bypass filters, use the following command in zone configuration mode:

```
show bypass-filters
```

Table 6-9 describes the fields in the **show bypass-filters** command output.

Table 6-9 Field Descriptions for the show bypass-filters Command

Field	Description
Row	Bypass filter priority.
Source IP	Source IP address of the traffic that the filter processes.
Source Mask	Source address subnet mask of the traffic that the filter processes.
Proto	Protocol number of the traffic that the filter processes.
DPort	Destination port of the traffic that the filter processes.

Table 6-9 Field Descriptions for the `show bypass-filters` Command (continued)

Field	Description
Frg	Fragmentation settings that the filter processes: <ul style="list-style-type: none"> yes—The filter processes fragmented traffic. no—The filter processes nonfragmented traffic. any—The filter processes both fragmented and nonfragmented traffic.
RxRate (pps)	Current traffic rate in packets per second that is measured for this filter.

The source IP address, source address mask, protocol number, and destination port may be nonspecific. An asterisk (*) indicates that the filter acts on all field values or that more than one value was matched for the filter.

Deleting Bypass Filters

To delete a bypass filter, enter the following command in zone configuration mode:

```
no bypass-filter row-num
```

The *row-num* argument specifies the bypass filter row number to be deleted. To delete all bypass filters, enter an asterisk (*). To display the list of bypass filters and identify the row number of the bypass filter that you want to delete, use the **show bypass-filters** command (see the [“Displaying Bypass Filters” section on page 6-12](#)). To delete all bypass filters, enter an asterisk (*) for the row number.

The following example shows how to delete a bypass filter:

```
user@GUARD-conf-zone-scannet# no bypass-filter 10
```

Configuring User Filters

User filters define the first actions that the Guard executes when it identifies abnormal or malicious traffic. User filters either apply the required protection level to the specified traffic flows or they drop the specified traffic.

Each zone configuration includes a default set of user filters that are configured for on-demand protection and can handle a wide range of attack types. You can modify user filters to customize the Guard protection capabilities and to set rules about how the Guard handles specific traffic flows when it suspects an attack.

During an attack on the zone, the Guard continuously analyzes the traffic going to the zone. When it detects abnormal traffic patterns, the user filters provide the first line of defense against the evolving DDoS attack. Once the Guard has had enough time to analyze the attack, it begins producing dynamic filters that define how to handle the attack.

The Guard examines both the user filters and the dynamic filters before deciding how to handle the specific traffic flow. It compares the first user filter that matches the flow with the dynamic filters and chooses the most severe protection measure suggested. It applies the appropriate protection level to the traffic flow to authenticate the traffic. See the [“Configuring Dynamic Filters” section on page 6-18](#) for more information about dynamic filters.

The dynamic filters and the user filters can take actions in these descending severity levels: drop, strong, basic, and permit (see [Table 6-10](#)). Dynamic filters with actions of redirect/zombie and block-unauthenticated are applied even if a user filter to handle the same type of traffic exists because dynamic filters affect the Guard traffic authentication mechanisms and do not directly affect the traffic flow.

User filters are activated in ascending row-number order. When you add a new user filter, it is important that you place it in the correct location in the list.

[Table 6-10](#) describes the actions that a user filter can take.

Table 6-10 User Filter Actions

Action	Description
basic/default	Authenticates non-TCP traffic flows.
basic/dns-proxy	Authenticates TCP DNS traffic flows.
basic/redirect	Authenticates applications over HTTP.
basic/reset	Authenticates applications over TCP. We recommend that you use an action of basic/redirect for HTTP traffic flows.
basic/safe-reset	Authenticates TCP application traffic flows that are not tolerant of TCP connection reset. We recommend that you use an action of basic/redirect for HTTP traffic flows.
basic/sip	Authenticates VoIP ¹ applications that use SIP ² over UDP to establish the VoIP sessions and RTP/RTCP ³ to transmit voice data between the SIP end points after sessions are established.
drop	Drops traffic flows.
permit	Prevents statistical analysis of the flow and the anti-spoofing or anti-zombie protection functions from handling this flow. We recommend that you set a rate and burst limit to this filter because it is not handled by other protection mechanisms.
strong	Enables strong authentication for a traffic flow. Use this filter when strong authentication is required or when the previous filters do not seem suitable for the application. Authentication is performed for every connection. For TCP incoming connections, the Guard serves as a proxy. Do not use the strong authentication action for connections if you use ACLs ⁴ , access policies, or load-balancing policies that are based on the incoming IP address in the network.

1. VoIP = Voice over IP
2. SIP = Session Initiation Protocol
3. RTP/RTCP = Real-Time Transport Protocol/Real-Time Control Protocol
4. ACL = Access Control List

This section contains the following topics:

- [Adding User Filters](#)
- [Displaying User Filters](#)
- [Deleting User Filters](#)

Adding User Filters

To add a user filter, perform the following steps:

- Step 1** Display the list of user filters and identify the location in the list in which you want to add the new filter. See the “[Displaying User Filters](#)” section on page 6-17 for more information.
- Step 2** If the current row numbers are consecutive, renumber the user filters in increments that allow you to insert the new user filters by entering the following command:

```
user-filter renumber [start [step]]
```

[Table 6-11](#) provides the arguments for the **user-filter renumber** command.

Table 6-11 Arguments for the user-filter renumber Command

Parameter	Description
<i>start</i>	(Optional) Integer from 1 to 10000 that denotes the new starting number of the user filter list. The default is 10.
<i>step</i>	(Optional) Integer from 1 to 1000 that defines the increment between the user filter row numbers. The default is 10.

- Step 3** Add a new user filter by entering the following command:

```
user-filter row-num filter-action src-ip [ip-mask] protocol dest-port [fragments-type]
[rate-limit rate burst units]
```

[Table 6-12](#) provides the arguments for the **user-filter** command.

Table 6-12 Arguments and Keywords for the user-filter Command

Parameter	Description
<i>row-num</i>	Unique number from 1 to 1000 that identifies the filter and defines priority among the user filters. The Guard operates the filters according to the ascending row-number order.
<i>filter-action</i>	Action that the filter performs on the specific traffic type. See Table 6-10 for more information.
<i>src-ip</i>	Traffic from a specific IP address. Use an asterisk (*) to indicate any IP address.
<i>ip-mask</i>	(Optional) Traffic from a specific subnet. The subnet mask can contain only Class C values. The default subnet is 255.255.255.255.
<i>protocol</i>	Traffic from a specific protocol. Use an asterisk (*) to indicate any protocol. Review possible protocol numbers at the Internet Assigned Numbers Authority (IANA) website: http://www.iana.org/assignments/protocol-numbers
<i>dest-port</i>	Traffic to a specific destination port. Use an asterisk (*) to indicate any destination port. Review possible port numbers at the Internet Assigned Numbers Authority (IANA) website: http://www.iana.org/assignments/port-numbers
<i>fragments-type</i>	(Optional) Type of traffic. The type can be one of the following: <ul style="list-style-type: none"> • no-fragments—Nonfragmented traffic • fragments—Fragmented traffic • any-fragments—Fragmented and nonfragmented traffic The default is no-fragments .
rate-limit <i>rate</i>	Specifies the rate limitation. The user filter limits the traffic to this rate. Enter an integer greater than 64. The units are specified by the <i>units</i> parameter. The default is not to limit the filter traffic rate. The <i>rate</i> limit can be up to 10 times greater than the <i>burst</i> limit.
<i>burst</i>	Integer greater than 64 that specifies the traffic burst limit. The units are bits, kilobits, kilopackets, megabits, and packets that correspond to the units that are specified by the <i>units</i> parameter. The <i>burst</i> limit can be up to eight times greater than the <i>rate</i> limit.
<i>units</i>	Rate limit units. The units can be one of the following: <ul style="list-style-type: none"> • bps—Bits per second • kbps—Kilobits per second • kpps—Kilo packets per second • mbps—Megabits per second • pps—Packets per second

The following example shows how to renumber the user filters starting from 10 in steps of 5. This example also shows how to add a user filter in row 12 that is aimed at traffic that is received from all source IP addresses of protocol 6 (TCP) and flows to destination port 25 (SMTP). The user filter limits the traffic flow rate to 600 pps and the burst size to 400 packets.

```
user@GUARD-conf-zone-scannet# user-filter renumber 10 5
user@GUARD-conf-zone-scannet# user-filter 12 permit * 6 25 rate-limit 600 400 pps
```

Displaying User Filters

You can display the user filters associated with a zone configuration by entering the **show** command or the **show running-config** command in zone configuration mode.



Tip

To display the user filter configuration at the beginning of the display, use the **show** command or the **show running-config** command with the **begin USER FILTERS** option.

Table 6-13 describes the user filter fields in the **show** command output.

Table 6-13 Field Descriptions for User Filter Fields in the show Command

Field	Description
Row	User filter priority.
Source IP	Source IP address of the traffic that the filter processes.
Source Mask	Source address mask of the traffic that the filter processes.
Proto	Protocol number of the traffic that the filter processes.
DPort	Destination port of the traffic that the filter processes.
Frg	Type of traffic that the filter processes. The type can be one of the following: <ul style="list-style-type: none"> • yes—The filter processes fragmented traffic. • no—The filter processes nonfragmented traffic. • any—The filter processes fragmented and nonfragmented traffic.
RxRate (pps)	Current traffic rate in packets per second that is measured for this filter.
Action	Action that the filter performs on the specific traffic type. See Table 6-10 for more information.
Rate	Limit on the traffic rate that the user filter can handle. The rate is displayed in the units specified by the <i>Units</i> field.
Burst	Traffic burst limit that the filter allows for the specific flow. The units are bits, kilobits, kilopackets, megabits, and packets, and correspond to the units specified in the <i>Units</i> field.
Units	Units by which the rate and the burst rate are displayed.

The source IP address, source address mask, protocol number, and destination port may be nonspecific. An asterisk (*) indicates that the filter acts on all field values or that more than one value was matched for the filter.

Deleting User Filters



Caution

If you delete all user filters when the policy action is set to **to-user-filter**, the Guard passes unprotected traffic to the zone. See the [“Configuring the Policy Action”](#) section on page 7-20 for more information.

To delete a user filter, enter the following command in zone configuration mode:

```
no user-filter row-num
```

The *row-num* argument specifies the user filter row number. To display the list of user filters and identify the row number of the user filter to delete, use the **show running-config** command (see the [“Displaying User Filters”](#) section on page 6-17). To delete all user filters, enter an asterisk (*) for the row number.

The following example shows how to delete all user filters:

```
user@GUARD-conf-zone-scannet# no user-filter *
```

Configuring Dynamic Filters

Dynamic filters apply the required protection level to traffic flow and define how the Guard mitigates the attack. The Guard creates dynamic filters when it identifies an anomaly in the zone traffic, which occurs when the flow exceeds the zone policy thresholds. The Guard creates new dynamic filters as changes occur to the zone traffic and the type of DDoS attack.

Dynamic filters have a limited life span and are deleted by the Guard when the attack ends. The Guard supports a maximum of 150,000 dynamic filters that are active concurrently in all zones. You can add or delete dynamic filters when the zone is under attack and zone protection is enabled.

When the Guard detects a traffic anomaly, it uses both user filters and dynamic filters to mitigate the attack. The user filters provide the first line of defense until the Guard can analyze the attack and begin creating dynamic filters with mitigation actions designed specifically for the attack. For more information about user filters and how the Guard uses them in combination with the dynamic filters, see the [“Configuring User Filters”](#) section on page 6-13.

The dynamic filters and the user filters can take actions in these descending severity levels: drop, strong, basic, and permit (see [Table 6-14](#)). Dynamic filters with actions of redirect/zombie and block-unauthenticated are applied even if a user filter to handle the same type of traffic exists because they affect the Guard authentication functions and do not directly affect the traffic flow.

[Table 6-14](#) describes the different actions dynamic filters can execute.

Table 6-14 Dynamic Filter Actions

Action	Description
drop	Drops the traffic.
strong	Applies the strong protection level anti-spoofing functions to the specific traffic.
to-user-filters	Forwards the traffic to the user filters. If you have modified the default user filters, you must make sure that there is a user filter to handle these dynamic filters.

Table 6-14 Dynamic Filter Actions (continued)

Action	Description
block-unauthenticated-basic	Enhances the basic protection level anti-spoofing functions so that they drop traffic flows that have not been authenticated.
block-unauthenticated-strong	Enhances the strong protection level anti-spoofing functions so that they drop traffic flows that have not been authenticated.
block-unauthenticated-dns	Drops the traffic that flows to DNS UDP servers (protocol=UDP, port=53) that were not authenticated by the DNS anti-spoofing functions.
redirect/zombie	Enhances authentication for all user filters with an action of basic/redirect .

Dynamic filters are configured to remain active for a specific amount of time. Depending on how the filter was created, the dynamic filter timeout parameter is configured in one of the following ways:

- Dynamic filters created by a zone policy—The dynamic filter timeout is set to the policy timeout. To modify the timeout of additional dynamic filters that are created by the policy, change the timeout of the policy that created the dynamic filter by entering the **timeout** command in policy configuration mode.
- User-defined dynamic filters—You define the dynamic filter timeout by configuring the *exp-time* argument of the **dynamic-filter** command.

When the dynamic filter timeout expires, the Guard determines whether or not the dynamic filter should be deactivated based on current traffic conditions. If the Guard determines that the dynamic filter should not be deactivated, the filter remains active for another time span. See the “[Deactivating Dynamic Filters](#)” section on page 6-23 for more information about deactivating dynamic filters.

This section contains the following topics:

- [Displaying Dynamic Filters](#)
- [Adding Dynamic Filters](#)
- [Deleting Dynamic Filters](#)
- [Preventing the Production of Dynamic Filters](#)
- [Deactivating Dynamic Filters](#)

Displaying Dynamic Filters

You can display the dynamic filters that the Guard created by using one of the following commands in zone configuration mode:

- **show dynamic-filters [details]**—Displays a list of all dynamic filters.
- **show dynamic-filters *dynamic-filter-id* [details]**—Displays a single dynamic filter.
- **show dynamic-filters sort {action | exp-time | id | filter-rate}**—Displays a sorted list of all dynamic filters.

Table 6-15 provides the arguments and keywords for the **show dynamic-filters** command.

Table 6-15 Arguments and Keywords for the show dynamic-filters Command

Parameter	Description
<i>dynamic-filter-id</i>	Identifier of the specific dynamic filter to display. This integer is assigned by the Guard. To identify the filter ID, display the complete list of dynamic filters.
details	(Optional) Displays dynamic filters in detail. The details consist of additional information on the attack flow, the triggering rate, and the policy that produced it.
action	Displays dynamic filters by their action, ranging from the most severe (drop) to the least severe (notify).
exp-time	Displays dynamic filters by their expiration time in ascending order.
id	Displays dynamic filters by the ascending ID number.
filter-rate	Displays dynamic filters by the triggering rate, measured in packets per second, in ascending order.

To display the pending dynamic filters, use the **show recommendations** command. See [Chapter 10, “Using Interactive Protect Mode,”](#) for more information about pending dynamic filters.



Note

The Guard displays a maximum of 1000 dynamic filters. When more than 1000 dynamic filters are active, examine the log file or zone report for a complete list of dynamic filters.

The following example shows how to display a dynamic filter in detail:

```
user@GUARD-conf-zone-scannet# show dynamic-filters 876 details
```

Table 6-16 describes the fields in the **show dynamic-filters** command output.

Table 6-16 Field Descriptions for show dynamic-filters Command Output

Field	Description
ID	Filter identification number.
Action	Action that the filter performs on the traffic flow. See Table 6-14 for more information.
Exp Time	Amount of time that the filter is active. After the time expires, the filter may be deleted according to the thresholds that you defined by using the filter-termination command.
Source IP	Source IP address of the traffic that the filter processes.
Source Mask	Source address mask of the traffic that the filter processes.
Proto	Protocol number of the traffic that the filter processes.
DPort	Destination port of the traffic that the filter processes.

Table 6-16 Field Descriptions for show dynamic-filters Command Output (continued)

Field	Description
Frg	Whether or not the filter processes fragmented traffic: <ul style="list-style-type: none"> • yes—The filter processes fragmented traffic. • no—The filter processes nonfragmented traffic. • any—The filter processes both fragmented and nonfragmented traffic.
RxRate (pps)	Current traffic rate in packets per seconds that is measured for this filter.

The source IP address, source address mask, protocol number, and destination port may be nonspecific. An asterisk (*) indicates that the filter acts on all field values or that more than one value was matched for the filter.

Table 6-17 describes the additional fields in the **show dynamic-filters details** command output.

Table 6-17 Field Descriptions for show dynamic-filters details Command

Field	Description
Attack flow	Mitigated attack flow characteristics. The mitigated attack flow, displayed in the dynamic filters table, might have a wider range than the attack flow. For example, a nonspoofed attack on port 80 blocks all TCP traffic from the originating source IP address, not from port 80 only. The attack flow contains the Source IP, Source Mask, Proto, DPort, and Frg fields that are described in Table 6-16.
Triggering Rate	Rate of the attack flow that exceeded a policy threshold.
Threshold	Policy threshold that was exceeded by the attack flow.
Policy	Policy that produced the dynamic filter. See Chapter 7, “Configuring Policy Templates and Policies,” for more information.

Adding Dynamic Filters

During an attack on the zone, you can add a dynamic filter to manipulate zone protection by using the following command in zone configuration mode:

```
dynamic-filter action { exp-time | forever } src-ip [ip-mask] protocol dest-port [fragments-type]
```

You can use multiple **dynamic-filter** commands to add multiple dynamic filters.

Table 6-18 provides the arguments and keywords for the **dynamic-filter** command.

Table 6-18 Arguments and Keywords for the dynamic-filter Command

Parameter	Description
<i>action</i>	Action that the filter performs on a specific traffic flow. See Table 6-14 for more information.
<i>exp-time</i>	Integer from 1 to 3,000,000 that specifies the time (in seconds) for the filter to be active.
forever	Activates the filter for an unlimited time. The filter is deleted when protection ends.

Table 6-18 Arguments and Keywords for the dynamic-filter Command (continued)

Parameter	Description
<i>src-ip</i>	Traffic from a specific source IP address. Enter the IP address in dotted-decimal notation (for example, enter an IP address of 192.168.100.1). Use an asterisk (*) to indicate any IP address.
<i>ip-mask</i>	(Optional) Traffic from a specific subnet. Enter the subnet mask in dotted-decimal notation (for example, enter 255.255.255.0). The subnet mask can contain only Class C values. The default subnet is 255.255.255.255.
<i>protocol</i>	Traffic from a specific protocol. Use an asterisk (*) to specify any protocol. Review possible protocol numbers at the Internet Assigned Numbers Authority (IANA) website: http://www.iana.org/assignments/protocol-numbers
<i>dest-port</i>	Traffic that is destined to a specific destination port. Use an asterisk (*) to specify any destination port. Review possible port numbers at the Internet Assigned Numbers Authority (IANA) website: http://www.iana.org/assignments/port-numbers
<i>fragments-type</i>	(Optional) Traffic type that the filter acts on. The three fragmented types are as follows: <ul style="list-style-type: none"> • no-fragments—nonfragmented traffic • fragments—fragmented traffic • any-fragments—fragmented and nonfragmented traffic The default is no-fragments .

The following example shows how to add a dynamic filter that directs the traffic to the user filters with an expiration time of 600 seconds:

```
admin@GUARD-conf-zone-scannet# dynamic-filter to-user-filters 600 192.128.30.45
255.255.255.252 6 88 no-fragments
```

Deleting Dynamic Filters

When you delete dynamic filters, the deletion is effective for a limited period of time because the Guard continues to configure new dynamic filters when zone protection is enabled. See the “[Preventing the Production of Dynamic Filters](#)” section on page 6-23 for information on how to prevent the Guard from producing a dynamic filter.

To delete a dynamic filter, use the following command in zone configuration mode:

```
no dynamic-filter dynamic-filter-id
```

The *dynamic-filter-id* argument specifies the dynamic filter ID. To display the list of dynamic filters and identify the ID of the dynamic filter to delete, use the **show dynamic-filters** command (see the “[Displaying Dynamic Filters](#)” section on page 6-19). To delete all zone dynamic filters, enter an asterisk (*) for the dynamic filter identifier.

The following example shows how to delete a dynamic filter:

```
user@GUARD-conf-zone-scannet# no dynamic-filter 876
```

Preventing the Production of Dynamic Filters

To prevent the Guard from producing unwanted dynamic filters, perform one of the following actions:

- Deactivate the policy that produces the dynamic filters (see the [“Changing the Policy State”](#) section on page 7-13 for more information). To determine which policy produced the unwanted dynamic filters, see the [“Displaying Dynamic Filters”](#) section on page 6-19.
- Configure a bypass filter for the desired traffic flow (see the [“Configuring Bypass Filters”](#) section on page 6-11).
- Increase the threshold of the policy that produces the undesired dynamic filter (see the [“Configuring the Policy Threshold”](#) section on page 7-14).

Deactivating Dynamic Filters

When the dynamic filter timeout expires, the Guard determines whether or not the dynamic filter should be deactivated based on current traffic conditions. If the Guard determines that the dynamic filter should not be deactivated, the filter remains active for another time span.

Dynamic filters are deactivated if one of the following conditions applies:

- The total zone malicious traffic rate, which equals the sum of the spoofed and dropped traffic, is less than or equal to the zone-malicious-rate termination threshold. See the following commands in this section.
- The dynamic filter measures the traffic rate (the filter rate counter does not display N/A) and the filter-rate termination threshold (see the following commands in this section) is equal to or greater than both of the following:
 - The dynamic filter current traffic rate.
 - The dynamic filter average traffic rate during a user-configured time span. This time span is defined by the policy timeout parameter. See the [“Configuring the Policy Timeout”](#) section on page 7-19 for more information.



Note Dynamic filters with an action of to-user-filters, block-unauthenticated, redirect/zombie, or notify do not measure the traffic rate.

To configure the zone malicious traffic threshold, use the following command in zone configuration mode:

```
filter-termination zone-malicious-rate threshold
```

The *threshold* argument specifies the zone malicious traffic threshold for a zone in packets per second (pps). This traffic consists of the sum of the spoofed and the dropped traffic. The default value is 50 pps.

To configure the dynamic filter-rate termination threshold, use the following command in zone configuration mode:

```
filter-termination { filter-rate | filter-rate-pph } threshold
```

Table 6-19 provides the arguments and keywords for the **filter-termination** command.

Table 6-19 Arguments and Keywords for the filter-termination Command

Parameter	Description
filter-rate	Specifies a filter rate of packets per second.
filter-rate-pph	Specifies a filter rate of packets per hour for filters created by PPH policies (see the “Understanding Zone Policies” section on page 7-1).
<i>threshold</i>	Dynamic filter traffic threshold in either packets per second or packets per hour depending on the keyword that you select. The default values are as follows: <ul style="list-style-type: none">• 2 pps• 120 pph

The following example shows how to configure the dynamic filter termination rates:

```
user@GUARD-conf-zone-scannet# filter-termination zone-malicious-rate 200  
user@GUARD-conf-zone-scannet# filter-termination filter-rate 50
```