



Detecting Specific Threats

You can use several preprocessors in a network analysis policy to detect specific threats to your monitored network, such as Back Orifice attacks, several portscan types, and rate-based attacks that attempt to overwhelm your network with excessive traffic. Note that when an intrusion rule or rule argument requires a disabled preprocessor, the system automatically uses it with its current configuration even though it remains disabled in the network analysis policy's user interface. For more information, see [Limitations of Custom Policies, page 18-11](#).

You can also use sensitive data detection, which you configure in an intrusion policy, to detect unsecured transmission of sensitive numerical data.

See the following sections for more information on detecting specific threats:

- [Detecting Back Orifice, page 28-1](#) explains detection of Back Orifice attacks.
- [Detecting Portscans, page 28-3](#) describes the different types of portscans and explains how you can use portscan detection to identify threats to your networks before they develop into attacks.
- [Preventing Rate-Based Attacks, page 28-9](#) explains how to limit denial of service (DoS) and SYN flood attacks.
- [Detecting Sensitive Data, page 28-19](#) explains how to detect and generate events on sensitive data such as credit card numbers and Social Security numbers in ASCII text.

Detecting Back Orifice

License: Protection

The ASA FirePOWER module provides a preprocessor that detects the existence of the Back Orifice program. This program can be used to gain admin access to your Windows hosts. The Back Orifice preprocessor analyzes UDP traffic for the Back Orifice magic cookie, "`*!*QWTY?`", which is located in the first eight bytes of the packet and is XOR-encrypted.

The Back Orifice preprocessor has a configuration page, but no configuration options. When it is enabled, you must also enable the preprocessor rules in the following table for the preprocessor to generate corresponding events.

Table 28-1 **Back Orifice GID:SDs**

Preprocessor rule GID:SID	Description
105:1	Back Orifice traffic detected
105:2	Back Orifice client traffic detected
105:3	Back Orifice server traffic detected
105:4	Back Orifice snort buffer attack detected

To view the Back Orifice Detection page:

-
- Step 1** Select **Configuration > ASA FirePOWER Configuration > Policies > Access Control Policy**.
The Access Control Policy page appears.
- Step 2** Click the edit icon () next to the access control policy you want to edit.
The access control policy editor appears.
- Step 3** Select the **Advanced** tab.
The access control policy advanced settings page appears.
- Step 4** Click the edit icon () next to **Network Analysis and Intrusion Policies**.
The Network Analysis and Intrusion Policies pop-up window appears.
- Step 5** Click **Network Analysis Policy List**.
The Network Analysis Policy List pop-up window appears.
- Step 6** Click the edit icon () next to the policy you want to edit.
If you have unsaved changes in another policy, click **OK** to discard those changes and continue. See [Resolving Conflicts and Committing Policy Changes, page 18-15](#) for information on saving unsaved changes in another policy.
The Policy Information page appears.
- Step 7** In the navigation panel on the left, click **Settings**.
The Settings page appears.
- Step 8** You have two choices, depending on whether **Back Orifice Detection** under **Specific Threat Detection** is enabled:
- If the preprocessor is enabled, click **Edit**.
 - If the preprocessor is disabled, click **Enabled**, then click **Edit**.
- The Back Orifice Detection page appears. A message at the bottom of the page identifies the intrusion policy layer that contains the configuration. See [Using Layers in a Network Analysis or Intrusion Policy, page 19-1](#) for more information.
- Step 9** Save your policy, continue editing, discard your changes, or exit while leaving your changes in the system cache. See [Resolving Conflicts and Committing Policy Changes, page 18-15](#) for more information.
-

Detecting Portscans

License: Protection

A portscan is a form of network reconnaissance that is often used by attackers as a prelude to an attack. In a portscan, an attacker sends specially crafted packets to a targeted host. By examining the packets that the host responds with, the attacker can often determine which ports are open on the host and, either directly or by inference, which application protocols are running on these ports.

Note that when portscan detection is enabled, you must enable rules on the intrusion policy Rules page with generator ID (GID) 122 for enabled portscan types for the portscan detector to generate portscan events. See [Setting Rule States, page 27-19](#) and [Table 28-5 on page 28-7](#) for more information.

By itself, a portscan is not evidence of an attack. In fact, some of the portscanning techniques used by attackers can also be employed by legitimate users on your network. Cisco's portscan detector is designed to help you determine which portscans might be malicious by detecting patterns of activity.

Attackers are likely to use several methods to probe your network. Often they use different protocols to draw out different responses from a target host, hoping that if one type of protocol is blocked, another may be available. The following table describes the protocols you can activate in the portscan detector.

Table 28-2 **Protocol Types**

Protocol	Description
TCP	Detects TCP probes such as SYN scans, ACK scans, TCP connect() scans, and scans with unusual flag combinations such as Xmas tree, FIN, and NULL
UDP	Detects UDP probes such as zero-byte UDP packets
ICMP	Detects ICMP echo requests (pings)
IP	Detects IP protocol scans. These scans differ from TCP and UDP scans because the attacker, instead of looking for open ports, is trying to discover which IP protocols are supported on a target host.

**Note**

For events generated by the portscan connection detector, the protocol number is set to 255. Because portscan does not have a specific protocol associated with it by default, the Internet Assigned Numbers Authority (IANA) does not have a protocol number assigned to it. IANA designates 255 as a reserved number, so that number is used in portscan events to indicate that there is not an associated protocol for the event.

Portscans are generally divided into four types based on the number of targeted hosts, the number of scanning hosts, and the number of ports that are scanned. The following table describes the kinds of portscan activity you can detect.

Table 28-3 Portscan Types

Type	Description
Portscan Detection	A one-to-one portscan in which an attacker uses one or a few hosts to scan multiple ports on a single target host. One-to-one portscans are characterized by: • a low number of scanning hosts • a single host that is scanned • a high number of ports scanned This option detects TCP, UDP, and IP portscans.
Port Sweep	A one-to-many portsweep in which an attacker uses one or a few hosts to scan a single port on multiple target hosts. Portsweeps are characterized by: • a low number of scanning hosts • a high number of scanned hosts • a low number of unique ports scanned This option detects TCP, UDP, ICMP, and IP portsweeps.
Decoy Portscan	A one-to-one portscan in which the attacker mixes spoofed source IP addresses with the actual scanning IP address. Decoy portscans are characterized by: • a high number of scanning hosts • a low number of ports that are scanned only once • a single (or a low number of) scanned hosts The decoy portscan option detects TCP, UDP, and IP protocol portscans.
Distributed Portscan	A many-to-one portscan in which multiple hosts query a single host for open ports. Distributed portscans are characterized by: • a high number of scanning hosts • a high number of ports that are scanned only once • a single (or a low number of) scanned hosts The distributed portscan option detects TCP, UDP, and IP protocol portscans.

The information that the portscan detector learns about a probe is largely based on seeing negative responses from the probed hosts. For example, when a web client tries to connect to a web server, the client uses port 80/tcp and the server can be counted on to have that port open. However, when an attacker probes a server, the attacker does not know in advance if it offers web services. When the portscan detector sees a negative response (that is, an ICMP unreachable or TCP RST packet), it records the response as a potential portscan. The process is more difficult when the targeted host is on the other side of a device such as a firewall or router that filters negative responses. In this case, the portscan detector can generate *filtered* portscan events based on the sensitivity level that you select.

The following table describes the three different sensitivity levels you can choose from.

Table 28-4 **Sensitivity Levels**

Level	Description
Low	Detects only negative responses from targeted hosts. Select this sensitivity level to suppress false positives, but keep in mind that some types of portscans (slow scans, filtered scans) might be missed. This level uses the shortest time window for portscan detection.
Medium	Detects portscans based on the number of connections to a host, which means that you can detect filtered portscans. However, very active hosts such as network address translators and proxies may generate false positives. Note that you can add the IP addresses of these active hosts to the Ignore Scanned field to mitigate this type of false positive. This level uses a longer time window for portscan detection.
High	Detects portscans based on a time window, which means that you can detect time-based portscans. However, if you use this option, you should be careful to tune the detector over time by specifying IP addresses in the Ignore Scanned and Ignore Scanner fields. This level uses a much longer time window for portscan detection.

See the following sections for more information:

- [Configuring Portscan Detection, page 28-5](#)
- [Understanding Portscan Events, page 28-7](#)

Configuring Portscan Detection

License: Protection

The portscan detection configuration options allow you to finely tune how the portscan detector reports scan activity.

Note that when portscan detection is enabled, you must enable rules on the Rules page with generator ID (GID) 122 for enabled portscan types for the portscan detector to generate portscan events. See [Setting Rule States, page 27-19](#) and the [Portscan Detection SIDs \(GID:122\)](#) table for more information.

To configure portscan detection:

Admin/Intrusion Admin

-
- Step 1** Select **Configuration > ASA FirePOWER Configuration > Policies > Access Control Policy**.
The Access Control Policy page appears.
- Step 2** Click the edit icon (✎) next to the access control policy you want to edit.
The access control policy editor appears.
- Step 3** Select the **Advanced** tab.
The access control policy advanced settings page appears.
- Step 4** Click the edit icon (✎) next to **Network Analysis and Intrusion Policies**.
The Network Analysis and Intrusion Policies pop-up window appears.

Step 5 Click **Network Analysis Policy List**.

The Network Analysis Policy List pop-up window appears.

Step 6 Click the edit icon (✎) next to the policy you want to edit.

If you have unsaved changes in another policy, click **OK** to discard those changes and continue. See [was Committing Intrusion Policy Changes; update xref] for information on saving unsaved changes in another policy.

The Policy Information page appears.

Step 7 In the navigation panel on the left, click **Settings**.

The Settings page appears.

Step 8 You have two choices, depending on whether **Portscan Detection** under **Specific Threat Detection** is enabled:

- If the configuration is enabled, click **Edit**.
- If the configuration is disabled, click **Enabled**, then click **Edit**.

The Portscan Detection page appears. A message at the bottom of the page identifies the intrusion policy layer that contains the configuration. See [Using Layers in a Network Analysis or Intrusion Policy, page 19-1](#) for more information.

Step 9 In the **Protocol** field, specify which of the following protocols you want to enable:

- TCP
- UDP
- ICMP
- IP

Use Ctrl or Shift while clicking to select multiple protocols or clear individual protocols. See the [Protocol Types](#) table for more information.

Note that you must ensure TCP stream processing is enabled to detect scans over TCP, and that UDP stream processing is enabled to detect scans over UDP.

Step 10 In the **Scan Type** field, specify which of the following portscans you want to detect:

- Portscan Detection
- Port Sweep
- Decoy Portscan
- Distributed Portscan

Use Ctrl or Shift while clicking to select or deselect multiple protocols. See the [Portscan Types](#) table for more information.

Step 11 In the **Sensitivity Level** list, select the level you want to use: low, medium, or high.

See the [Sensitivity Levels](#) table for more information.

Step 12 Optionally, in the **Watch IP** field, specify which host you want to watch for signs of portscan activity, or leave the field blank to watch all network traffic.

You can specify a single IP address or address block, or a comma-separated lists of either or both. For information on using IPv4 and IPv6 address blocks, see [IP Address Conventions, page 1-4](#).

Step 13 Optionally, in the **Ignore Scanners** field, specify which hosts you want to ignore as scanners. Use this field to indicate hosts on your network that are especially active. You may need to modify this list of hosts over time.

You can specify a single IP address or address block, or a comma-separated lists of either or both. For information on using IPv4 and IPv6 address blocks, see [IP Address Conventions, page 1-4](#).

- Step 14** Optionally, in the **Ignore Scanned** field, specify which hosts you want to ignore as the target of a scan. Use this field to indicate hosts on your network that are especially active. You may need to modify this list of hosts over time.

You can specify a single IP address or address block, or a comma-separated lists of either or both. For information on using IPv4 and IPv6 address blocks, see [IP Address Conventions, page 1-4](#).

- Step 15** Optionally, clear the **Detect Ack Scans** check box to discontinue monitoring of sessions picked up in mid-stream.



Note

Detection of mid-stream sessions helps to identify ACK scans, but may cause false events, particularly on networks with heavy traffic and dropped packets.

- Step 16** Save your policy, continue editing, discard your changes, or exit while leaving your changes in the system cache. See [Resolving Conflicts and Committing Policy Changes, page 18-15](#) for more information.

Understanding Portscan Events

License: Protection

When portscan detection is enabled, you must enable rules with generator ID (GID) 122 and a Snort® ID (SID) from among SIDs 1 through 27 to generate events for each enabled portscan type. See [Setting Rule States, page 27-19](#) for more information. The **Preprocessor Rule SID** column in the following table lists the SID for the preprocessor rule you must enable for each portscan type.

Table 28-5 Portscan Detection SIDs (GID:122)

Portscan Type	Protocol:	Sensitivity Level	Preprocessor Rule SID
Portscan Detection	TCP	Low	1
		Medium or High	5
	UDP	Low	17
		Medium or High	21
	ICMP	Low	Does not generate events.
		Medium or High	Does not generate events.
	IP	Low	9
		Medium or High	13
Port Sweep	TCP	Low	3, 27
		Medium or High	7
	UDP	Low	19
		Medium or High	23
	ICMP	Low	25
		Medium or High	26
	IP	Low	11
		Medium or High	15

Table 28-5 Portscan Detection SIDs (GID:122) (continued)

Portscan Type	Protocol:	Sensitivity Level	Preprocessor Rule SID
Decoy Portscan	TCP	Low	2
		Medium or High	6
	UDP	Low	18
		Medium or High	22
	ICMP	Low	Does not generate events.
		Medium or High	Does not generate events.
	IP	Low	10
		Medium or High	14
Distributed Portscan	TCP	Low	4
		Medium or High	8
	UDP	Low	20
		Medium or High	24
	ICMP	Low	Does not generate events.
		Medium or High	Does not generate events.
	IP	Low	12
		Medium or High	16

When you enable the accompanying preprocessor rules, the portscan detector generates intrusion events that you can view just as you would any other intrusion event. However, the information presented on the packet view is different from the other types of intrusion events. This section describes the fields that appear on the packet view for a portscan event and how you can use that information to understand the types of probes that occur on your network.

Begin by using the intrusion event views to drill down to the packet view for a portscan event.

Note that you cannot download a portscan packet because single portscan events are based on multiple packets; however, the portscan packet view provides all usable packet information.

**Note**

For events generated by the portscan connection detector, the protocol number is set to 255. Because portscan does not have a specific protocol associated with it by default, the Internet Assigned Numbers Authority (IANA) does not have a protocol number assigned to it. IANA designates 255 as a reserved number, so that number is used in portscan events to indicate that there is not an associated protocol for the event.

The following table describes the information provided in the packet view for portscan events.

Table 28-6 Portscan Packet View

Information	Description
Device	The device that detected the event.
Time	The time when the event occurred.
Message	The event message generated by the preprocessor.
Source IP	The IP address of the scanning host.
Destination IP	The IP address of the scanned host.

Table 28-6 *Portscan Packet View (continued)*

Information	Description
Priority Count	The number of negative responses (for example, TCP RSTs and ICMP unreachable) from the scanned host. The higher the number of negative responses, the higher the priority count.
Connection Count	The number of active connections on the hosts. This value is more accurate for connection-based scans such as TCP and IP.
IP Count	The number of times that the IP addresses that contact the scanned host changes. For example, if the first IP address is 10.1.1.1, the second IP is 10.1.1.2, and the third IP is 10.1.1.1, then the IP count is 3. This number is less accurate for active hosts such as proxies and DNS servers.
Scanner/Scanned IP Range	The range of IP addresses for the scanned hosts or the scanning hosts, depending on the type of scan. For portsweeps, this field shows the IP range of scanned hosts. For portscans, this shows the IP range of the scanning hosts.
Port/Proto Count	For TCP and UDP portscans, the number of times that the port being scanned changes. For example, if the first port scanned is 80, the second port scanned is 8080, and the third port scanned is again 80, then the port count is 3. For IP protocol portscans, the number of times that the protocol being used to connect to the scanned host changes.
Port/Proto Range	For TCP and UDP portscans, the range of the ports that were scanned. For IP protocol portscans, the range of IP protocol numbers that were used to attempt to connect to the scanned host.
Open Ports	The TCP ports that were open on the scanned host. This field appears only when the portscan detects one or more open ports.

Preventing Rate-Based Attacks

License: Protection

Rate-based attacks are attacks that depend on frequency of connection or repeated attempts to perpetrate the attack. You can use rate-based detection criteria to detect a rate-based attack as it occurs and respond to it when it happens, then return to normal detection settings after it stops. For more information on configuring rate-based detection, see the following topics:

- [Understanding Rate-Based Attack Prevention, page 28-9](#)
- [Rate-Based Attack Prevention and Other Filters, page 28-12](#)
- [Configuring Rate-Based Attack Prevention, page 28-17](#)
- [Understanding Dynamic Rule States, page 27-28](#)
- [Setting a Dynamic Rule State, page 27-29](#)

Understanding Rate-Based Attack Prevention

License: Protection

You can configure your network analysis policy to include rate-based filters that detect excessive activity directed at hosts on your network. You can use this feature on a device deployed in inline mode to block rate-based attacks for a specified time, then revert to only generating events and not drop traffic.

Rate-based attack prevention identifies abnormal traffic patterns and attempts to minimize the impact of that traffic on legitimate requests. Rate-based attacks usually have one of the following characteristics:

- any traffic containing excessive incomplete connections to hosts on the network, indicating a SYN flood attack

To configure SYN attack detection, see [Preventing SYN Attacks, page 28-11](#).

- any traffic containing excessive complete connections to hosts on the network, indicating a TCP/IP connection flood attack

To configure simultaneous connection detection, see [Controlling Simultaneous Connections, page 28-12](#).

- excessive rule matches in traffic going to a particular destination IP address or addresses or coming from a particular source IP address or addresses.

To configure source or destination-based dynamic rule states, see [Setting a Dynamic Rule State, page 27-29](#).

- excessive matches for a particular rule across all traffic.

To configure rule-based dynamic rule states, see [Setting a Dynamic Rule State, page 27-29](#).

In a network analysis policy, you can either configure SYN flood or TCP/IP connection flood detection for the entire policy; in an intrusion policy, you can set rate-based filters for individual intrusion or preprocessor rules. Note that manually adding a rate-based filter to rules 135:1 and 135:2 has no effect. Rules with GID:135 use the client as the source value and the server as the destination value. See [Preventing SYN Attacks, page 28-11](#) and [Controlling Simultaneous Connections, page 28-12](#) for more information.

Each rate-based filter contains several components:

- for policy-wide or rule-based source or destination settings, the network address designation
- the rule matching rate, which you configure as a count of rule matches within a specific number of seconds
- a new action to be taken when the rate is exceeded

When you set a rate-based setting for the entire policy, the system generates events when it detects a rate-based attack, and optionally can drop the traffic in an inline deployment. When setting rate-based actions for individual rules, you have three available actions: Generate Events, Drop and Generate Events, and Disable.

- the duration of the action, which you configure as a timeout value

Note that when started, the new action occurs until the timeout is reached, even if the rate falls below the configured rate during that time period. When the timeout period expires, if the rate has fallen below the threshold, the action for the rule reverts to the action initially configured for the rule. For policy-wide settings, the action reverts to the action of each rule the traffic matches or stops if it does not match any rules.

You can configure rate-based attack prevention in an inline deployment to block attacks, either temporarily or permanently. Without rate-based configuration, rules set to Generate Events create events, but the system does not drop packets for those rules. However, if the attack traffic matches rules that have rate-based criteria configured, the rate action may cause packet dropping to occur for the period of time that the rate action is active, even if those rules are not initially set to Drop and Generate Events.


**Note**

Rate-based actions cannot enable disabled rules or drop traffic that matches disabled rules. However, if you set a rate-based filter at the policy level, you can generate events on or generate events on and drop traffic that contains an excessive number of SYN packets or SYN/ACK interactions within a designated time period.

You can define multiple rate-based filters on the same rule. The first filter listed in the intrusion policy has the highest priority. Note that when two rate-based filter actions conflict, the system implements the action of the first rate-based filter. Similarly, policy-wide rate-based filters override rate-based filters set on individual rules if the filters conflict.

The following diagram shows an example where an attacker is attempting to access a host. Repeated attempts to find a password trigger a rule which has rate-based attack prevention configured. The rate-based settings change the rule attribute to Drop and Generate Events after rule matches occur five times in a 10-second span. The new rule attribute times out after 15 seconds.

After the timeout, note that packets are still dropped in the rate-based sampling period that follows. If the sampled rate is above the threshold in the current or previous sampling period, the new action continues. The new action reverts to generating events only after a sampling period completes where the sampled rate is below the threshold rate.



Preventing SYN Attacks

License: Protection

The SYN attack prevention option helps you protect your network hosts against SYN floods. You can protect individual hosts or whole networks based on the number of packets seen over a period of time. If your device is deployed passively, you can generate events. If your device is placed inline, you can also drop the malicious packets. After the timeout period elapses, if the rate condition has stopped, the event generation and packet dropping stops.

For example, you could configure a setting to allow a maximum of 10 SYN packets from any one IP address, and block further connections from that IP address for 60 seconds.

Enabling this option also activates rule 135:1. Manually activating this rule has no effect. The rule state is always displayed as Disabled, and never changes. The rule generates events when this option is enabled and a defined rate condition is exceeded.

Controlling Simultaneous Connections

License: Protection

You can also limit TCP/IP connections to or from hosts on your network to prevent denial of service (DoS) attacks or excessive activity by users. When the system detects the configured number of successful connections to or from a specified IP address or range of addresses, it generates events on additional connections. The rate-based event generation continues until the timeout period elapses without the rate condition occurring. In an inline deployment you can choose to drop packets until the rate condition times out.

For example, you could configure a setting to allow a maximum of 10 successful simultaneous connections from any one IP address, and block further connections from that IP address for 60 seconds.

Enabling this option also activates rule 135:2. Manually activating this rule has no effect. The rule state is always displayed as Disabled, and never changes. The rule generates events when this option is enabled and a defined rate condition is exceeded.

Rate-Based Attack Prevention and Other Filters

License: Protection

The `detection_filter` keyword and the thresholding and suppression features provide other ways to filter either the traffic itself or the events that the system generates. You can use rate-based attack prevention alone or in any combination with thresholding, suppression, or the `detection_filter` keyword.

See the following examples for more information:

- [Rate-Based Attack Prevention and Detection Filtering, page 28-12](#)
- [Dynamic Rule States and Thresholding or Suppression, page 28-13](#)
- [Policy-Wide Rate-Based Detection and Thresholding or Suppression, page 28-15](#)
- [Rate-Based Detection with Multiple Filtering Methods, page 28-16](#)

Rate-Based Attack Prevention and Detection Filtering


License: Protection

The `detection_filter` keyword prevents a rule from triggering until a threshold number of rule matches occur within a specified time. When a rule includes the `detection_filter` keyword, the system tracks the number of incoming packets matching the pattern in the rule per timeout period. The system can count hits for that rule from particular source or destination IP addresses. After the rate exceeds the rate in the rule, event notification for that rule begins.

The following example shows an attacker attempting a brute-force login. Repeated attempts to find a password trigger a rule that also includes the `detection_filter` keyword, with a count set to 5. This rule has rate-based attack prevention configured. The rate-based settings change the rule attribute to Drop and Generate Events for 20 seconds when there are five hits on the rule in a 10-second span.

As shown in the diagram, the first five packets matching the rule do not generate events because the rule does not trigger until the rate exceeds the rate indicated by the `detection_filter` keyword. After the rule triggers, event notification begins, but the rate-based criteria do not trigger the new action of Drop and Generate Events until five more packets pass.

After the rate-based criteria are met, events are generated and the packets are dropped until the rate-based timeout period expires and the rate falls below the threshold. After twenty seconds elapse, the rate-based action times out. After the timeout, note that packets are still dropped in the rate-based sampling period that follows. Because the sampled rate is above the threshold rate in the previous sampling period when the timeout happens, the rate-based action continues.



Note that although the example does not depict this, you can use the Drop and Generate Events rule state in combination with the `detection_filter` keyword to start dropping traffic when hits for the rule reach the specified rate. When deciding whether to configure rate-based settings for a rule, consider whether setting the rule to Drop and Generate Events and including the `detection_filter` keyword would achieve the same result, or whether you want to manage the rate and timeout settings in the intrusion policy. For more information, see [Setting Rule States, page 27-19](#).

Dynamic Rule States and Thresholding or Suppression

License: Protection


You can use thresholding and suppression to reduce excessive events by limiting the number of event notifications for a rule or by suppressing notifications altogether for that rule. For more information on the available options for thresholding and suppression, see [Configuring Event Thresholding, page 27-21](#) and [Configuring Suppression Per Intrusion Policy, page 27-25](#).

If you apply suppression to a rule, the system suppresses event notifications for that rule for all applicable IP addresses even if a rate-based action change occurs. However, the interaction between thresholding and rate-based criteria is more complex.

The following example shows an attacker attempting a brute-force login. Repeated attempts to find a password trigger a rule that has rate-based attack prevention configured. The rate-based settings change the rule attribute to Drop and Generate Events for 15 seconds when there are five hits on the rule in 10 seconds. In addition, a limit threshold limits the number of events the rule can generate to 10 events in 23 seconds.

As shown in the diagram, the rule generates events for the first five matching packets. After five packets, the rate-based criteria trigger the new action of Drop and Generate Events, and for the next five packets the rule generates events and the system drops the packet. After the tenth packet, the limit threshold has been reached, so for the remaining packets the system does not generate events but does drop the packets.

After the timeout, note that packets are still dropped in the rate-based sampling period that follows. If the sampled rate is above the threshold rate in the current or previous sampling period, the new action continues. The new action reverts to Generate Events only after a sampling period completes where the sampled rate is below the threshold rate.



372203

Note that although it is not shown in this example, if a new action triggers because of rate-based criteria *after* a threshold has been reached, the system generates a single event to indicate the change in action. So, for example, when the limit threshold of 10 is reached and the system stops generating events and the action changes from Generate Events to Drop and Generate Events on the 14th packet, the system generates an eleventh event to indicate the change in action.

Policy-Wide Rate-Based Detection and Thresholding or Suppression

License: Protection


You can use thresholding and suppression to reduce excessive events by limiting the number of event notifications for a source or destination or by suppressing notifications altogether for that rule. For more information on the available options for thresholding and suppression, see [Configuring Global Thresholds, page 29-3](#), [Configuring Event Thresholding, page 27-21](#), and [Configuring Suppression Per Intrusion Policy, page 27-25](#).

If suppression is applied to a rule, event notifications for that rule for all applicable IP addresses are suppressed even if a rate-based action change occurs because of a policy-wide or rule-specific rate-based setting. However, the interaction between thresholding and rate-based criteria is more complex.

The following example shows an attacker attempting denial of service (DoS) attacks on hosts in your network. Many simultaneous connections to hosts from the same sources trigger a policy-wide Control Simultaneous Connections setting. The setting generates events and drops malicious traffic when there are five connections from one source in 10 seconds. In addition, a global limit threshold limits the number of events any rule or setting can generate to 10 events in 20 seconds.

As shown in the diagram, the policy-wide setting generates events for the first ten matching packets and drops the traffic. After the tenth packet, the limit threshold is reached, so for the remaining packets no events are generated but the packets are dropped.

After the timeout, note that packets are still dropped in the rate-based sampling period that follows. If the sampled rate is above the threshold rate in the current or previous sampling period, the rate-based action of generating events and dropping traffic continues. The rate-based action stops only after a sampling period completes where the sampled rate is below the threshold rate.



372200

Note that although it is not shown in this example, if a new action triggers because of rate-based criteria *after* a threshold has been reached, the system generates a single event to indicate the change in action. So, for example, if the limit threshold of 10 has been reached and the system stops generating events and the action changes to Drop and Generate events on the 14th packet, the system generates an eleventh event to indicate the change in action.

Rate-Based Detection with Multiple Filtering Methods


License: Protection

You may encounter situations where the `detection_filter` keyword, thresholding or suppression, and rate-based criteria all apply to the same traffic. When you enable suppression for a rule, events are suppressed for the specified IP addresses even if a rate-based change occurs.

The following example shows an attacker attempting a brute force login, and describes a case where a `detection_filter` keyword, rate-based filtering, and thresholding interact. Repeated attempts to find a password trigger a rule which includes the `detection_filter` keyword, with a count set to 5. This rule also has rate-based attack prevention settings that change the rule attribute to Drop and Generate Events for 30 seconds when there are five rule hits in 15 seconds. In addition, a limit threshold limits the rule to 10 events in 30 seconds.

As shown in the diagram, the first five packets matching the rule do not cause event notification because the rule does not trigger until the rate indicated in the `detection_filter` keyword is exceeded. After the rule triggers, event notification begins, but the rate-based criteria do not trigger the new action of Drop and Generate Events until five more packets pass. After the rate-based criteria are met, the system generates events for packets 11-15 and drops the packets. After the fifteenth packet, the limit threshold has been reached, so for the remaining packets the system does not generate events but does drop the packets.

After the rate-based timeout, note that packets are still dropped in the rate-based sampling period that follows. Because the sampled rate is above the threshold rate in the previous sampling period, the new action continues.



Configuring Rate-Based Attack Prevention

License: Protection

You can configure rate-based attack prevention at the policy level to stop SYN flood attacks. You can also stop excessive connections from a specific source or to a specific destination.

To configure rate-based attack prevention:

Admin/Intrusion Admin

- Step 1** Select **Configuration > ASA FirePOWER Configuration > Policies > Access Control Policy**.
The Access Control Policy page appears.
- Step 2** Click the edit icon (✎) next to the access control policy you want to edit.
The access control policy editor appears.
- Step 3** Select the **Advanced** tab.
The access control policy advanced settings page appears.

- Step 4** Click the edit icon (✎) next to **Network Analysis and Intrusion Policies**.
The Network Analysis and Intrusion Policies pop-up window appears.
- Step 5** Click **Network Analysis Policy List**.
The Network Analysis Policy List pop-up window appears.
- Step 6** Click the edit icon (✎) next to the policy you want to edit.
If you have unsaved changes in another policy, click **OK** to discard those changes and continue. See [Resolving Conflicts and Committing Policy Changes, page 18-15](#) for information on saving unsaved changes in another policy.
The Policy Information page appears.
- Step 7** In the navigation panel on the left, click **Settings**.
The Settings page appears.
- Step 8** You have two choices, depending on whether **Rate-Based Attack Prevention** under **Specific Threat Detection** is enabled:
- If the configuration is enabled, click **Edit**.
 - If the configuration is disabled, click **Enabled**, then click **Edit**.
- The Rate-Based Attack Prevention page appears. A message at the bottom of the page identifies the intrusion policy layer that contains the configuration. See [Using Layers in a Network Analysis or Intrusion Policy, page 19-1](#) for more information.
- Step 9** You have two options:
- To prevent incomplete connections intended to flood a host, click **Add** under **SYN Attack Prevention**.
The SYN Attack Prevention dialog box appears.
 - To prevent excessive numbers of connections, click **Add** under **Control Simultaneous Connections**.
The Control Simultaneous Connections dialog box appears.
- Step 10** Select how you want to track traffic:
- To track all traffic from a specific source or range of sources, select **Source** from the **Track By** drop-down list and type a single IP address or address block in the **Network** field.
 - To track all traffic to a specific destination or range of destinations, select **Destination** from the **Track By** drop-down list and type an IP address or address block in the **Network** field.
- Note that the system tracks traffic separately for each IP address included in the Network field. Traffic from an IP address that exceeds the configured rate results in generated events only for that IP address. As an example, you might set a source CIDR block of 10.1.0.0/16 for the network setting and configure the system to generate events when there are ten simultaneous connections open. If eight connections are open from 10.1.4.21 and six from 10.1.5.10, the system does not generate events, because neither source has the triggering number of connections open. However, if eleven simultaneous connections are open from 10.1.4.21, the system generates events only for the connections from 10.1.4.21.
- For information on using CIDR notation and prefix lengths, see [IP Address Conventions, page 1-4](#).
- Step 11** Indicate the triggering rate for the rate tracking setting:
- For SYN attack configuration, indicate the number of SYN packets per number of seconds in the **Rate** fields.
 - For simultaneous connection configuration, indicate the number of connections in the **Count** field.
- Step 12** To drop packets matching the rate-based attack prevention settings, select **Drop**.

- Step 13** In the **Timeout** field, indicate the time period after which to stop generating events, and if applicable, dropping, for traffic with the matching pattern of SYNs or simultaneous connections.

**Caution**

Timeout values can be integers from 1 to 1,000,000. However, setting a high timeout value may entirely block connection to a host in an inline deployment.

- Step 14** Save your policy, continue editing, discard your changes, or exit while leaving your changes in the system cache. See [Resolving Conflicts and Committing Policy Changes, page 18-15](#) for more information.

Detecting Sensitive Data

License: Protection

Sensitive data such as Social Security numbers, credit card numbers, driver's license numbers, and so on may be leaked onto the Internet, intentionally or accidentally. The system provides a sensitive data preprocessor that can detect and generate events on sensitive data in ASCII text, which can be particularly useful in detecting accidental data leaks.

The system does not detect encrypted or obfuscated sensitive data, or sensitive data in a compressed or encoded format such as a Base64-encoded email attachment. For example, the system would detect the phone number (555)123-4567, but not an obfuscated version where each number is separated by spaces, as in (5 5 5) 1 2 3 - 4 5 6 7, or by intervening HTML code, such as `<b>(555)</b>-<i>123-4567</i>`. However, the system would detect, for example, the HTML coded number `<b>(555)-123-4567</b>` where no intervening codes interrupt the numbering pattern.

**Tip**

The sensitive data preprocessor can detect sensitive data in unencrypted Microsoft Word files that are uploaded and downloaded using FTP or HTTP; this is possible because of the way Word files group ASCII text and formatting commands separately.

The system detects sensitive data per TCP session by matching individual data types against traffic. You can modify the default settings for each data type and for global options that apply to all data types in your intrusion policy. Cisco provides predefined, commonly used data types. You can also create custom data types.

A sensitive data preprocessor rule is associated with each data type. You enable sensitive data detection and event generation for each data type by enabling the corresponding preprocessor rule for the data type. A link on the configuration page takes you to a filtered view of sensitive data rules on the Rules page, where you can enable and disable rules and configure other rule attributes.

When you save changes to your intrusion policy, you are given the option to automatically enable the sensitive data preprocessor if the rule associated with a data type is enabled and sensitive data detection is disabled.

See the following sections for more information:

- [Deploying Sensitive Data Detection, page 28-20](#)
- [Selecting Global Sensitive Data Detection Options, page 28-20](#)
- [Selecting Individual Data Type Options, page 28-21](#)
- [Using Predefined Data Types, page 28-22](#)

- [Configuring Sensitive Data Detection, page 28-23](#)
- [Selecting Application Protocols to Monitor, page 28-25](#)
- [Special Case: Detecting Sensitive Data in FTP Traffic, page 28-26](#)
- [Using Custom Data Types, page 28-27](#)

Deploying Sensitive Data Detection

License: Protection

Because sensitive data detection can have a high impact on the performance of your system, Cisco recommends that you adhere to the following guidelines:

- Select the No Rules Active default policy as your base intrusion policy; see [Understanding System-Provided Base Policies, page 19-3](#) for more information.
- Ensure that the following settings are enabled in the corresponding network analysis policy:
 - **FTP and Telnet Configuration** under **Application Layer Preprocessors**
 - **IP Defragmentation** and **TCP Stream Configuration** under **Transport/Network Layer Preprocessors**.
- Apply the access control policy that includes the intrusion policy containing your sensitive data configuration to a device reserved for sensitive data detection; see [Deploying Configuration Changes, page 4-12](#) for more information.

Selecting Global Sensitive Data Detection Options

License: Protection

Global sensitive data preprocessor options control how the preprocessor functions. You can modify global options that specify the following:

- whether the preprocessor replaces all but the last four credit card or Social Security numbers in triggering packets
- which destination hosts on your network to monitor for sensitive data
- how many total occurrences of all data types in a single session result in an event

Note that global sensitive data options are policy-specific and apply to all data types.

You can configure the following global sensitive data detection options.

Mask

Replaces with Xs all but the last four digits of credit card numbers and Social Security numbers in the triggering packet. The masked numbers appear in the intrusion event packet view in the user interface and in downloaded packets.

Networks

Specifies the destination host or hosts to monitor for sensitive data. You can specify a single IP address, address block, or a comma-separated list of either or both. The system interprets a blank field as *any*, meaning any destination IP address. For information on using IPv4 and IPv6 address blocks, see [IP Address Conventions, page 1-4](#).

Global Threshold

Specifies the total number of all occurrences of all data types during a single session that the preprocessor must detect in any combination before generating a global threshold event. You can specify 1 through 65535.

Cisco recommends that you set the value for this option higher than the highest threshold value for any individual data type that you enable in your policy. See [Selecting Individual Data Type Options, page 28-21](#) for more information.

Note the following points regarding global thresholds:

- You must enable preprocessor rule 139:1 to detect and generate events on combined data type occurrences. See [Setting Rule States, page 27-19](#) for information on enabling rules in your intrusion policy.
- The preprocessor generates up to one global threshold event per session.
- Global threshold events are independent of individual data type events; that is, the preprocessor generates an event when the global threshold is reached, regardless of whether the event threshold for any individual data type has been reached, and vice versa.

Selecting Individual Data Type Options

License: Protection

Individual data types identify the sensitive data you can detect and generate events on in your specified destination network traffic. You can modify default settings for data type options that specify the following:

- a threshold that must be met for a detected data type to generate a single per-session event
- the destination ports to monitor for each data type
- the application protocols to monitor for each data type

At a minimum, each data type must specify an event threshold and at least one port or application protocol to monitor.

Each predefined data type provided by Cisco uses an otherwise inaccessible `sd_pattern` keyword to define a built-in data pattern to detect in traffic. See [Table 28-8 on page 28-23](#) for a listing of predefined data types. You can also create custom data types for which you use simple regular expressions to specify your own data patterns. See [Using Custom Data Types, page 28-27](#) for more information.

Note that data type names and patterns are system-wide; all other data type options are policy-specific.

The following table describes the data type options you can configure.

Table 28-7 *Individual Data Type Options*

Option	Description
Data Type	Displays the unique name for the data type.
Threshold	Specifies the number of occurrences of the data type when the system generates an event. You receive an error message when you save the policy if you do not set a threshold for an enabled data type. You can specify 1 through 255. Note that the preprocessor generates one event for a detected data type per session. Note also that global threshold events are independent of individual data type events; that is, the preprocessor generates an event when the data type event threshold is reached, regardless of whether the global event threshold has been reached, and vice versa.
Destination Ports	Specifies destination ports to monitor for the data type. You can specify a single port, a comma-separated list of ports, or <code>any</code> , meaning any destination port. You receive an error message when you save the policy if you enable the rule for a data type without setting at least one port or application protocol for the data type.
Application Protocols Note that this feature requires a Control license.	Specifies up to eight application protocols to monitor for the data type. You receive an error message when you save the policy if you enable the rule for a data type without setting at least one port or application protocol for the data type. See Selecting Application Protocols to Monitor, page 28-25 for detailed instructions for selecting application protocols for data types.
Pattern	For a custom data type, the specified pattern to detect (data patterns for data types provided by Cisco are predefined). See Using Custom Data Types, page 28-27 for more information. The user interface does not display built-in patterns for predefined data types. Note that custom and predefined data patterns are system-wide.

Using Predefined Data Types

License: Protection

Each intrusion policy includes predefined data types for detecting commonly used data patterns such as credit card numbers, email addresses, U.S. phone numbers, and U.S. Social Security numbers with and without dashes. Each predefined data type is associated with a single sensitive data preprocessor rule that has a generator ID (GID) of 138. You must enable the associated sensitive data rule in the intrusion policy to enable detection, and event generation, for each data type you want to use in your policy. See [Setting Rule States, page 27-19](#) for information on enabling rules in an intrusion policy.

To help you enable sensitive data rules, a link on the configuration page takes you to a filtered view of the Rules page that displays all predefined and custom sensitive data rules. You can also display only predefined sensitive data rules by selecting the sensitive-data rule filtering category on the Rules page. See [Filtering Rules in an Intrusion Policy, page 27-9](#) for more information. Predefined sensitive data rules are also listed on the Rule Editor page (**Policies > Intrusion > Rule Editor**), where you can view but not edit them under the sensitive-data rule category.

The following table describes each data type and lists the corresponding preprocessor rule that you must enable to enable detection and event generation for the data type.

Table 28-8 **Sensitive Data Types**

Data Type	Description	Preprocessor Rule GID:SID
Credit Card Numbers	Matches Visa®, MasterCard®, Discover® and American Express® fifteen- and sixteen-digit credit card numbers, with or without their normal separating dashes or spaces; also uses the Luhn algorithm to verify credit card check digits.	138:2
Email Addresses	Matches email addresses.	138:5
U.S. Phone Numbers	Matches U.S. phone numbers adhering to the pattern <code>(\d{3}) ?\d{3}-\d{4}.</code>	138:6
U.S. Social Security Numbers Without Dashes	Matches 9-digit U.S. Social Security numbers that have valid 3-digit area numbers, valid 2-digit group numbers, and do not have dashes.	138:4
U.S. Social Security Numbers With Dashes	Matches 9-digit U.S. Social Security numbers that have valid 3-digit area numbers, valid 2-digit group numbers, and dashes.	138:3
Custom	Matches a user-defined data pattern in the specified traffic. See Using Custom Data Types, page 28-27 for more information.	138:>999999

To reduce false positives from 9-digit numbers other than Social Security numbers, the preprocessor uses an algorithm to validate the 3-digit area number and 2-digit group number that precede the 4-digit serial number in each Social Security number. The preprocessor validates Social Security group numbers through November 2009.

Configuring Sensitive Data Detection

License: Protection

You can modify default global settings and settings for individual data types. You must also enable the preprocessor rule for each data type you want to detect.

If you enable sensitive data preprocessor rules in your policy without enabling sensitive data detection, you are prompted to enable sensitive data detection when you save changes to your policy. See [Resolving Conflicts and Committing Policy Changes, page 18-15](#) for more information.

The following table describes actions you can take on the Sensitive Data Detection page.

Table 28-9 **Sensitive Data Configuration Actions**

To...	You can...
modify global settings	see Table 28-6 on page 28-8 for information on the global settings you can modify.
modify data type options	click the data type name in the Targets page area. The Configuration page area updates to display the current settings for the data type. See the Individual Data Type Options table for information on the options you can modify.

Table 28-9 Sensitive Data Configuration Actions (continued)

To...	You can...
add or remove application protocols to monitor for a data type Note that this feature requires a Control license.	click inside the Application Protocols field, or click Edit next to the field. The Application Protocols pop-up window appears: - To add up to eight application protocols to monitor, select one or more application protocols from the Available list on the left, then click the right arrow (>) button. - To remove an application protocol, select it from the Enabled list on the right, then click the left arrow (<) button. Use Ctrl or Shift while clicking to select multiple application protocols. You can also click and drag to select multiple adjacent application protocols. Note To detect sensitive data in FTP traffic, you must add the <code>Ftp data</code> application protocol. See Special Case: Detecting Sensitive Data in FTP Traffic, page 28-26 for more information.
create a custom data type	click the + sign next to Data Types on the left side of the page. The Add Data Type pop-up window appears. Specify a unique data type name and the pattern you want to detect with this data type and click OK, or click Cancel to abandon your edits. See Using Custom Data Types, page 28-27 for more information.
display sensitive data preprocessor rules	click the Configure Rules for Sensitive Data Detection link above the Global Settings page area. A listing of all sensitive data preprocessor rules appears in a filtered display of the Rules page. Optionally, you can enable or disable any of the listed rules. Note that you must enable the sensitive data preprocessor rule for each data type that you want to use in your intrusion policy. See Setting Rule States, page 27-19 for more information. You can also configure sensitive data rules for any of the other actions available on the Rules page, such as rule suppression, rate-based attack prevention, and so on; see Tuning Intrusion Policies Using Rules, page 27-1 for more information. Click Back to return to the Sensitive Data Detection page.

To configure sensitive data detection:

-
- Step 1** Select **Configuration > ASA FirePOWER Configuration > Policies > Intrusion Policy**.
The Intrusion Policy page appears.
- Step 2** Click the edit icon (✎) next to the policy you want to edit.
If you have unsaved changes in another policy, click **OK** to discard those changes and continue. See [Resolving Conflicts and Committing Policy Changes, page 18-15](#) for information on saving unsaved changes in another policy.
The Policy Information page appears.
- Step 3** Click **Advanced Settings** in the navigation panel on the left.
The Advanced Settings page appears.
- Step 4** You have two choices, depending on whether **Sensitive Data Detection** under **Specific Threat Detection** is enabled:
- If the configuration is enabled, click **Edit**.
 - If the configuration is disabled, click **Enabled**, then click **Edit**.

The Sensitive Data Detection page appears. A message at the bottom of the page identifies the intrusion policy layer that contains the configuration. See [Using Layers in a Network Analysis or Intrusion Policy, page 19-1](#) for more information.

- Step 5** You can take any of the actions described in the [Sensitive Data Configuration Actions](#) table.
- Step 6** Save your policy, continue editing, discard your changes, or exit while leaving your changes in the system cache. See [Resolving Conflicts and Committing Policy Changes, page 18-15](#) for more information.

Selecting Application Protocols to Monitor

License: Control

You can specify up to eight application protocols to monitor for each data type.

You must specify at least one application protocol or port to monitor for each data type. However, except in the case where you want to detect sensitive data in FTP traffic, Cisco recommends for the most complete coverage that you specify corresponding ports when you specify application protocols. For example, if you specify HTTP, you might also configure the well-known HTTP port 80. If a new host on your network implements HTTP, the system will monitor port 80 during the interval when it is discovering the new HTTP application protocol.

In the case where you want to detect sensitive data in FTP traffic, you must specify the `FTP data` application protocol; there is no advantage in specifying a port number. See [Special Case: Detecting Sensitive Data in FTP Traffic, page 28-26](#) for more information.

To modify application protocols to detect sensitive data:

Admin/Intrusion Admin

- Step 1** Select **Configuration > ASA FirePOWER Configuration > Policies > Intrusion Policy**.
The Intrusion Policy page appears.
- Step 2** Click the edit icon () next to the policy you want to edit.
If you have unsaved changes in another policy, click **OK** to discard those changes and continue. See [Resolving Conflicts and Committing Policy Changes, page 18-15](#) for information on saving unsaved changes in another policy.
The Policy Information page appears.
- Step 3** Click **Advanced Settings** in the navigation panel on the left.
The Advanced Settings page appears.
- Step 4** You have two choices, depending on whether **Sensitive Data Detection** under **Specific Threat Detection** is enabled:
- If the configuration is enabled, click **Edit**.
 - If the configuration is disabled, click **Enabled**, then click **Edit**.
- The Sensitive Data Detection page appears.
A message at the bottom of the page identifies the intrusion policy layer that contains the configuration. See [Using Layers in a Network Analysis or Intrusion Policy, page 19-1](#) for more information.
- Step 5** Click the data type name under **Data Types** to select the data type you want to modify.

The Configuration area updates to display the current settings for the selected data type.

Step 6 Click inside the **Application Protocols** field, or click **Edit** next to the field.

The Application Protocols pop-up window appears.

Step 7 You have two choices:

- To add up to eight application protocols to monitor, select one or more application protocols from the **Available** list on the left, then click the right arrow (>) button.
- To remove an application protocol, select it from the **Enabled** list on the right, then click the left arrow (<) button.

Use Ctrl or Shift while clicking to select multiple application protocols. You can also click and drag to select multiple adjacent application protocols.



Note

To detect sensitive data in FTP traffic, you must add the `FTP data` application protocol. See [Special Case: Detecting Sensitive Data in FTP Traffic, page 28-26](#) for more information.

Step 8 Click **OK** to add the application protocols.

The Sensitive Data Detection page is displayed and the application protocols are updated.

Special Case: Detecting Sensitive Data in FTP Traffic

License: Control

You usually determine which traffic to monitor for sensitive data by specifying the ports to monitor or, optionally, specifying application protocols in deployments. However, specifying ports or application protocols is not sufficient for detecting sensitive data in FTP traffic. Sensitive data in FTP traffic is found in traffic for the FTP application protocol, which occurs intermittently and uses a transient port number, making it difficult to detect. To detect sensitive data in FTP traffic, you **must** include the following in your configuration:

- Specify the `FTP data` application protocol.

Specifying the `FTP data` application protocol enables detection of sensitive data in FTP traffic. See [Selecting Application Protocols to Monitor, page 28-25](#) for more information.

In the special case of detecting sensitive data in FTP traffic, specifying the `FTP data` application protocol does not invoke detection; instead, it invokes the rapid processing of the FTP/Telnet processor to detect sensitive data in FTP traffic. See [Decoding FTP and Telnet Traffic, page 22-18](#) for more information.

- Ensure that your configuration includes at least one port to monitor for sensitive data.

Note that it is not necessary to specify an FTP port except in the unlikely case where you only want to detect sensitive data in FTP traffic. Most sensitive data configurations will include other ports such as HTTP or email ports. In the case where you do want to specify only one FTP port and no other ports to monitor, Cisco recommends that you specify the FTP command port 23. See [Configuring Sensitive Data Detection, page 28-23](#) or more information.

Using Custom Data Types

License: Protection

You can create and modify custom data types to detect data patterns that you specify. For example, a hospital might create a data type to protect patient numbers, or a university might create a data type to detect student numbers that have a unique numbering pattern.

Each custom data type you create also creates a single sensitive data preprocessor rule that has a generator ID (GID) of 138 and a Snort ID of 1000000 or greater, that is, a SID for a local rule. You must enable the associated sensitive data rule to enable detection, and event generation, for each custom data type you want to use in your policy. See [Setting Rule States, page 27-19](#) for information on enabling rules in an intrusion policy.

To help you enable sensitive data rules, a link on the configuration page takes you to a filtered view of the Rules page that displays all predefined and custom sensitive data rules. You can also display custom sensitive data rules along with any local custom rules by selecting the local rule filtering category on the Rules page. See [Filtering Rules in an Intrusion Policy, page 27-9](#) for more information. Note that custom sensitive data rules are not listed on the Rule Editor page.

Custom data types you create are added to all intrusion policies. You must enable the associated sensitive data rule in any policy that you want to use to detect and generate events for a particular custom data type.

Note that you must use the Sensitive Data Detection configuration page to create data types and their associated rules. You cannot use the rule editor to create sensitive data rules.

See the following sections for more information:

- [Defining Data Patterns in Custom Data Types, page 28-27](#)
- [Configuring Custom Data Types, page 28-29](#)
- [Editing Custom Data Type Names and Detection Patterns, page 28-30](#)

Defining Data Patterns in Custom Data Types

License: Protection

You define the data pattern for a custom data type using a simple set of regular expressions comprised of the following:

- three metacharacters
- escaped characters that allow you to use the metacharacters as literal characters
- six character classes

Metacharacters are literal characters that have special meaning within regular expressions. The following table describes the metacharacters you can use when defining a custom data pattern.

Table 28-10 Sensitive Data Pattern Metacharacters

Metacharacter	Description	Example
?	Matches zero or one occurrence of the preceding character or escape sequence; that is, the preceding character or escape sequence is optional.	colou?r matches color or colour
{ <i>n</i> }	Matches the preceding character or escape sequence <i>n</i> times.	For example, \ d {2} matches 55, 12, and so on; \ 1 {3} matches AbC, www, and so on; \ w {3} matches a1B, 25C, and so on; x {5} matches xxxxx
\	Allows you to use metacharacters as actual characters and is also used to specify a predefined character class. See Table 28-12 on page 28-28 for a description of the character classes you can use in sensitive data patterns.	\? matches a question mark, \\ matches a backslash, \d matches numeric characters, and so on

You must use a backslash to escape the characters in the following table for the sensitive data preprocessor to interpret them correctly as literal characters.

Table 28-11 Escaped Sensitive Data Pattern Characters

Use this escaped character...	To represent this literal character...
\?	?
\{	{
\}	}
\\	\

The following table describes the character classes you can use when defining a custom sensitive data pattern.

Table 28-12 Sensitive Data Pattern Character Classes

Character Class	Description	Character Class Definition
\d	Matches any numeric ASCII character 0-9	0-9
\D	Matches any byte that is not a numeric ASCII character	not 0-9
\l (lowercase “ell”)	Matches any ASCII letter	a-zA-Z
\L	Matches any byte that is not an ASCII letter	not a-zA-Z
\w	Matches any ASCII alphanumeric character Note that, unlike PCRE regular expressions, this does not include an underscore (_).	a-zA-Z0-9
\W	Matches any byte that is not an ASCII alphanumeric character	not a-zA-Z0-9

The preprocessor treats characters entered directly, instead of as part of a regular expression, as literal characters. For example, the data pattern 1234 matches 1234.

The following data pattern example, which is used in predefined sensitive data rule 138:4, uses the escaped digits character class, the multiplier and option-specifier metacharacters, and the literal dash (-) and left and right parentheses () characters to detect U.S. phone numbers:

```
(\d{3}) ?\d{3}-\d{4}
```

Exercise caution when creating custom data patterns. Consider the following alternative data pattern for detecting phone numbers which, although using valid syntax, could cause many false positives:

```
(?\d{3})? ?\d{3}-?\d{4}
```

Because the second example combines optional parentheses, optional spaces, and optional dashes, it would detect, among others, phone numbers in the following desirable patterns:

- (555) 123-4567
- 555123-4567
- 5551234567

However, the second example pattern would also detect, among others, the following potentially invalid patterns, resulting in false positives:

- (555 1234567
- 555) 123-4567
- 555) 123-4567

Consider finally, for illustration purposes only, an extreme example in which you create a data pattern that detects the lowercase letter a using a low event threshold in all destination traffic on a small company network. Such a data pattern could overwhelm your system with literally millions of events in only a few minutes.

Configuring Custom Data Types

License: Protection

You configure essentially the same data type options for custom data types that you configure for predefined data types. See [Selecting Individual Data Type Options, page 28-21](#) for information on setting options that are common to all data types. In addition, you must also specify the name and data pattern for custom data types.

Note that creating a custom data type also creates an associated custom sensitive data preprocessing rule, which you must enable in each policy where you want to use that data type. See [Setting Rule States, page 27-19](#) for information on enabling rules in your intrusion policy.

To create or modify a custom data type:

Admin/Intrusion Admin

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies > Intrusion Policy**.

The Intrusion Policy page appears.

Step 2 Click the edit icon (✎) next to the policy you want to edit.

If you have unsaved changes in another policy, click **OK** to discard those changes and continue. See [Resolving Conflicts and Committing Policy Changes, page 18-15](#) for information on saving unsaved changes in another policy.

The Policy Information page appears.

Step 3 Click **Advanced Settings** in the navigation panel on the left.

The Advanced Settings page appears.

- Step 4** You have two choices, depending on whether **Sensitive Data Detection** under **Specific Threat Detection** is enabled:

- If the configuration is enabled, click **Edit**.
- If the configuration is disabled, click **Enabled**, then click **Edit**.

The Sensitive Data Detection page appears.

A message at the bottom of the page identifies the intrusion policy layer that contains the configuration. See [Using Layers in a Network Analysis or Intrusion Policy, page 19-1](#) for more information.

- Step 5** You have the following options:

- To create a custom data type, click the + sign next to **Data Types** on the left side of the page. The Add Data Type pop-up window appears.

Specify a unique data type name and the pattern you want to detect with this data type and click **OK**, or click **Cancel** to abandon your edits. See [Editing Custom Data Type Names and Detection Patterns, page 28-30](#) for more information.

The Sensitive Data Detection page appears. If you clicked **OK**, the page updates to display your changes.

- To modify any of the options that are common to predefined and custom data types, click the data type name in the **Targets** page area.

The Configuration page area updates to display the current settings for the data type. See [Configuring Sensitive Data Detection, page 28-23](#) for more information.

- To edit the system-wide name and data pattern for a custom data type, see [Editing Custom Data Type Names and Detection Patterns, page 28-30](#).
- To delete a custom data type, click the delete icon () next to the data type you want to remove and then click **OK**, or click **Cancel** to abandon deleting the data type.

Note that you cannot delete a data type when the sensitive data rule for that data type is enabled in any intrusion policy. Deleting a custom data type deletes it from all intrusion policies.

Editing Custom Data Type Names and Detection Patterns

License: Protection

You can modify the system-wide name and detection pattern for custom sensitive data rules. Note that changing these settings changes them in all other policies on the system. Note also that you must reapply any applied access control policies that include intrusion policies that use custom data types that you modify.

Except for custom data type names and data patterns, all data type options are policy-specific for both custom and predefined data types. See [Selecting Individual Data Type Options, page 28-21](#) for information on modifying options other than the name and data pattern in your custom data types.

To edit custom data type names and data patterns:

Admin/Intrusion Admin

- Step 1** Select **Configuration > ASA FirePOWER Configuration > Policies > Intrusion Policy**.

The Intrusion Policy page appears.

- Step 2** Click the edit icon () next to the policy you want to edit.
- If you have unsaved changes in another policy, click **OK** to discard those changes and continue. See [Resolving Conflicts and Committing Policy Changes, page 18-15](#) for information on saving unsaved changes in another policy.
- The Policy Information page appears.
- Step 3** Click **Advanced Settings** in the navigation panel on the left.
- The Advanced Settings page appears.
- Step 4** You have two choices, depending on whether **Sensitive Data Detection** under **Specific Threat Detection** is enabled:
- If the configuration is enabled, click **Edit**.
 - If the configuration is disabled, click **Enabled**, then click **Edit**.
- The Sensitive Data Detection page appears.
- A message at the bottom of the page identifies the intrusion policy layer that contains the configuration. See [Using Layers in a Network Analysis or Intrusion Policy, page 19-1](#) for more information.
- Step 5** In the **Targets** page area, click the name of the custom data type you want to modify.
- The page updates to show the current settings for the data type, and the **Edit Data Type Name and Pattern** link appears in the upper right of the Configuration page area.
- Step 6** Click the **Edit Data Type Name and Pattern** link.
- The Edit Data Type pop-up window appears.
- Step 7** Modify the data type name, pattern, or both and click **OK**, or click **Cancel** to abandon your edits. See [Defining Data Patterns in Custom Data Types, page 28-27](#) for information on specifying the data pattern.
- The Sensitive Data Detection page appears. If you clicked **OK**, the page displays your changes.
-

