Intrusion Policies

The following topics explain intrusion policies and the closely associated network analysis policies (NAP). Intrusion policies include rules that check traffic for threats and block traffic that appears to be an attack. Network analysis policies control traffic preprocessing, which prepares traffic to be further inspected by normalizing traffic and identifying protocol anomalies.

Because preprocessing and intrusion inspection are so closely related, the network analysis and intrusion policies examining a single packet must complement each other.

- About Intrusion and Network Analysis Policies, on page 1
- License Requirements for Intrusion Policies, on page 6
- Applying Intrusion Policies in Access Control Rules, on page 7
- Configuring Syslog for Intrusion Events, on page 7
- Configuring the Network Analysis Policy (Snort 3), on page 8
- Managing Intrusion Policies (Snort 3), on page 12
- Monitoring Intrusion Policies, on page 25
- Examples for Intrusion Policies, on page 25

About Intrusion and Network Analysis Policies

Network analysis and intrusion policies work together to detect and prevent intrusion threats.

- A network analysis policy (NAP) governs how traffic is decoded and preprocessed so that it can be further evaluated, especially for anomalous traffic that might signal an intrusion attempt.
- An intrusion policy uses intrusion and preprocessor rules, which are collectively known as intrusion rules, to examine the decoded packets for attacks based on patterns. The rules can either prevent (drop) the threatening traffic and generate an event, or simply detect (alert) it and generate an event only.

As the system analyzes traffic, the network analysis decoding and preprocessing phase occurs before and separately from the intrusion prevention phase. Together, network analysis and intrusion policies provide broad and deep packet inspection. They can help you detect, alert on, and protect against network traffic that could threaten the availability, integrity, and confidentiality of hosts and their data.

System-Defined Network Analysis and Intrusion Policies

The system includes several pairs of same-named network analysis and intrusion policies that complement and work with each other. For example there are both NAP and intrusion policies named “Balanced Security and Connectivity," which are meant to be used together. The system-provided policies are configured by the Cisco Talos Intelligence Group (Talos). For these policies, Talos sets the intrusion and preprocessor rule states and provides the initial configurations for preprocessors and other advanced settings.

As new vulnerabilities become known, Talos releases intrusion rule updates. These rule updates can modify any system-provided network analysis or intrusion policy, and can provide new and updated intrusion rules and preprocessor rules, modified states for existing rules, and modified default policy settings. Rule updates might also delete rules from system-provided policies and provide new rule categories, as well as modify the default variable set.

You can manually update the rules database, or configure a regular update schedule. You must deploy an update for it to take effect. For more information on updating system databases, see Updating System Databases.

The following are the system-provided policies:

Balanced Security and Connectivity network analysis and intrusion policies

These policies are built for both speed and detection. Used together, they serve as a good starting point for most networks and deployment types. The system uses the Balanced Security and Connectivity network analysis policy as the default.

Connectivity Over Security network analysis and intrusion policies

These policies are built for networks where connectivity, the ability to get to all resources, takes precedence over network infrastructure security. The intrusion policy enables far fewer rules than those enabled in the Security over Connectivity policy. Only the most critical rules that block traffic are enabled.

Security Over Connectivity network analysis and intrusion policies

These policies are built for networks where network infrastructure security takes precedence over user convenience. The intrusion policy enables numerous network anomaly intrusion rules that could alert on or drop legitimate traffic.

Maximum Detection network analysis and intrusion policies

These policies are built for networks where network infrastructure security is given even more emphasis than is given by the Security Over Connectivity policies, with the potential for even greater operational impact. For example, the intrusion policy enables rules in a large number of threat categories including malware, exploit kit, old and common vulnerabilities, and known in-the-wild exploits.

Inspection Mode: Prevention vs. Detection

By default, all intrusion policies operate in Prevention mode to implement an Intrusion Prevention System (IPS). In the Prevention inspection mode, if a connection matches an intrusion rule whose action is to drop traffic, the connection is actively blocked.

If you instead want to test the effect of the intrusion policy on your network, you can change the mode to Detection, which implements an Intrusion Detection System (IDS). In this inspection mode, drop rules are treated like alert rules, where you are notified of matching connections, but the action result becomes Would Have Blocked, and connections are never in fact blocked.

You change the inspection mode per intrusion policy, so you can have a mix of prevention and detection.

The Snort 3 network analysis policy (NAP) also has an inspection mode. Unlike the intrusion policy, the NAP policy is global, so you must run all NAP processing in either Prevention or Detection mode. You should use the same mode you use for your intrusion policies. If you have a mix of Prevention and Detection policies, select Prevention to match your most restrictive intrusion policies.

Intrusion and Preprocessor Rules

An intrusion rule is a specified set of keywords and arguments that the system uses to detect attempts to exploit vulnerabilities in your network. As the system analyzes network traffic, it compares packets against the conditions specified in each rule, and triggers the rule if the data packet meets all the conditions specified in the rule.

The system includes the following types of rules created by Cisco Talos Intelligence Group (Talos):

- Intrusion rules, which are subdivided into shared object rules and standard text rules
- Preprocessor rules, which are rules associated with preprocessors and packet decoder detection options in the network analysis policy. Most preprocessor rules are disabled by default.

The following topics explain intrusion rules in more depth.

Intrusion Rule Attributes

When you view an intrusion policy, you see a list of all the intrusion rules available for identifying threats.

The list of rules for each policy is the same. The difference is in the action configured for each rule. Because there are over 30,000 rules, scrolling through the list will take time. Rules are revealed as you scroll through the list.

Following are the attributes that define each rule:

> (Signature Description)

Click the > button in the left column to open the signature description. The description is the actual code used by the Snort inspection engine to match traffic against the rule. Explaining the code is out of scope for this document, but it is explained in detail in *Management Center Configuration Guide*; select the book for your software version from http://www.cisco.com/c/en/us/support/security/defense-center/products-installation-and-configuration-guides-list.html. Look for information on intrusion rule editing.

The signatures contain variables for certain items. For more information, see Default Intrusion Variable Set, on page 4.

GID

Generator Identifier (ID). This number indicates which system component evaluates the rule and generates events. A 1 indicates a standard text intrusion rule, a 3 indicates a shared object intrusion rule. (The difference in these rule types is not meaningful for a Firewall Device Manager user.) These are the main rules of interest when configuring an intrusion policy. For information on the other GIDs, see Generator Identifiers, on page 5.

SID

Snort Identifier (ID), also called signature ID. Snort IDs lower than 1000000 were created by the Cisco Talos Intelligence Group (Talos).

Action

The state of this rule in the selected intrusion policy. For each rule, “(Default)” is added to the action that is the default action for the rule within this policy. To return a rule to its default setting, you select this action. Possible actions are:

- **Alert**—Create an event when this rule matches traffic, but do not drop the connection.
- **Drop**—Create an event when this rule matches traffic, and also drop the connection.

- **Disabled**—Do not match traffic against this rule. No events are generated.

Message

This is the name of the rule, which also appears in events triggered by the rule. The message typically identifies the threat that the signature matches. You can search the Internet for more information on each threat.

Default Intrusion Variable Set

The intrusion rule signatures contain variables for certain items. Following are the default values for the variables, with $HOME_NET and $EXTERNAL_NET being the most commonly used variables. Note that the protocol is specified separately from port numbers, so port variables are numbers only.

- $DNS_SERVERS = $HOME_NET (meaning any IP address).
- $EXTERNAL_NET = any IP address.
- $FILE_DATA_PORTS = $HTTP_PORTS, 143, 110.
- $FTP_PORTS = 21, 2100, 3535.
- $GTP_PORTS = 3386, 2123, 2152.
- $HOME_NET = any IP address.
- $HTTP_PORTS = 144 ports numbered: 36, 80-90, 311, 383, 443, 555, 591, 593, 631, 666, 801, 808, 818, 901, 972, 1158, 1212, 1220, 1414, 1422, 1533, 1741, 1830, 1942, 2231, 2301, 2381, 2578, 2809, 2980, 3029, 3037, 3057, 3128, 3443, 3507, 3702, 4000, 4343, 4848, 5000, 5117, 5222, 5250, 5450, 5600, 5814, 6080, 6173, 6767, 6988, 7000, 7001, 7005, 7071, 7080, 7144, 7145, 7510, 7770, 7777-7779, 8000, 8001, 8008, 8014, 8015, 8020, 8028, 8040, 8060, 8080-8082, 8085, 8088, 8118, 8123, 8161, 8180-8182, 8222, 8243, 8280, 8300, 8333, 8344, 8400, 8443, 8500, 8509, 8787, 8800, 8888, 8899, 8983, 9000, 9002, 9060, 9080, 9090, 9091, 9111, 9290, 9443, 9447, 9710, 9788, 9999, 10000, 11371, 12601, 13014, 15489, 19980, 23472, 29991, 33300, 34412, 34443, 34444, 40007, 41080, 44449, 50000, 50002, 51423, 53331, 55252, 55555, 56712.
- $HTTP_SERVERS = $HOME_NET (meaning any IP address).
- $ORACLE_PORTS = any
- $SHELLCODE_PORTS = 180.
- $SIP_PORTS = 5060, 5061, 5600
- $SIP_SERVERS = $HOME_NET (meaning any IP address).
- $SMTP_SERVERS = $HOME_NET (meaning any IP address).
- $SNMP_SERVERS = $HOME_NET (meaning any IP address).
- $SQL_SERVERS = $HOME_NET (meaning any IP address).
- $SSH_PORTS = 22.
- $SSH_SERVERS = $HOME_NET (meaning any IP address).
- $TELNET_SERVERS = $HOME_NET (meaning any IP address).

Generator Identifiers

The generator identifier (GID) identifies the subsystem that evaluates an intrusion rule and generates events. Standard text intrusion rules have a generator ID of 1, and shared object intrusion rules have a generator ID of 3. There are also several sets of rules for various preprocessors. The following table explains the GIDs.

Table 1: Generator IDs

ID	Component
1	Standard Text Rule.
2	Tagged Packets. (Rules for the Tag generator, which generates packets from a tagged session.)
3	Shared Object Rule.
102	HTTP Decoder.
105	Back Orifice Detector.
106	RPC Decoder.
116	Packet Decoder.
119, 120	HTTP Inspect Preprocessor. (GID 120 rules relate to server-specific HTTP traffic.)
122	Portscan Detector.
123	IP Defragmentor.
124	SMTP Decoder. (Exploits against SMTP verbs.)
125	FTP Decoder.
126	Telnet Decoder.
128	SSH Preprocessor.
129	Stream Preprocessor.
131	DNS Preprocessor.
133	DCE/RPC Preprocessor.
134	Rule Latency, Packet Latency. (Events for these rules are generated when rule latency suspends (SID 1) or re-enables (SID 2) a group of intrusion rules, or when the system stops inspecting a packet because the packet latency threshold is exceeded (SID 3).)

ID	Component
135	Rate-Based Attack Detector. (Excessive connections to hosts on the network.)
137	SSL Preprocessor.
138, 139	Sensitive Data Preprocessor.
140	SIP Preprocessor.
141	IMAP Preprocessor.
142	POP Preprocessor.
143	GTP Preprocessor.
144	Modbus Preprocessor.
145	DNP3 Preprocessor.

Network Analysis Policies

Network analysis policies control traffic preprocessing. Preprocessors prepare traffic to be further inspected by normalizing traffic and identifying protocol anomalies. Network analysis-related preprocessing occurs after Security Intelligence drops and SSL decryption, but before access control and intrusion or file inspection.

By default, the system uses the Balanced Security and Connectivity network analysis policy to preprocess all traffic handled by the access control policy. However, if you configure an intrusion policy on any access control rule, the system uses the network analysis policy that matches the most aggressive intrusion policy applied. For example, if you use both Security over Connectivity and Balanced policies in your access control rules, the system uses the Security over Connectivity NAP for all traffic. For Snort 3 custom intrusion policies, this assignment is done according to the base template policy assigned to the intrusion policy.

When using Snort 3, you can select a policy explicitly and optionally customize its settings. We recommend that you select the policy whose name matches the intrusion policy that you use for most traffic that goes through the device, whether you use the intrusion policy directly, or you use it as the base policy in your custom intrusion policies. You can then change the inspection mode, or adjust specific inspector or binder settings to account for the traffic on your network.

In addition, consider whether you have enabled preprocessor rules in the intrusion policy. If you enable rules that require a preprocessor, ensure that you also enable the corresponding inspector in the NAP. For each inspector, you can also adjust the attributes for the inspector, including the ports examined (the binders), to customize inspector behavior for your network.

License Requirements for Intrusion Policies

You must enable the **IPS** license to apply intrusion policies in access control rules. For information on configuring licenses, see Enabling or Disabling Optional Licenses.

No extra license is needed for network analysis policies.

Applying Intrusion Policies in Access Control Rules

To apply intrusion policies to network traffic, you select the policy within an access control rule that allows traffic. You do not directly assign intrusion policies.

You can assign different intrusion policies to provide variable intrusion protection based on the relative risks of the networks you are protecting. For example, you might use the more stringent Security over Connectivity policy for traffic between your inside network and external networks. On the other hand, you might apply the more lenient Connectivity over Security policy for traffic between inside networks.

You can also simplify your configuration by using the same policy for all networks. For example, the Balanced Security and Connectivity policy is design to provide good protection without excessively impacting connectivity.

Procedure

Step 1 Select **Policies** > **Access Control**.

Step 2 Either create a new rule, or edit an existing rule, that **allows** traffic.

If the default action is allow, you can also specify an intrusion policy in the default action.

You cannot apply intrusion policies to rules that trust or block traffic.

Step 3 Click the **Intrusion Policy** tab.

Step 4 Select **Intrusion Policy** > **On** and select the intrusion inspection policy to use on matching traffic.

Configuring Syslog for Intrusion Events

You can configure an external syslog server for intrusion policies to send intrusion events to your syslog server. You must configure the syslog server on the intrusion policy to get intrusion events sent to the server. Configuring a syslog server on an access rule sends connection events only to the syslog server, not intrusion events.

If you select multiple syslog servers, events are sent to each of the servers.

Intrusion events have the message ID 430001.

Procedure

Step 1 Select **Policies** > **Intrusion**.

Step 2 Click the **Intrusion Policy Settings** button (⚙) to configure syslog.

Step 3 Click the **+** button under **Send Intrusion Events To** and select the server objects that define the syslog servers. If the required objects do not already exist, click **Create New Syslog Server** and create them.

Step 4 Click **OK**.

Configuring the Network Analysis Policy (Snort 3)

The Network Analysis Policy (NAP) is applied to all allowed connections on the device. The NAP determines which inspectors are enabled, and the values of the attributes used by the inspectors. The binders determine the ports and protocols that should be associated with the various inspectors.

Coordinate the NAP with the intrusion policies you apply in access control rules:

- If you use a single intrusion policy in your access control rules, select the same-named NAP. Then, make adjustments to inspectors and attributes based on the settings in your intrusion policy. For example, if you enable intrusion rules for a particular inspector, such as CIP, ensure that you enable that inspector in the NAP.
- If you use multiple intrusion policies, select the NAP that matches the most strict intrusion policy that you use.
- If you use custom intrusion policies, make your NAP selection based on the base intrusion policy for your custom intrusion policies.
- If you do not need to customize any inspectors or binders, consider configuring the system to automatically select the most appropriate NAP based on your intrusion policy usage. This is the default option.

Before you begin

Unless you prevent it, the system regularly downloads LSP updates to the inspection rules. These updates can add or remove inspectors and attributes, and change default settings for attributes. If you have made overrides to removed inspectors, these overrides are preserved and you will see warnings that the inspector is no longer supported. In this case, delete the inspector and make any other flagged adjustments to ensure your NAP is completely valid.

Procedure

Step 1 Choose **Policies** > **Intrusion**.

Verify that the Snort version shown above the table is 3.x.

Step 2 Click the **Intrusion Policy Settings** button (⚙).

Step 3 In **Default Network Analysis Policy**, select one of the following:

- **Auto**—Automatically select the NAP that matches the most-used intrusion policy (or base policy for custom rules) applied in access control rules. If you do not apply any intrusion policies, the Balanced Security and Connectivity NAP is used. The NAP is run in Prevention mode, and you cannot customize intrusion or binder settings. The remainder of this procedure does not apply when running in automatic mode.

- **Custom**—Explicitly select the NAP that should be used. Click the **Edit** link next to the policy name to select a different policy. You can then select the inspection mode, and customize inspector and binder settings, as described in the following steps.

Step 4 In the Edit Network Analysis Policy dialog box, select the policy and configure its settings.

a) In **Network Analysis Policy**, select the policy that should apply globally to all allowed connections.

b) Select the **Inspection Mode**.

The inspection mode determines how non-compliant traffic is handled. Use the same inspection mode as you use in your intrusion policies for optimal results.

- **Prevention**—Block any decoder, normalization, or protocol anomalies based on the settings in the policy. You must use this option if you enable the SSL Decryption policy, or if you enable the **TLS Server Identity Discovery** option in the access control policy settings.
- **Detection**—Simply issue alerts for decoder, normalization, or protocol anomalies. Do not block any traffic.

c) (Optional.) Configure and manage overrides to the inspectors and binders:

- To edit overrides, see Configuring Inspector and Binder Overrides, on page 9.
- To download the schema or overrides, see Downloading Overrides and the Schema, on page 11.
- To upload overrides, see Uploading Overrides, on page 12.
- To reset all the overrides, click the **Reset Inspector/Binder Overrides** link above the NAP file. You are asked to confirm the reset. The deletion is limited to inspectors or binders, as indicated in the command name. For example, deleting all binder overrides leaves your inspector overrides unchanged.
- To reverse all changes to the selected inspector, click **Reset *Inspector* to Defaults**.
- To filter the view so you see only those inspectors that have overrides, click **Show Only Overrides**. Click **Show All Inspectors** to remove the filter.

d) Click **OK**.

Configuring Inspector and Binder Overrides

When you select a base NAP, you are selecting inspector settings contained in that baseline policy. In most cases, these are the appropriate settings.

However, you can override the settings in the selected NAP. For example, you can enable or disable individual inspectors, or change the value for an attribute or a binder.

The following procedure explains how to configure overrides directly. Alternatively, you can download the schema, make changes offline, then upload your overrides. You can also upload the overrides you downloaded from another device.

Before you begin

Explaining each inspector, binder, and attribute is beyond the scope of this document. For detailed information, including examples, see *Snort 3 Inspector Reference* at https://www.cisco.com/c/en/us/td/docs/security/secure-firewall/snort3-inspectors/snort-3-inspector-reference.html.

Procedure

Step 1 Choose **Policies** > **Intrusion**, click the **Intrusion Policy Settings** button (⚙), select **Custom** for the NAP settings, then click the **Edit** link next to the policy name.

Step 2 Click the tab that contains the setting you want to change:

- **Inspectors**—The inspectors examine specific types of traffic, such as FTP, for protocol anomalies.
- **Binders**—The Binder inspector determines when a service inspector needs to be used to inspect traffic. The configurations in the binder inspector include the ports, hosts, CIDRs, and services that define when another inspector in the network analysis policy needs to inspect traffic.

Step 3 Edit the settings as needed.

- Use the following to control the view in the JSON editor:
  - Use the **Filter** edit box to do a full-text search of the JSON file.
  - Click the **Expand All Fields** button (⇕) to open all of the folders in the JSON file.
  - Click the **Collapse All Fields** (⇕) button to close all folders in the JSON file.
  - Click the **Undo Last Action** (↶) button to reverse your most recent change.
  - Click the **Redo** (↷) button to redo your last reversed change.
  - Select **Tree** to see a formatted view of the JSON file, which includes action menus, error flags, and other features to guide your edits.
  - Select **Code** to see the raw JSON file.
- In Tree view, click the **Menu** (⋮) button to manipulate the contents of the file. You can:
  - **Insert** attributes. Use Auto to have the editor determine the appropriate data type. Otherwise, add an Array, Object, or String. If you add an invalid attribute, the system will mark the inspector or binder as having a problem that you must resolve.
  - **Append** attributes. This action does the same as Insert, but puts the attribute at the end of the section.
  - **Duplicate** the selected attribute.
  - **Remove** (delete) the selected attribute. When editing an attribute, a popup message might also provide a **Delete** command.
- To enable an inspector that is currently disabled, or change the setting of any Boolean attribute, click the checkbox in front of the attribute value. For example, to enable an inspector, change the **enabled : false** attribute to:

enabled : ☑ true

- To change the value of a string or numeric attribute, click in the attribute and edit the value as needed. If your entry violates the rules for the field, error messages explain the discrepancy. For example, a numeric value indicates the valid range of values if you enter a value outside the range.
- To reset overrides:
  - Click **Reset Inspector/Binder Overrides** to remove all of your changes to all inspectors or binders and return to default values. The deletion is limited to inspectors or binders, as indicated in the command name. For example, deleting all binder overrides leaves your inspector overrides unchanged.
  - Click **Reset *Inspector* to Defaults** to reverse all changes to the selected inspector only.
- To filter the view so you see only those inspectors that have overrides, click **Show Only Overrides**. Click **Show All Inspectors** to remove the filter.
- If an inspector is no longer supported, the inspector is flagged with a message. Click the **Delete Inspector** link in the message to remove the inspector.

Step 4 Click **OK** when finished.

Downloading Overrides and the Schema

You can download the NAP schema, or download the overrides you have configured for the policy.

Downloading overrides is recommended whenever you change the base NAP, in case you want to go back to your previous settings. In addition, you could use the JSON editor on one device to implement the overrides you want to use on all devices, download the overrides, then upload that override file to other devices.

Downloading the schema is useful if you want to edit the file offline, then upload your overrides to this device or to multiple devices. You should copy/paste just the sections you need to change, rather than upload the entire file, to ensure only those changes you make are considered overrides.

Procedure

Step 1 Choose **Policies** > **Intrusion**, click the **Intrusion Policy Settings** button (⚙), select **Custom** for the NAP settings, then click the **Edit** link next to the policy name.

Step 2 Do one of the following:

- To download the schema for the currently-select NAP, click the gear icon (⚙) and select **Download** > **Policy Schema**.
- To download the saved set of overrides, as they existed prior to the current editing session, click the gear icon (⚙) and select **Download** > **Last Saved Overrides**. The file includes overridden attributes plus their containing objects.

- To download the overrides you have created in the current editing session, click the gear icon (⚙) and select **Download** > **Current Unsaved Overrides**. The file includes overridden attributes plus their containing objects.

Uploading Overrides

Rather than editing attributes using the embedded JSON editor, you can download the NAP policy schema, edit the file offline, then upload the file. Any overrides configured in the uploaded file are then applied to the selected NAP.

You can also upload a file you downloaded after configuring overrides on another device.

By uploading your overrides, you can upload the same file to multiple devices and easily apply the same overrides.

Before you begin

For overriding an inspector configuration in the network analysis policy, you should upload only the changes that you require. You should not upload the entire configuration because it makes the overrides sticky in nature and therefore, any subsequent changes to the default values or configuration as part of the LSP updates would not be applied. Ensure that your uploaded overrides are tightly focused on just those attributes you want to change.

Procedure

Step 1 Choose **Policies** > **Intrusion**, click the **Intrusion Policy Settings** button (⚙), select **Custom** for the NAP settings, then click the **Edit** link next to the policy name.

Step 2 Click the gear icon (⚙) and select **Upload** > **Overrides**.

Step 3 (Optional.) Click one of the **Download** links to save a copy of your existing overrides.

You can either download the last saved overrides (those made before the current editing session) or the current unsaved overrides (those made during the current editing session).

Step 4 Click **Yes** on the Confirm Upload Overrides dialog box to confirm you want to continue.

Step 5 Click **Browse**, or drag and drop, to select the JSON file that contains your overrides, and click **OK**.

Managing Intrusion Policies (Snort 3)

When you use Snort 3 as the inspection engine, you can create your own intrusion policies and customize them for your purposes. The system comes with pre-defined policies that are based on the same-named Cisco Talos Intelligence Group (Talos)-defined policies. Although you can edit these policies, a better practice is to create your own policy based on the underlying Talos policy and change that if you need to adjust rule actions.

Each of these pre-defined policies includes the same list of intrusion rules (also known as signatures), but they differ in the actions taken for each rule. For example, a rule might be enabled in one policy, but disabled in another policy.

If you find that a particular rule is giving you too many false positives, where the rule is blocking traffic that you do not want blocked, you can disable the rule without needing to switch to a less-secure intrusion policy. You could alternatively change it to alert on matches without dropping traffic.

Conversely, if you know you need to protect against a specific attack, but the related rule is disabled in your chosen intrusion policy, you can enable the rule without changing to a more secure policy.

Use the intrusion related dashboards, and the Event Viewer (both on the **Monitoring** page), to evaluate how intrusion rules are impacting traffic. Keep in mind that you will see intrusion events and intrusion data only for traffic that matches intrusion rules set to alert or drop; disabled rules are not evaluated.

Procedure

Step 1 Choose **Policies** > **Intrusion**.

Verify that the Snort version shown above the table is 3.x.

Step 2 Do any of the following:

- Use the **Search/Filter** box to find a policy. You can search by name only.
- Click the gear icon (⚙) to enable logging to a syslog server. See Configuring Syslog for Intrusion Events, on page 7.
- Click the gear icon (⚙) to configure the network analysis policy (NAP). See Configuring the Network Analysis Policy (Snort 3), on page 8.
- Click + to create a new policy. See Configuring a Custom Intrusion Policy (Snort 3), on page 13.
- Click the edit icon (✎) to see the properties and rules in the policy, and to edit them. See Viewing or Editing Intrusion Policy Properties (Snort 3), on page 14.
- Click the delete icon (🗑) to delete a policy.

Configuring a Custom Intrusion Policy (Snort 3)

You can create new intrusion policies to customize rule behavior if the pre-defined policies do not fit your needs. In general, it is good practice to create custom policies based on the pre-defined policies rather than to alter those policies. This ensures you can easily implement one of the Cisco Talos defined policies if you find that your customizations do not deliver the results that you need.

Procedure

Step 1 Choose **Policies** > **Intrusion**.

Step 2 Do one of the following:

- To create a new policy, click **+**.
- To edit an existing policy, click the edit icon (✎) for the policy. When you are shown the policy details, click the **Edit** link in the policy properties section at the top of the page.

Step 3 Enter a **Name**, and optionally, a description for the policy.

Step 4 Configure the **Inspection Mode** for the policy.

- **Prevention**—Intrusion rule actions are always applied. Connections that match a drop rule are blocked.
- **Detection**—Intrusion rules generate alerts only. A connection that matches a drop rule will generate alert messages, but the connection will not be blocked.

Step 5 Select the **Base Template** for the policy.

The base templates are provided by Cisco Talos. Click the info icon for each to see more information about the policies. Note that policy names can change, and new policies appear, when a new rules package gets installed.

- **Maximum Detection (Cisco Talos)**—This policy places all emphasis on security. Network connectivity and throughput is not guaranteed and false positives are likely. This policy should only be used for high security areas and security monitors must be prepared to investigate alerts to determine their validity.
- **Security Over Connectivity (Cisco Talos)**—This policy places an emphasis on security, at the possible expense of network connectivity and throughput. Traffic is inspected more deeply, more rules are evaluated, and both false positives and increased latency are expected but within reason.
- **Balanced Security and Connectivity (Cisco Talos)**—(Default.) This policy attempts to strike the delicate balance between network connectivity and throughput and the needs of security. While not as strict as Security Over Connectivity, this policy attempts to keep users secure while being less obtrusive about normal traffic.
- **Connectivity Over Security (Cisco Talos)**—This policy places an emphasis on network connectivity and throughput, at the possible expense of security. Traffic is inspected less deeply, and fewer rules are evaluated.
- **No Rules Active (Cisco Talos)**—This policy is a basic policy that configures typical preprocessor settings but does not have any rules or built-in alerts enabled. Use this policy as a base if you want to ensure that only those policies you want to apply are enabled.

Step 6 Click **OK**.

You are returned to the intrusion policy list. You can now view the new policy and adjust rule actions as needed.

Viewing or Editing Intrusion Policy Properties (Snort 3)

The Intrusion Policy page shows a list of the policies, including both pre-defined and user-defined policies, and their descriptions. To edit a policy, you must first view the policy's properties.

Procedure

Step 1 Choose **Policies** > **Intrusion**.

Step 2 Click the edit icon (✎) for a policy.

The policy contains the following sections:

- **Policy Name** drop-down list.
  - You can easily switch to a different policy by selecting it from the drop-down list, or return to the list of policies by clicking the back button (←).
  - You can delete this policy by clicking the delete icon next to the policy name (🗑).
- **General properties**. This section shows the intrusion mode, base policy, and description. Click **Edit** to change these properties or the policy name.
- **Rule Group** table of contents. This list displays all the rule groups that have active rules in the policy. The groups have a hierarchy, with parent groups containing child groups that organize subsets of rules within the larger parent group. Each group is a logical collection of rules, and a given rule can appear in more than one group.
  - To add a group that currently has no active rules in the policy, click **+** > **Add Existing Rule Group** and select the group. See Adding or Removing Rule Groups in an Intrusion Policy (Snort 3), on page 17.
  - To change the security level of a group, select the child group in the list. The rule list changes to show the security level at the top, with the rules in the group listed below. Click the **Edit** link next to the security level and select a new level. Click **View Description** when editing to get information about each security level. Note that changing the level can change which rules are active, and also the action for a given rule, with more secure levels tending to have more active rules and more rules with the Drop action. Click **OK** to confirm the change. (Security Level does not apply to custom rule groups.)

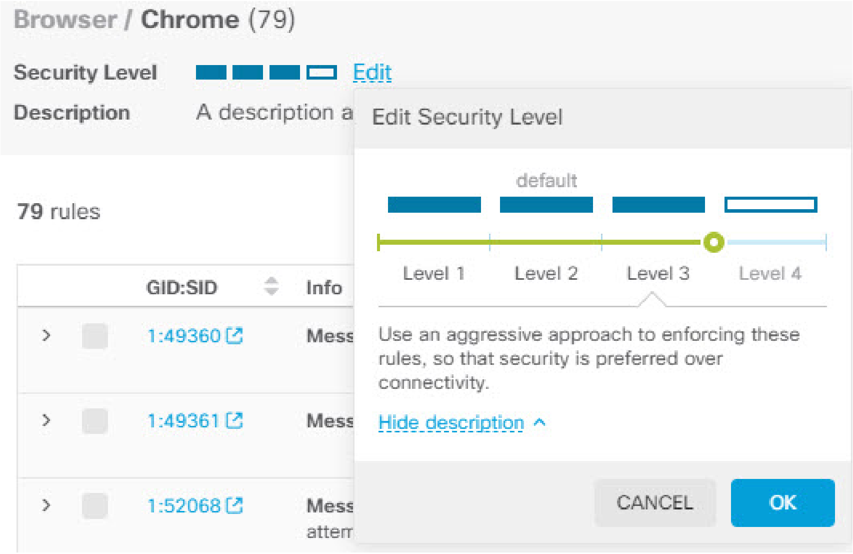

Browser / Chrome (79)

Security Level ■■■□ Edit

Description A description a

79 rules

Edit Security Level

Use an aggressive approach to enforcing these rules, so that security is preferred over connectivity.

  - To remove all the rules in a group, select the child group in the list. Then, click the **Exclude** link to the far right of the group name and confirm that you want to exclude the group. Excluding the group simply disables all the rules in the group. It does not delete the group.

    However, if the group includes rules that are shared with other groups that are enabled, the shared rules retain whatever actions are applied by the still-active group. In all cases, we retain your most aggressive setting for an individual rule, regardless of group membership.

  - To add a new custom rule group of custom rules, click **+** > **Upload Custom Rules**. For details, see Uploading Custom Intrusion Rules, on page 21.

  - To change the name or description of a custom rule group, click **Edit**.

  - To delete a custom rule group, click **Delete**. For more information, see Managing Custom Intrusion Rules and Rule Groups, on page 20.

  - To add a new custom rule in a custom rule group, click **+** above the rule table. See Configuring Individual Custom Intrusion Rules, on page 23.

  - To edit, duplicate, delete, or manage group membership for a custom rule, mouse over the right of the rule and click the appropriate button or command. For more information, see Configuring Individual Custom Intrusion Rules, on page 23.

- **List of rules**. You can use the search field to help you find rules using full-text search. You can also select filtering items to search on any combination of GID or SID, show only user defined rules (those you added), show only those rules whose actions are overridden, or simply show rules based on their actions (disabled, alert, drop). The rules are lazy-loaded, so it will take quite a bit of time to scroll through the entire unfiltered list. When filtering the list, click the refresh button to reload the filtered view.

  - To change the action for a rule, click the **Action** cell for the rule and select the new action, to **Alert** only, to **Block** matching traffic, or to **Disable** the rule. The default action for each rule is indicated.

  - To change the action for more than one rule at a time, click the check box in the left column of the rules you want to change, then select the new action from the **Action** drop-down list above the rules table. Click the check box in the GID:SID header to select all the rules in the list. You can change up to 5000 rules at a time.

  - To update the rules within a custom rule group, click **Upload Rule File**. For more information, see Uploading Custom Intrusion Rules, on page 21.

  - To get more information about a rule, click the link in the **GID:SID** cell. The link takes you to Snort.org.

  - To change the rules listed, you can click a child group from the rule group table of contents (not a parent group). You can return to the all-rules list by clicking **ALL RULES** at the top of the rule group list.

  - To change the sorting order, click the table header for a column. The default sorting for the rules is overridden rules first, then drop rules, then alert rules.

  - To see what changes where made in an intrusion rule (LSP) update, select **LSP Update** in the filter field, then select the updates whose changes you want to see, and specify whether you want to see all changes, or just additions or changes to the rules.

Adding or Removing Rule Groups in an Intrusion Policy (Snort 3)

Intrusion rules are organized in local groups. There is a hierarchy to the groups, with parent groups containing related child groups. The rules themselves appear in child groups only: parent groups are simply an organizational construct. A given rule can appear in more than one group.

Any custom rule groups you create are in the User Defined Groups folder. Custom rule groups do not have a hierarchy.

The easiest way to add or remove rules in an intrusion policy is to add or remove groups. Because the rules in a group are logically related, it is highly likely that you will want to use most if not all rules within a given group.

The following procedure explains how to add groups and change the group's security level.

Procedure

Step 1 Choose **Policies** > **Intrusion**.

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

Step 3 (Adding groups.) If the group is not shown in the list of rule groups, click **+** > **Add Existing Rule Group** and do the following:

a) Find the child group.

- A check mark next to a parent group name indicates that all child groups in the parent group are already selected.
- A minus mark next to a parent group name indicates that one or more child group has no enabled rules for this policy. These are the groups you can add.
- A check mark next to a child group name indicates the group is already selected.

b) Select the group you want to add (that is, check its check box).

c) (Optional, does not apply to custom rule groups.) Each group has a default security level depending on the base policy used for the custom policy. If you want to change it, click the security level icon, select a new level, and click **OK**.

Level 1 is the least secure posture, emphasizing connectivity over security, whereas level 4 is the most aggressive, providing maximum security. You can click **View Description** to see an explanation of each level as you select it.

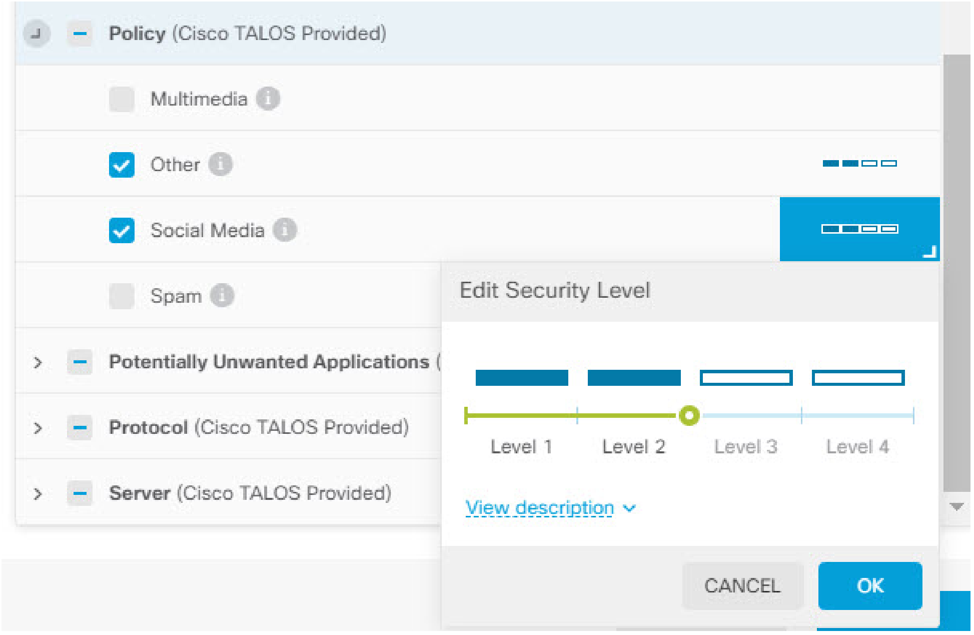
Policy (Cisco TALOS Provided)
□ Multimedia
☑ Other
☑ Social Media
□ Spam
Potentially Unwanted Applications (
Protocol (Cisco TALOS Provided)
Server (Cisco TALOS Provided)
Edit Security Level
Level 1
Level 2
Level 3
Level 4
View description
CANCEL
OK

d) Continue selecting (or deselecting) groups until you have made all of your changes.

e) Click **OK**.

Step 4 (Removing groups.) If you want to disable all the rules within a group, you can use any of the following methods:

- Select the group, then click the **Exclude** link to the far right of the group name, above the list of rules.
- Use the method for adding a group, but instead, deselect the unwanted group (that is, uncheck its check box), and click **OK**.
- You can delete a custom rule group to remove it entirely from the system and from all intrusion policies that use it. Select the group, then click **Delete**.

Changing Intrusion Rule Actions (Snort 3)

Each intrusion policy has the same rules. The difference is the action taken for each rule can be different from policy to policy.

By changing the rule action, you can disable rules that are giving you too many false positives, or you can change whether the rule alerts on or drops matching traffic. You can also enable disabled rules to alert or drop matching traffic.

The easiest method to change rule actions is to change the security level of a rule group. When you change a group's security level, the action of the rules within the group change. This can mean that some rules become enabled (or disabled), or the action can change between alert and drop, based on the security posture you select. However, you can change an individual rule action if that is what you need.

Note The default action for a given rule is based on the overall selection of group and severity. Changing the severity of the group, or excluding the group, can change the default action for the rule.

Before you begin

Custom rule groups do not have security levels. You cannot use the security level technique to change rule actions for custom rules.

Procedure

Step 1 Choose **Policies** > **Intrusion**.

Step 2 Click the view icon (👁) for the policy whose rule actions you want to change.

Step 3 (Recommended method.) Change the security level for a rule group.

a) Click the child rule group in the rule group list.

b) Above the list of rules, click **Edit** next to the security level for the group.

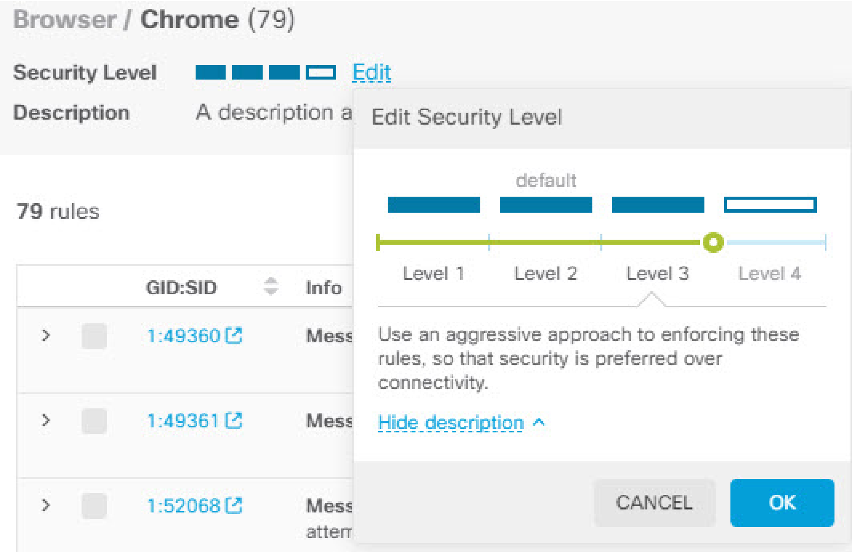

Browser / Chrome (79)

Use an aggressive approach to enforcing these rules, so that security is preferred over connectivity.

Note
If you want to disable all the rules in the group, do not click **Edit**. Instead, click **Exclude** and confirm that you want to exclude the group. The group is not deleted, its rules are simply disabled. Skip the remaining steps.

c) Select the new level for the group. Click **View Description** to see an explanation of each level as you pick it.

Level 1 is the least secure posture, emphasizing connectivity over security, whereas level 4 is the most aggressive, providing maximum security.

d) Click **OK**.

Step 4 (Manual method.) Change the action for one or more rule.

a) Find the rule whose action you want to change.

Use the **Search/Filter** box to search on strings within the rule information. You can also select filtering items to search on any combination of GID or SID, or simply show rules based on their actions (disabled, alert, drop). The rules are lazy-loaded, so it will take quite a bit of time to scroll through the entire unfiltered list. When filtering the list, click the refresh button to reload the filtered view.

Ideally, you can get the Snort identifier (SID) and generator identifier (GID) from an event or from Cisco Technical Support, if you are working with them on an issue. You can then search precisely for the rule.

b) To change the action, do one of the following:

- Change one rule at a time—Click the **Action** column for the rule and select the required action:
  - **Alert**—Create an event when this rule matches traffic, but do not drop the connection.
  - **Drop**—Create an event when this rule matches traffic, and also drop the connection.
  - **Disabled**—Do not match traffic against this rule. No events are generated.
- Change multiple rules at one time—Click the check boxes for the rules you want to change, then click the **Bulk** drop-down above the table and pick the desired action. Click the check box in the GID:SID header to select all the rules in the list. You can change up to 5000 rules at a time.

Managing Custom Intrusion Rules and Rule Groups

The system comes with thousands of intrusion rules defined by Cisco Talos Intelligence Group (Talos). If you know of additional attacks, you can create and upload custom intrusion rules to screen for those attacks, and either alert or drop them. You can also create, edit, and delete rules one at a time.

For uploaded rules, you create the rules offline using a text editor. We recommend that you include a group of custom rules in each text file you upload. You can then easily upload changes to your rules, and either merge new rules into your custom rule groups, or replace your rules with new, edited copies.

Explaining how to create these rules is outside the scope of this document. For detailed information on how to write intrusion rules for Snort, including how to convert Snort 2 rules to Snort 3 format, see the guides on https://snort.org/documents. For example, *Rules Authors Introduction to Writing Snort 3 Rules* at https://snort.org/documents/rules-writers-guide-to-snort-3-rules.

Before you begin

You create custom rule groups during the process of uploading custom rules, as described in Uploading Custom Intrusion Rules, on page 21, or when creating individual rules or managing rule membership. After you create the group, you can manage the group and its contents.

Note that custom groups are available for all intrusion policies, not just the policy you were editing when you created the group. Thus, changes you make to a group are made for all policies. For example, if you delete a custom rule group, it is deleted from all policies and is no longer available for any of them.

Procedure

Step 1 Choose **Policies** > **Intrusion**.

Step 2 Click the edit icon (✎) for a policy.

We recommend that you add custom rules to a custom intrusion policy rather than one of the built-in policies.

Step 3 Do any of the following:

- To create a group, click **+** > **Upload Custom Rules**. See Uploading Custom Intrusion Rules, on page 21.
- To edit a group's name or description, select the group in the group table of contents in the User Defined Groups folder. You can then click **Edit** and make your changes.
- To exclude the group and its rules from the policy, select the group in the group table of contents in the User Defined Groups folder. You can then click **Exclude** to remove the group.
- To delete the group from the system and all policies that use it, select the group in the group table of contents in the User Defined Groups folder. Then click **Delete**. Note that if a rule exists only in the deleted group, it is also deleted from the system. However, if a rule also exists in other custom rule groups that you are not deleting, the rule will remain in those groups.
- To replace or update the rules in a group in bulk, select the group in the group table of contents in the User Defined Groups folder. Then, click **Upload Rule File** next to the Action drop-down list above the group's rule table. The process is the same as that described in Uploading Custom Intrusion Rules, on page 21.
- To create and manage individual rules, and their assignment to rule groups, see Configuring Individual Custom Intrusion Rules, on page 23.

Uploading Custom Intrusion Rules

If you know of attacks that are not currently covered by other rules, you can create and upload custom intrusion rules to screen for those attacks, and either alert or drop them. The action of the imported rules must be either alert or drop, and the default action for the rule is defined by the action in the imported file. Once imported, you can change a rule action and disable a rule if necessary.

You must create these rules offline. In the Firewall Device Manager, you are simply uploading a rules file, you are not configuring the rules directly. The rules file should be a text file. You can use line returns to format the rules to be readable, or put a rule on a single line, and empty lines are allowed. Rule format is explained at snort.org.

For example, an upload file of three rules could look like the following:

```
alert tcp $HOME_NET any -> $EXTERNAL_NET $HTTP_PORTS (
   msg: "My Custom Rule: EXPLOIT-KIT Styx exploit kit landing page request";
   flow:to_server,established;
   http_raw_uri;
   bufferlen:>100;
   http_uri;
   content:"/i.html?",depth 8; pcre:"/\/i\.html\?[a-z0-9]+\=[a-zA-Z0-9]{25}/";
   flowbits:set,styx_landing;
   metadata: copied from talos sid 29452;
   service:http;
   classtype:trojan-activity;
   gid:1;
   sid:1000000;
   rev:1;
)

alert tcp $HOME_NET 8811 -> $EXTERNAL_NET any (
   msg:"My Custom rule: MALWARE-BACKDOOR fear1.5/aciddrop1.0 runtime detection - initial
connection";
   flow:to_client,established;
   flowbits:isset,Fear15_conn.2;
   content:"Drive",nocase;
   metadata:copied from talos sid 7710;
   classtype:trojan-activity;
   gid:1;
   sid:1000001;
   rev:1;
)

alert tcp $EXTERNAL_NET $FILE_DATA_PORTS -> $HOME_NET any (
   msg:"My Custom Rule: INDICATOR-COMPROMISE download of a Office document with embedded
PowerShell";
   flow:to_client,established;
   flowbits:isset,file.doc;
   file_data;
   content:"powershell.exe",fast_pattern,nocase;
   metadata:copied from talos sid 37244;
   classtype:trojan-activity;
   gid:1;
   sid:1000002;
   rev:1;
)
```

Procedure

Step 1 Choose **Policies** > **Intrusion**.

Step 2 Click the edit icon (✎) for a policy.

We recommend that you add custom rules to a custom intrusion policy rather than one of the built-in policies.

Step 3 Do one of the following:

- Above the list of groups, click **+** > **Upload Custom Rules**.
- If you are uploading rules into a custom rule group that you already created, you can select the custom rule group, then click **Upload Rule File** next to the **Action** drop-down list above the group's rule table.

Step 4 Click **Browse** and select your custom rule file, or drag and drop the file into the Upload File dialog box.

Wait for the upload to complete.

Step 5 Select how you want to handle conflicts:

A conflict occurs when a rule you are adding is the same as one already in the system. This should happen only if you are uploading the same rules or edited versions of rules that you previously uploaded.

Select one of these options:

Note
Merge and **Replace** are basically the same thing. Uploaded rules must have higher revision numbers than the ones you already uploaded for any changes to be made to the existing rules. The only difference is that if the upload file is missing rules that are in the targeted custom rule group, the **Replace** option will delete those rules from the rule group. The **Merge** option will leave those "missing" rules in place.

- **Merge**—Any changed rules in the uploaded file that also exist in the selected group will have those changes merged, if the rule in the uploaded file has a higher revision number. Any unchanged rules, or rules in the group that do not have corresponding rules in the upload, will be left unchanged. Any new rules in the upload will be added. This is the default option.
- **Replace**—The rules in the uploaded file will replace the rules in the selected group, if the revision number on the uploaded rule is higher. Any existing rules not in the uploaded file will be deleted from the group. Existing rules whose uploaded version has a revision number that is the same or lower will be left unchanged. Any new rules in the upload will be added.

Step 6 Click + and select the custom rule group for the uploaded rules.

If the custom rule group you want to use does not yet exist, click **Create New Group** and create it now. The new group needs a name and optionally, a description. Then, you can select the new group.

If you are replacing rules, you can select a single group only. If you are merging them, you can select multiple groups.

Step 7 Click **OK**.

The files are uploaded and placed in the new group. You should see a summary of how many rules were uploaded, and how many rules were updated, deleted, or ignored.

If there are errors in the file, the upload will fail. You can click the **Download Error File** link to get more information on the errors.

The group is automatically activated in this intrusion policy. The group and new rules are available to be added to other policies, but the group and rules are not automatically enabled in any other policy. For information on adding groups to other policies, see Adding or Removing Rule Groups in an Intrusion Policy (Snort 3), on page 17.

Configuring Individual Custom Intrusion Rules

You can configure custom intrusion rules one at a time rather than in bulk through file uploads. This method works well when you need to make a quick adjustment to a rule, or you need to create or modify just a few rules at a time.

When configuring intrusion rules, keep the following in mind:

- The GID for all custom rules should be 1.
- The SID for a rule must be unique across all rules in the system. It also must be one million (1000000) or higher.
- If you edit a rule, you must change the rule version. Normally, you increment the version number by 1.

- You can duplicate a Cisco Talos Intelligence Group (Talos) rule to create your own version of the rule, but you still must change the SID of the duplicate to make it unique.

The system will do some validity checking to ensure the rule is well formed, and you will see error messages about any problems. However, the system cannot determine whether the rule is sensible.

For detailed information on how to write intrusion rules for Snort, including how to convert Snort 2 rules to Snort 3 format, see the guides on https://snort.org/documents. For example, *Rules Authors Introduction to Writing Snort 3 Rules* at https://snort.org/documents/rules-writers-guide-to-snort-3-rules.

Procedure

Step 1 Choose **Policies** > **Intrusion**.

Step 2 Click the edit icon (✎) for a policy.

We recommend that you add custom rules to a custom intrusion policy rather than one of the built-in policies.

Step 3 Do one of the following:

- To add an intrusion rule, click the **Add New Intrusion Rule** button (+) above the rule table. When adding a rule, you must select one or more custom rule groups to contain the new rule. You can create new groups while adding the rule if necessary.
- To add a rule by duplicating and editing an existing rule, mouse over the right end of the rule and click the Duplicate (⧉) button. The button appears only on mouse-over. For custom rules, the **Duplicate** command is under the more options (**...**) button.
- To edit a custom rule, find the rule in a custom rule group and click the edit (✎) button for the rule. Your edits apply across all groups in which the rule resides. Make sure you increment the rule version number by at least 1 when making changes.
- To delete a custom rule, click the delete (🗑) button for the rule. The rule is deleted from all rule groups that contain it. If you just want to remove a rule from a group, use the **Manage Group Assignments** option instead of deleting the rule.
- To change which groups contain the rule, click the more options (**...**) button and select **Manage Group Assignments**. You can then add or remove groups. Your changes simply affect group membership, it does not change the rule or delete it.

Step 4 For new rules and groups, add the rule to the policy.

When you create a new group when creating a new rule, or editing an existing rule, that group is not added to your policy automatically, nor is the rule enabled automatically. You are prompted to add the group to the policy you are editing. If you do not add the group while adding or editing the rule, you can add the group later using the following process:

a) Click **+** > **Add Existing Rule Group** above the group table of contents.
b) Find the group under the User Defined Groups folder, select it, and click **OK**.
c) Select the group in the table of contents and verify that the new rule is in the group and has the desired action.

Monitoring Intrusion Policies

You can find intrusion policy statistics on the **Attackers** and **Targets** dashboards on the **Monitoring** page. You must apply an intrusion policy to at least one access control rule to see any information on these dashboards. See Monitoring Traffic and System Dashboards.

To see intrusion events, select **Monitoring** > **Events**, then click the **Intrusion** tab. You can hover over an event and click the **View Details** link to get more information. From the details page, you can click the **View IPS Rule** to go to the rule in the relevant intrusion policy, where you can change the rule action. This can help you reduce the impact of false positives, where a rule is blocking too many good connections, by changing the action from drop to alert. Conversely, you can change an alert rule into a drop rule if you are seeing a lot of attack traffic for a rule.

If you configure a syslog server for the intrusion policy, intrusion events have the message ID 430001.

Examples for Intrusion Policies

The use case chapter includes the following examples of implementing intrusion policies.

- How to Block Threats
- How to Passively Monitor the Traffic on a Network