



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](#)
- [Switching Between Snort 2 and Snort 3, on page 7](#)
- [Configuring Syslog for Intrusion Events, on page 9](#)
- [Managing Intrusion Policies \(Snort 3\), on page 9](#)
- [Managing Intrusion Policies \(Snort 2\), on page 16](#)
- [Monitoring Intrusion Policies, on page 18](#)
- [Examples for Intrusion Policies, on page 18](#)

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.

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 FDM 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.

Status

For Snort 2 rules, Status is a separate column. If you change the default action for a rule, this column shows “Overridden.” Otherwise, the column is empty.

For Snort 3 rules, the “Overridden” status is shown at the bottom of the Action attribute, if you changed it.

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).

- `SSSH_PORTS = 22`.
- `SSSH_SERVERS = $HOME_NET` (meaning any IP address).
- `STELNET_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.

ID	Component
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).)
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.

License Requirements for Intrusion Policies

You must enable the **Threat** 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.
-

Switching Between Snort 2 and Snort 3

Snort is the main inspection engine for the product. Although you can switch Snort versions freely, some intrusion rules in Snort 2.0 might not exist in Snort 3.0, and vice versa. If you changed the rule action for one of these rules, that change is not preserved if you switch to Snort 3 and then back to Snort 2, or back again to Snort 3. Your changes to rule actions for rules that exist in both versions are preserved. Note that the mapping between rules in Snort 3 vs. Snort 2 can be one-to-one or one-to-many, so preservation of changes is done on a best effort basis.

Snort 3.0 does not support all features that Snort 2.0 supports. Following are the unsupported features:

- Virtual routers. You cannot switch to Snort 3.0 if you have configured virtual routers.
- Time-based access control rules. You might have added time range objects to an access control rule using the FTD API. You cannot configure these kinds of rule through the FDM.
- The decryption of TLS 1.1 or lower connections using the SSL Decryption policy.
- The decryption of the following protocols using the SSL Decryption policy: FTPS, SMTPS, IMAPS, POP3S.

If you change the Snort version, the system will perform an automatic deployment to implement the change. You can see the progress in the Task list. The tasks are Snort Version Change and Automatic

Deployment—Snort version toggle. Because of the deployment and the fact that Snort must be stopped and restarted, all existing connections, including VPN, are dropped and must be reestablished, which will cause a momentary traffic loss.



Note If you try to switch Snort versions and the switch fails, you will have pending changes that you cannot discard, and a subsequent toggle attempt will not be allowed. If this happens, you must complete the switch using the ToggleInspectionEngine API, which you can use from the API Explorer. You must set the bypassPendingChangeValidation attribute to TRUE.

Before you begin

To determine which Snort version is currently enabled, use this procedure, or choose **Policies > Intrusion**. Look for the **Snort Version** line above the table. The current version is the first number in the complete version number. For example, 2.9.17-95 is a Snort 2 version.

If the device is in an air-gapped network, consider manually uploading the latest rules package for the new version before switching.

If you downgrade to 2.0, any custom intrusion policies that you created are converted to the base policy used in the custom policy. As far as possible, rule action overrides are retained. If more than one custom policy uses the same base policy, the overrides of the custom policy that is used in the most access control policies are retained, and the overrides for the other custom policies are lost. Access control rules that used these "duplicate" policies will now use the base policy created from your most-used custom policy. All custom policies are deleted. If you want to preserve the custom policies so that you can import them later, after switching back to Snort 3, use the FTD API to export the configuration.

You must deploy any pending changes before you can switch Snort versions.

Procedure

- Step 1** Choose **Device**, then click **View Configuration** in the Updates summary.
Look at the **Intrusion Rule** group. The current Snort version is indicated.
- Step 2** In the **Intrusion Rule** group, you can change the Snort version by clicking **Upgrade to Snort 3.0** or **Downgrade to Snort 2.0**.
- Step 3** When prompted to confirm your action, select the option to get the latest intrusion rules package, then click **Yes**.

We recommend that you get the latest rules package. The system downloads packages for the active Snort version only, so it is unlikely that you have the latest package installed for the Snort version you are switching to.

You must wait until the task to switch versions completes before you can edit intrusion policies.

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 **Edit Logging 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**.
-

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.



Note If you switch to Snort 2, you cannot create custom policies, and the use of intrusion policies differs slightly. Instead of this topic, see [Managing Intrusion Policies \(Snort 2\)](#), on page 16.

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 9](#).
 - Click + to create a new policy. See [Configuring a Custom Intrusion Policy \(Snort 3\), on page 10](#).
 - 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 11](#).
 - 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

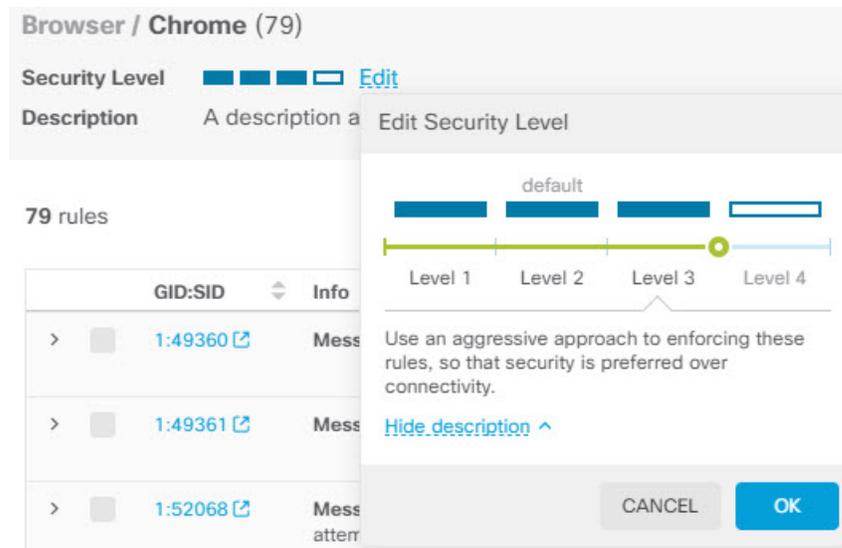
Step 1 Choose **Policies > Intrusion**.

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 + and select the group. See [Adding or Removing Rule Groups in an Intrusion Policy \(Snort 3\)](#), on page 13.
 - 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.



- 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.
- **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, 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 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.

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.

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

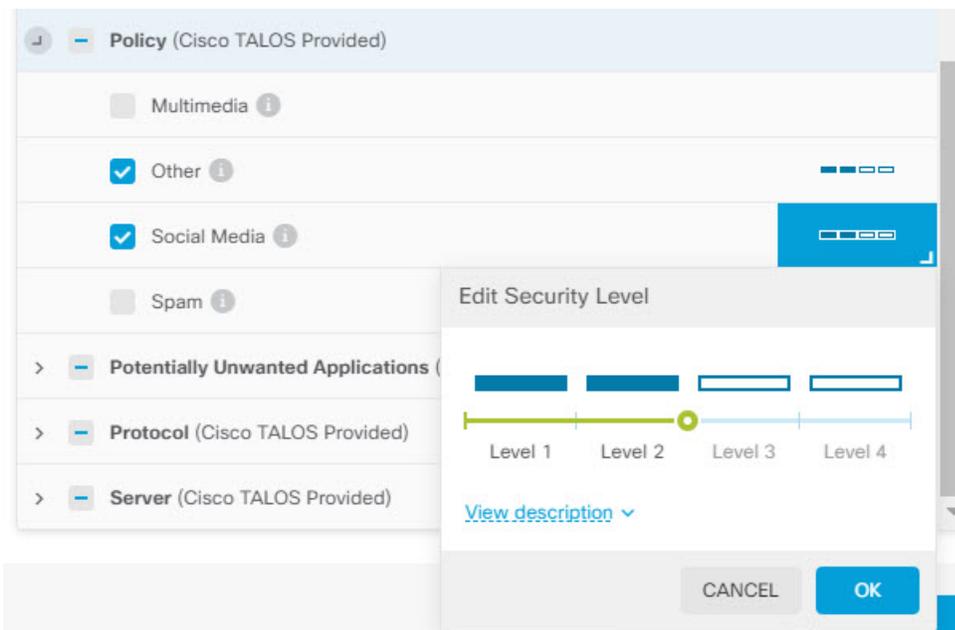
Step 1 Choose **Policies > Intrusion**.

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 + 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.) 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.



- 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**.

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.

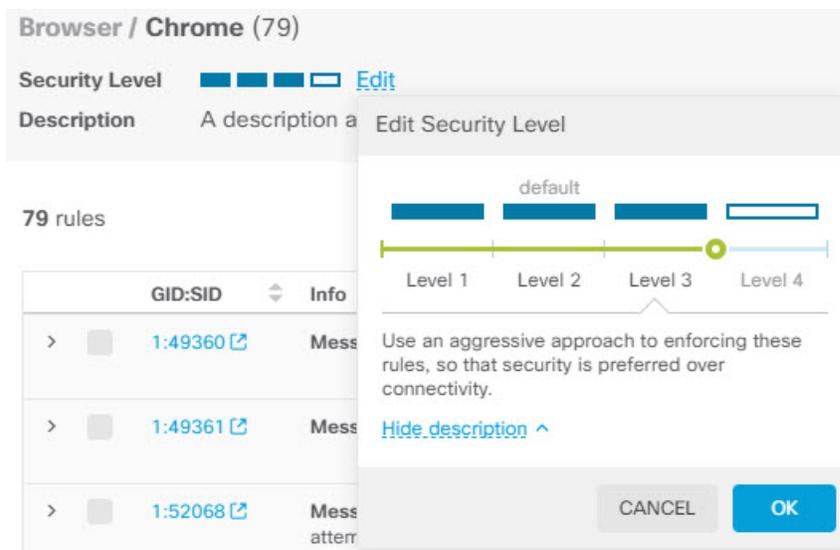
The easiest method to change rule actions is to change the security level of a rule group. When you change a group's security level, the action of the rules within the group change. This can mean that some rules become enabled (or disabled), or the action can change between alert and drop, based on the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

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.
- Click the child rule group in the rule group list.
 - Above the list of rules, click **Edit** next to the security level for the group.



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.

- 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.

- Click **OK**.

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

- 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.

- 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 Intrusion Policies (Snort 2)

You can apply any of the pre-defined intrusion policies. Each of these 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 active 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.

The following topics explain intrusion policies and rule tuning in more detail.

Configure the Inspection Mode for an Intrusion Policy (Snort 2)

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.

Procedure

Step 1 Select **Policies > Intrusion**.

Step 2 Click the tab for the intrusion policy whose inspection mode you want to change.

The **Inspection Mode** is indicated above the rules table.

Step 3 Click the **Edit** link next to the inspection mode, change the mode for the policy, and click **OK**.

The options are:

- **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.

Changing Intrusion Rule Actions (Snort 2)

Each pre-defined 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.

Procedure

Step 1 Select **Policies > Intrusion**.

Step 2 Click the tab for the intrusion policy whose rule actions you want to change.

The pre-defined policies are:

- Connectivity over Security
- Balanced Security and Connectivity
- Security over Connectivity
- Maximum Detection

Step 3 Find the rule whose action you want to change.

The rules are sorted with the overridden ones listed first, and sorted by action within the group of overridden rules. Otherwise, the rules are sorted by GID, then SID.

Use the search box to locate the rule you want to change. 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.

For detailed information about the elements of each rule, see [Intrusion Rule Attributes, on page 3](#).

To search the list:

- Click in the **Search** box to open the search attributes dialog box.
- Enter a combination of Generator ID (**GID**), Snort ID (**SID**), or rule **Action**, and click **Search**.

For example, you could select **Action = Drop** to see all the rules in the policy that will drop matching connections. The text beside the search box indicates how many rules match your criteria, for example, “8937 of 9416 rules found.”

To clear a search criteria, click the x for the criteria in the search box.

Step 4 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.

The default action for the rule is indicated by “(Default)” added to the action. If you change the default, the status column indicates “Overridden” for that rule.

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](#)