



Get Started with Snort 3 Intrusion Policies

This chapter provides information on managing Snort 3 intrusion policies and access control rule configurations for intrusion detection and prevention.

- [Overview of Intrusion Policies, on page 1](#)
- [Prerequisites for Network Analysis and Intrusion Policies, on page 2](#)
- [Create a Custom Snort 3 Intrusion Policy , on page 2](#)
- [Edit Snort 3 Intrusion Policies, on page 3](#)
- [Change the Base Policy of an Intrusion Policy, on page 7](#)
- [Manage Intrusion Policies, on page 8](#)
- [Access Control Rule Configuration to Perform Intrusion Prevention, on page 9](#)

Overview of Intrusion Policies

Intrusion policies are defined sets of intrusion detection and prevention configurations that inspect traffic for security violations and, in inline deployments, can block or alter malicious traffic. Intrusion policies are invoked by your access control policy and are the system's last line of defense before traffic is allowed to its destination.

At the heart of each intrusion policy are the intrusion rules. An enabled rule causes the system to generate intrusion events for (and optionally block) traffic matching the rule. Disabling a rule stops processing of the rule.

The system delivers several base intrusion policies, which enable you to take advantage of the experience of the Cisco Talos Intelligence Group (Talos). For these policies, Talos sets intrusion and inspector rule states (enabled or disabled), as well as provides the initial configurations for other advanced settings.



Tip System-provided intrusion and network analysis policies are similarly named but contain different configurations. For example, the Balanced Security and Connectivity network analysis policy and the Balanced Security and Connectivity intrusion policy work together and can both be updated in intrusion rule updates. However, the network analysis policy governs mostly preprocessing options, whereas the intrusion policy governs mostly intrusion rules.

If you create a custom intrusion policy, you can:

- Tune detection by enabling and disabling rules, as well as by writing and adding your own rules.

- Use Secure Firewall recommendations to associate the operating systems, servers, and client application protocols detected on your network with rules specifically written to protect those assets.

An intrusion policy can drop matching packets and generate intrusion events. To configure an intrusion or preprocessor drop rule, set its state to Block.

When tailoring your intrusion policy, especially when enabling and adding rules, keep in mind that some intrusion rules require that traffic first be decoded or preprocessed in a certain way. Before an intrusion policy examines a packet, the packet is preprocessed according to configurations in a network analysis policy. If you disable a required inspector, the system automatically uses it with its current settings, although the inspector remains disabled in the network analysis policy web interface.



Caution Because preprocessing and intrusion inspection are so closely related, the network analysis and intrusion policies examining a single packet **must** complement each other. Tailoring preprocessing, especially using multiple custom network analysis policies, is an **advanced** task.

After you configure a custom intrusion policy, you can use it as part of your access control configuration by associating the intrusion policy with one or more access control rules or an access control policy's default action. This forces the system to use the intrusion policy to examine certain allowed traffic before the traffic passes to its final destination. A variable set that you pair with the intrusion policy allows you to accurately reflect your home and external networks and, as appropriate, the servers on your network.

Note that by default, the system disables intrusion inspection of encrypted payloads. This helps reduce false positives and improve performance when an encrypted connection matches an access control rule that has intrusion inspection configured.

Refer to the video for additional support and information - [Snort 3 Intrusion Policy Overview](#).

Prerequisites for Network Analysis and Intrusion Policies

To allow the Snort inspection engine to process traffic for intrusion and malware analysis, you must have the IPS license enabled for the Threat Defense device.

You must be an Admin user to manage network analysis, intrusion policies, and perform migration tasks.

Create a Custom Snort 3 Intrusion Policy

-
- Step 1** Choose **Policies > Intrusion**.
 - Step 2** Click **Create Policy**.
 - Step 3** Enter a unique **Name** and, optionally, a **Description**.
 - Step 4** Choose the **Inspection Mode**.

The selected action determines whether intrusion rules block and alert (**Prevention** mode) or only alert (**Detection** mode).

Note Before selecting the prevention mode, you might want block rules to alert only so you can identify rules that cause a lot of false positives.

- Step 5** Choose the **Base Policy**.
You can use either a system-provided policy or an existing policy as your base policy.
- Step 6** Click **Save**.
The new policy has the same settings as its base policy.

What to do next

To customize the policy, see [Edit Snort 3 Intrusion Policies, on page 3](#).

Edit Snort 3 Intrusion Policies

While editing a Snort 3 policy, all the changes are saved instantaneously. No additional action is required to save the changes.

-
- Step 1** Choose **Policies > Intrusion**.
- Step 2** Ensure the **Intrusion Policies** tab is selected.
- Step 3** Click **Snort 3 Version** next to the intrusion policy you want to configure.
- Step 4** Edit your policy:

- Change the mode—Click the **Mode** drop-down to change the inspection mode.

Caution The inspection mode is changed only for the Snort 3 version of the policy. The existing inspection mode is retained in the Snort 2 version as is, which means that your Snort 2 and Snort 3 versions of the policy will have different inspection modes. We recommend you to use this option with caution.

- **Prevention**—Triggered Block rules create an event (alert) and drop the connection.
- **Detection**—Triggered Block rules create an alert.



You can choose the detection mode before going for prevention. For example, before choosing the prevention mode, you might want block rules to alert only, so that you can identify rules that cause a lot of false positives.

- Step 5** Click the **Base Policy** layer that defines the intrusion policy's default settings.
- Search rules—Use the search field to filter the display. You can enter the GID, SID, rule message, or reference info. For example, GID:1; SID:9621—to display only rule 1:962, SID:9621,9622,9623—to display multiple rules with different SIDs. You can also click inside the Search text box to choose any of the following options:
 - apply the filters **Action = Alert**, or **Action: Block**
 - apply the **Disabled Rules** filter
 - show **Custom/User Defined Rules**
 - filter by GID, SID, or GID:SID
 - filter by CVE
 - filter by comment

- View filtered rules—Click any of the **Presets** to view rules that are set to alert, block, disabled, and so on.

Overridden rules indicate the rules where the rule action has been changed from the default action to a different action. Note that, once changed, the rule action status is Overridden even if you change it back to its original default action. However, if you select **Revert to default** from the **Rule Action** drop-down list, the Overridden status is removed.

Advanced Filters provides filter options based on the Lightweight Security Package (LSP) releases, Classifications of intrusions, and Microsoft Vulnerabilities.

- View rule documentation—Click the rule ID or the **Rule Documentation** icon to display Talos documentation for the rule.
- View a rule details—Click the **Expand Arrow** () icon in a rule row to view the rule details.
- Add rule comments—Click **Comment** () under the Comments column to add comments for a rule.

Step 6

Group Overrides—Click the **Group Overrides** layer that lists all the categories of rule groups. The top level parent rule groups with Description, Overrides and Enabled Groups, and so on is displayed. Parent rule groups cannot be updated and are read-only. Only the leaf rule groups can be updated. In each rule group, you can traverse up to the last leaf group. Across each group, you can override, include, and exclude rule groups. In the leaf rule groups, you can:

- Search rule groups—Use the search field to enter keywords and search for rule groups.
- In the left panel, you can choose any of the preset filter options to search for rule groups:
 - All—For displaying all rule groups.
 - Excluded—For excluded groups.
 - Included—For included groups.
 - Overridden—For rule group configuration that is overridden.
- Set the security level for a rule group—Navigate to the required rule group on the left pane and click it. Click **Edit** next to the **Security Level** of the rule group to increase or decrease the security level based on system-defined rule settings.

In the **Edit Security Level** dialog box, you have the option to click **Revert to Default**, which reverts the changes you made.

The management center automatically changes the action for the rules of the rule group for the configured security level. In the **Rule Overrides** layer, notice the count of Block Rules and Disabled Rules in the **Presets** every time you change the security level.

- You can make bulk changes to the security level to change the security level of all rule groups within a particular rule category. Bulk security level applies to rule groups that have more than one rule group. After a bulk update of rule groups, you can still update the security level of any of the associated rule groups within it.

There can be **mixed** security levels within rule groups; **mixed** indicates that the child groups contain a mix of security levels within the parent rule group.

- Include or exclude rule groups—The rule groups displayed are the default rule groups associated with the system-provided base intrusion policy. You can include and exclude rule groups from the intrusion policy. An excluded rule group is removed from the intrusion policy and its rules are not applied on the traffic. For information on uploading custom rules in management center, see [Add Custom Rules to Rule Groups](#).

To exclude a rule group:

- a. Navigate the Rule Groups pane and choose the rule group that you want to exclude.
- b. Click the **Exclude** hyperlink on the right-pane.
- c. Click **Exclude**.

To include a new rule group or multiple rule groups with the uploaded custom rules or a previously excluded rule group:

- a. Click **Add** (+) next to the rule group filter dropdown list.
 - b. Choose all the rule groups you want to add by checking the check box next to it.
 - c. Click **Save**.
- For a leaf rule group, click the icon under the **Override** column header to see the rule action trail, which describes the sequence of overridden rule actions that can be assigned due to the base policy and group overrides for an intrusion rule. Rule actions can be obtained from either the base policy configurations or the user group override. The user group override takes the priority between the two; priority refers to the final overridden action that is assigned to the rule group.
 - Click the rule count (number) under the **Rule Count** column header to see a summary of rules that are part of the rule group.

Step 7 Recommendations—Click the **Recommendations** layer if you want to generate and apply Cisco recommended rules. Recommendations use the host database to enable or disable rules, based on known vulnerabilities.

Step 8 Rule Overrides—Click the **Rule Overrides** layer to choose any of the presets to view rules, which are set to alert, block, disabled, overridden, rewrite, pass, drop, or reject.

- The **Set By** column shows the default set by state (Base Policy) or modified rule state by Group Overrides, Rule Overrides, or Recommendations. The **Set By** column in **All Rules** (in the left pane) shows the trail of rule action override actions based on priority order. The priority order of rule actions is Rule Override > Recommendations > Group Override > Base Policy.
- Modify **Rule Action**—To modify rule actions, choose either of the following:
 - Bulk edit—Choose one or more rules, then choose the required action from the **Rule Action** drop-down list; and click **Save**.
Note Bulk rule action changes are supported only for the first 500 rules.
 - Single rule edit—Choose the action for the rule from the drop-down list in the **Rule Action** column.

Rule actions are:

- **Block**—Generates event, blocks current matching packet and all the subsequent packets in this connection.
- **Alert**—Generates only events for matching packet and does not drop packet or connection.
- **Disable**—Does not match traffic against this rule. No events are generated.
- **Revert to default**—Reverts to the system default action.
- **Pass**—No events are generated, allows packet to pass without further evaluation by any subsequent Snort rules.

Note The Pass action is available only for custom rules and not for system-provided rules.

- **Drop**— Generates event, drops matching packet and does not block further traffic in this connection.
- **Reject**— Generates event, drops matching packet, blocks further traffic in this connection and sends TCP reset if it is a TCP protocol to source and destination hosts.

Behavior of reject in different firewall modes and IP address or source or destination in relation to Client or Server: Snort sends RST packets to both client and server in cases of routed, inline, and bridged interfaces. Snort sends two RST packets. RST packet in clients directions will have source set to server's IP and destination set to client's IP. RST packet in servers direction will have source set to client's IP and destination set to server's IP.

- **Rewrite**— Generates event and overwrites packet contents based on the replace option in the rule.

For IPS rule action logging, see [Rule Action Logging, on page 7](#).

If there is a **React** rule, it is converted to an alert action.

Step 9 Click the **Summary** layer for a holistic view of the current changes to the policy. The policy summary page contains the following information:

- Rule distribution of the policy, that is, active rules, disabled rules, and so on.
- Option to export policy and generate report of the intrusion policy.
- Base policy details.
- Option to generate recommendations.
- Group overrides that shows the list of groups that you have overridden.
- Rule overrides that shows the list of rules that you have overridden.
- In the **Summary** layer, click the ? icon to open a popup window of the Snort helper guide that explains the Snort layering concepts.

To change the base policy, see [Change the Base Policy of an Intrusion Policy, on page 7](#).

Note You can navigate to **Objects > Intrusion Rules** and click the **Snort 3 All Rules** tab and traverse through all the intrusion rule groups. The parent rule group lists the associated child groups and rule count.

What to do next

Deploy configuration changes; see [Deploy Configuration Changes](#).

Rule Group Reporting

The rule groups are reflected in the intrusion events generated and MITRE tactics and techniques are also called out. There are columns for MITRE tactics and techniques and for non-MITRE rule groups for intrusion events. To access the intrusion events, in management center, go to **Analysis > Intrusions > Events**, and click the **Table View of Events** tab. You can also view the intrusion event fields in the **Unified Events** viewer. In the **Analysis** tab, click **Unified Events**.

In the **Intrusion Events** page, the following fields are added for rule group reporting. Note that you must explicitly enable the mentioned columns.

- MITRE ATT&CK
- Rule Group

For information about these fields, see the section *Intrusion Event Fields* in the *Cisco Secure Firewall Management Center Administration Guide, 7.3*.

Rule Action Logging

From Management Center 7.2.0 onwards, in the **Intrusion Events** page, the event in the **Inline Result** column displays the same name as the IPS action applied to the rule, so that you can see the action that was applied on the traffic matching the rule.

For the IPS actions, the following table shows the events that are displayed in the **Inline Result** column of the **Intrusion Events** page and **Action** column for **Intrusion Event Type** in the **Unified Events** page.

IPS Action for Snort 3	Inline Result - Management Center 7.1.0 and earlier	Inline Result -Management Center 7.2.0 onwards
Alert	Pass	Alert
Block	Dropped/Would Have Dropped/Partially Dropped	Block/Would Block/Partial Block
Drop	Dropped/Would have dropped	Drop/Would drop
Reject	Dropped/Would have dropped	Reject/Would reject
Rewrite	Allow	Rewrite



Important

- In case of a rule without the “Replace” option, the **Rewrite** action is displayed as **Would Rewrite**.
- The **Rewrite** action would also be displayed as **Would Rewrite** if the "Replace" option is specified, but the IPS policy is in Detection mode or the device is in Inline-TAP/Passive mode.



Note

In case of backward compatibility (Management Center 7.2.0 managing a Threat Defense 7.1.0 device), the events mentioned are applicable only to the Alert IPS action where **Pass** is displayed as **Alert** for events. For all the other actions, the events for Management Center 7.1.0 are applicable.

Change the Base Policy of an Intrusion Policy

You can choose a different system-provided or custom policy as your base policy.

You can chain up to five custom policies, with four of the five using one of the other four previously created policies as its base policy; the fifth must use a system-provided policy as its base.

-
- Step 1** Choose **Policies > Intrusion**.
- Step 2** Click **Edit** (✎) next to the intrusion policy you want to configure.
- Step 3** Choose a policy from the **Base Policy** drop-down list.
- Step 4** Click **Save**.
-

What to do next

Deploy configuration changes; see [Deploy Configuration Changes](#).

Manage Intrusion Policies

On the Intrusion Policy page (**Policies > Intrusion**) you can view your current custom intrusion policies, along with the following information:

- Number of access control policies and devices are using the intrusion policy to inspect traffic
- In a multidomain deployment, the domain where the policy was created

In a multidomain deployment, the system displays policies created in the current domain, which you can edit. It also displays policies created in ancestor domains, which you cannot edit. To view and edit policies created in a lower domain, switch to that domain.

-
- Step 1** Choose **Policies > Intrusion**.
- Step 2** Manage your intrusion policy:
- Create — Click **Create Policy**; see [Create a Custom Snort 3 Intrusion Policy](#), on page 2.
 - Delete — Click **Delete** (■) next to the policy you want to delete. The system prompts you to confirm and informs you if another user has unsaved changes in the policy. Click **OK** to confirm.
If the controls are dimmed, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.
 - Edit intrusion policy details — Click **Edit** (✎) next to the policy you want to edit. You can edit the **Name**, **Inspection Mode**, and the **Base Policy** of the intrusion policy.
 - Edit intrusion policy settings — Click **Snort 3 Version**; see [Edit Snort 3 Intrusion Policies](#), on page 3.
 - Export — If you want to export an intrusion policy to import on another management center, click **Export**; see the *Exporting Configurations* topic in the latest version of the *Cisco Secure Firewall Management Center Configuration Guide*.
 - Deploy — Choose **Deploy > Deployment**; see [Deploy Configuration Changes](#).
 - Report — Click **Report**; see the *Generating Current Policy Reports* topic in the latest version of the *Cisco Secure Firewall Management Center Configuration Guide*. Generates two reports, one for each policy version.
-

Access Control Rule Configuration to Perform Intrusion Prevention

An access control policy can have multiple access control rules associated with intrusion policies. You can configure intrusion inspection for any Allow or Interactive Block access control rule, which permits you to match different intrusion inspection profiles against different types of traffic on your network before it reaches its final destination.

Whenever the system uses an intrusion policy to evaluate traffic, it uses an associated *variable set*. Variables in a set represent values commonly used in intrusion rules to identify source and destination IP addresses and ports. You can also use variables in intrusion policies to represent IP addresses in rule suppressions and dynamic rule states.



Tip Even if you use system-provided intrusion policies, Cisco **strongly** recommends you configure the system's intrusion variables to accurately reflect your network environment. At a minimum, modify default variables in the default set.

Understanding System-Provided and Custom Intrusion Policies

Cisco delivers several intrusion policies with the system. By using system-provided intrusion policies, you can take advantage of the experience of the Cisco Talos Intelligence Group (Talos). For these policies, Talos sets intrusion and preprocessor rule states, as well as provides the initial configurations for advanced settings. You can use system-provided policies as-is, or you can use them as the base for custom policies. Building custom policies can improve the performance of the system in your environment and provide a focused view of the malicious traffic and policy violations occurring on your network.

Connection and Intrusion Event Logging

When an intrusion policy invoked by an access control rule detects an intrusion and generates an intrusion event, it saves that event to the Management Center. The system also automatically logs the end of the connection where the intrusion occurred to the Management Center database, regardless of the logging configuration of the access control rule.

Access Control Rule Configuration and Intrusion Policies

The number of unique intrusion policies you can use in a single access control policy depends on the model of the target devices; more powerful devices can handle more. Every unique **pair** of intrusion policy and variable set counts as one policy. Although you can associate a different intrusion policy-variable set pair with each Allow and Interactive Block rule (as well as with the default action), you cannot deploy an access control policy if the target devices have insufficient resources to perform inspection as configured.

Configure an Access Control Rule to Perform Intrusion Prevention

You must be an Admin, Access Admin, or Network Admin to perform this task.

-
- Step 1** In the access control policy editor, create a new rule or edit an existing rule; see the *Access Control Rule Components* topic in the latest version of the *Cisco Secure Firewall Management Center Configuration Guide*.
- Step 2** Ensure the rule action is set to **Allow**, **Interactive Block**, or **Interactive Block with reset**.
- Step 3** Click **Inspection**.
- Step 4** Choose a system-provided or a custom intrusion policy, or choose **None** to disable intrusion inspection for traffic that matches the access control rule.
- Step 5** If you want to change the variable set associated with the intrusion policy, choose a value from the **Variable Set** drop-down list.
- Step 6** Click **Save** to save the rule.
- Step 7** Click **Save** to save the policy.
-

What to do next

Deploy configuration changes; see [Deploy Configuration Changes](#).