



Getting Started with Intrusion Policies

The following topics explain how to get started with intrusion policies:

- [Intrusion Policy Basics, on page 1](#)
- [License Requirements for Intrusion Policies, on page 2](#)
- [Requirements and Prerequisites for Intrusion Policies, on page 3](#)
- [Managing Intrusion Policies, on page 3](#)
- [Custom Intrusion Policy Creation, on page 4](#)
- [Editing Snort 2 Intrusion Policies, on page 5](#)
- [Access Control Rule Configuration to Perform Intrusion Prevention, on page 6](#)
- [Drop Behavior in an Inline Deployment, on page 7](#)
- [Drop Behavior in a Dual System Deployment, on page 8](#)
- [Intrusion Policy Advanced Settings, on page 9](#)
- [Optimizing Performance for Intrusion Detection and Prevention, on page 10](#)

Intrusion Policy Basics

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 Talos Intelligence Group. For these policies, Talos sets intrusion and preprocessor 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 Cisco recommendations to associate the operating systems, servers, and client application protocols detected on your network with rules specifically written to protect those assets.
- Configure various advanced settings such as external alerting, sensitive data preprocessing, and global rule thresholding.
- Use layers as building blocks to efficiently manage multiple intrusion policies.

In an inline deployment, an intrusion policy can block and modify traffic:

- *Drop rules* can drop matching packets and generate intrusion events. To configure an intrusion or preprocessor drop rule, set its state to Drop and Generate Events.
- Intrusion rules can use the `replace` keyword to replace malicious content.

For intrusion rules to affect traffic, you must correctly configure drop rules and rules that replace content, as well as correctly deploy managed devices inline, that is, with inline interface sets. Finally, you must enable the intrusion policy's *drop behavior*, or **Drop when Inline** setting.

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 preprocessor, the system automatically uses it with its current settings, although the preprocessor 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.

License Requirements for Intrusion Policies

Threat Defense License

IPS

Classic License

Protection

Requirements and Prerequisites for Intrusion Policies

Model Support

Any.

Supported Domains

Any

User Roles

- Admin
- Intrusion Admin

Managing Intrusion Policies

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

- the time and date the policy was last modified (in local time) and the user who modified it
- whether the **Drop when Inline** setting is enabled, which allows you to drop and modify traffic in an inline deployment. An inline deployment could be configurations that are deployed to devices using routed, switched, or transparent interfaces, or inline interface pairs.
- which access control policies and devices are using the intrusion policy to inspect traffic
- whether a policy has unsaved changes, as well as information about who (if anyone) is currently editing the policy
- 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.

Procedure

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


Step 2 Manage your intrusion policy:



- Compare—Click **Compare Policies**; see [Comparing policies](#).
- Create — Click **Create Policy**; see:
 - [Creating a Custom Snort 2 Intrusion Policy, on page 4](#) for Snort 2 policies.
 - [Creating a Custom Snort 3 Intrusion Policy](#) topic in the latest version of the [Cisco Secure Firewall Management Center Snort 3 Configuration Guide](#) for Snort 3 policies.

- **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** — Choose:
 - **Snort 2 Version**; see [Editing Snort 2 Intrusion Policies, on page 5](#).
 - **Snort 3 Version**; see *Editing Snort 3 Intrusion Policies* topic in the latest version of the [Cisco Secure Firewall Management Center Snort 3 Configuration Guide](#).

If **View** () appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

- **Export** — If you want to export an intrusion policy to import on another Secure Firewall Management Center, click **YouTube EDU** (); see *Exporting Configurations* in the [Cisco Secure Firewall Management Center Administration Guide](#).
- **Deploy**—Choose **Deploy > Deployment**; see [Deploy Configuration Changes](#).
- **Report**—Click **Report** (); see [Generate Current Policy Reports](#).

Custom Intrusion Policy Creation

When you create a new intrusion policy you must give it a unique name, specify a base policy, and specify drop behavior.

The base policy defines the intrusion policy's default settings. Modifying a setting in the new policy overrides—but does not change—the settings in the base policy. You can use either a system-provided or custom policy as your base policy.

Creating a Custom Snort 2 Intrusion Policy

Procedure

- Step 1** Choose **Policies > Access Control > Intrusion**.
- Step 2** Click **Create Policy**. If you have unsaved changes in another policy, click **Cancel** when prompted to return to the Intrusion Policy page.
Ensure the **Intrusion Policies** tab is selected.
- 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**).

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

Related Topics

- [Intrusion Rules in Layers](#)
- [Conflicts and Changes: Network Analysis and Intrusion Policies](#)

Editing Snort 2 Intrusion Policies

Procedure

- Step 1** Choose **Policies > Access Control > Intrusion**.
- Step 2** Ensure the **Intrusion Policies** tab is selected.
- Step 3** Click **Snort 2 Version** next to the intrusion policy you want to configure.
- Step 4** Edit your policy:
- Change the base policy—Choose a base policy from the **Base Policy** drop-down list; see [Changing the Base Policy](#).
 - Configure advanced settings—Click **Advanced Settings** in the navigation panel; see [Intrusion Policy Advanced Settings, on page 9](#).
 - Configure Cisco recommended intrusion rules—Click **Cisco Recommendations** in the navigation panel; see [Generating and Applying Cisco Recommendations](#).
 - Drop behavior in an inline deployment—Check or clear **Drop when Inline**; see [Setting Drop Behavior in an Inline Deployment, on page 8](#).
 - Filter rules by recommended rule state—After you generate recommendations, click **View** next to each recommendation type. Click **View Recommended Changes** to view all recommendations.
 - Filter rules by current rule state—Click **View** next to each rule state type (generate events, drop and generate events); see [Intrusion Rule Filters in an Intrusion Policy](#).
 - Manage policy layers—Click **Policy Layers** in the navigation panel; see [Layer Management](#).
 - Manage intrusion rules—Click **Manage Rules**; see [Viewing Intrusion Rules in an Intrusion Policy](#).
 - View settings in base policy—Click **Manage Base Policy**; see [The Base Layer](#).
- Step 5** To save changes you made in this policy since the last policy commit, choose **Policy Information**, then click **Commit Changes**.
- If you leave the policy without committing changes, changes since the last commit are discarded if you edit a different policy.

What to do next

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

Related Topics

- [Generating and Applying Cisco Recommendations](#)
- [Configuring Intrusion Rules in Layers](#)
- [Conflicts and Changes: Network Analysis and Intrusion Policies](#)

Intrusion Policy Changes

When you create a new intrusion policy, it has the same intrusion rule and advanced settings as its base policy.

The system caches one intrusion policy per user. While editing an intrusion policy, if you choose any menu or other path to another page, your changes stay in the system cache even if you leave the page.

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 Talos Intelligence Group. 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 Secure Firewall Management Center. The system also automatically logs the end of the connection where the intrusion occurred to the Secure Firewall Management Center database, regardless of the logging configuration of the access control rule.

Related Topics

- [Predefined Default Variables](#)

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.

Configuring an Access Control Rule to Perform Intrusion Prevention

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

Procedure

- Step 1** In the access control policy editor, create a new rule or edit an existing rule; see [Access Control Rule Components](#).
 - 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 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](#).

Related Topics

- [Variable Set](#)
- [Snort Restart Scenarios](#)

Drop Behavior in an Inline Deployment

If you want to assess how your configuration would function in an inline deployment (that is, where relevant configurations are deployed to devices using routed, switched, or transparent interfaces, or inline interface pairs) without actually affecting traffic, you can disable drop behavior. In this case, the system generates intrusion events but does not drop packets that trigger the drop rules. When you are satisfied with the results, you can enable drop behavior.

Note that in passive or inline deployments in tap mode, the system cannot affect traffic regardless of the drop behavior. In a passive deployment, rules set to **Drop and Generate Events** behave identically to rules set to **Generate Events**. The system generates intrusion events but cannot drop packets.



Note Suppose a file Block action causes a Block or Pending file policy verdict on a packet, and later, an IPS event is generated on the same packet. In that case, the IPS event is marked as Dropped instead of Would have dropped even if the IPS policy is in detection mode (IDS).



Note To block the transfer of malware over FTP, you must not only correctly configure malware defense, but also enable **Drop when Inline** in your access control policy's default intrusion policy.

When you view intrusion events, workflows can include the *inline result*, which indicates whether traffic was actually dropped, or whether it only would have dropped.

Setting Drop Behavior in an Inline Deployment

Procedure

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

Step 2 Click **Snort 2 Version** next to the policy you want to edit.

If **View** (👁) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 3 Set the policy's drop behavior:

- Check the **Drop when Inline** check box to allow intrusion rules to affect traffic and generate events.
- Clear the **Drop when Inline** check box to prevent intrusion rules from affecting traffic while still generating events.

Step 4 Click **Commit Changes** to save changes you made in this policy since the last policy commit.

If you leave the policy without committing changes, changes since the last commit are discarded if you edit a different policy.

What to do next

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

Drop Behavior in a Dual System Deployment

When there are two systems connected back to back in a network, it is normal to see the first system drop events and still record a drop or "would have dropped" event on the second system. The first system decides to drop the packets by the time it scans the last packet of the file, while the second system also investigates and identifies the traffic as "to be dropped".

For example, a 5 packet HTTP GET request whose first packet triggers a rule is blocked by the first system and only the last packet is dropped. The second system receives only 4 packets and the connection gets dropped, but when the second system finally flushes the partial GET request while it is pruning the session, it triggers the same rule with "would have dropped" as the inline result.

Intrusion Policy Advanced Settings

An intrusion policy's *advanced settings* require specific expertise to configure. The base policy for your intrusion policy determines which advanced settings are enabled by default and the default configuration for each.

When you choose **Advanced Settings** in the navigation panel of an intrusion policy, the policy lists its advanced settings by type. On the Advanced Settings page, you can enable or disable advanced settings in your intrusion policy, as well as access advanced setting configuration pages. An advanced setting must be enabled for you to configure it.

When you disable an advanced setting, the sublink and **Edit** link no longer appear, but your configurations are retained. Note that some intrusion policy configurations (sensitive data rules, SNMP alerts for intrusion rules) require enabled and correctly configured advanced settings.

Modifying the configuration of an advanced setting requires an understanding of the configuration you are modifying and its potential impact on your network.

Specific Threat Detection

The sensitive data preprocessor detects sensitive data such as credit card numbers and Social Security numbers in ASCII text.

Note that other preprocessors that detect specific threats (back orifice attacks, several portscan types, and rate-based attacks that attempt to overwhelm your network with excessive traffic) are configured in network analysis policies.

Intrusion Rule Thresholds

Global rule thresholding can prevent your system from being overwhelmed with a large number of events by allowing you to use thresholds to limit the number of times the system logs and displays intrusion events.

External Responses

In addition to the various views of intrusion events in the web interface, you can enable logging to system log (syslog) facilities or send event data to an SNMP trap server. Per policy, you can specify intrusion event notification limits, set up intrusion event notification to external logging facilities, and configure external responses to intrusion events.

Note that in addition to these per-policy alerting configurations, you can globally enable or disable email alerting on intrusion events for each rule or rule group. Your email alert settings are used regardless of which intrusion policy processes a packet.

Related Topics

[Sensitive Data Detection Basics](#)

[Global Rule Thresholding Basics](#)

Optimizing Performance for Intrusion Detection and Prevention

If you want the system to perform intrusion detection and prevention but do not need to take advantage of discovery data, you can optimize performance by disabling new discovery as described below.

Before you begin

To perform this task, you must have one of the following user roles:

- Admin, Access Admin, or Network Admin for access control.
- Admin or Discovery Admin for network discovery.

Procedure

- Step 1** Modify or delete rules associated with the access control policy deployed at the target device. None of the access control rules associated with that device can have user, application, or URL conditions; see [Create and Edit Access Control Rules](#).
- Step 2** Delete all rules from the network discovery policy for the target device; see [Configuring Network Discovery Rules](#).
- Step 3** Deploy the changed configuration to the target device; see [Deploy Configuration Changes](#).
-