

# Access Control Policies

These topics describe how to work with access control policies:


- Access control policies, on page 1
- Requirements and prerequisites for access control policies, on page 6
- Manage access control policies, on page 6
- History for access control policies, on page 22


# Access control policies

An access control policy is a hierarchical policy that allows you to specify, inspect, and log network traffic. You can use the characteristics of each connection to allow, trust, block, or monitor the connection.

Each managed device can be assigned to one access control policy. The data that the policy's assigned (also known as targeted) devices collect about your network traffic can be used to filter and control that traffic based on:

- simple, easily determined transport and network layer characteristics: source and destination, port, protocol, and so on
- the latest contextual information on the traffic, including characteristics such as reputation, risk, business relevance, application used, or URL visited
- realm, user, user group, or ISE attribute
- custom Security Group Tag (SGT)
- characteristics of encrypted traffic; you can also decrypt this traffic for further analysis
- whether unencrypted or decrypted traffic contains a prohibited file, detected malware, or intrusion attempt
- time and day (on supported devices)

Each type of traffic inspection and control occurs where it makes the most sense for maximum flexibility and performance. For example, reputation-based blocking uses simple source and destination data, so it can block prohibited traffic early in the process. In contrast, detecting and blocking intrusions and exploits is a last-line of defense.

# Access control policy components

This list describes the main components of an access control policy.

- Name and description: Each access control policy must have a unique name. A description is optional.
- Inheritance settings: Policy inheritance allows you to create a hierarchy of access control policies. A parent (or base) policy defines and enforces default settings for its descendants. A policy's inheritance settings allow you to select its base policy. You can also lock settings in the current policy to force any descendants to inherit them. Descendant policies can override unlocked settings.
- Policy assignment: Every access control policy identifies which devices use it. Each device can use only one access control policy.
- Rules: Access control rules provide a granular method of handling network traffic. Rules are numbered, starting at 1, including rules inherited from ancestor policies. The system matches traffic to rules from the top down, ascending by rule number.

  Traffic is handled by the first rule for which all conditions match. Conditions can be simple or complex, and often depend on licensing.
- Default action: The default action determines how the system handles and logs traffic that is not handled by any other access control configuration. It can block or trust all traffic without further inspection, or inspect traffic for intrusions and discovery data. Although an access control policy can inherit its default action from an ancestor policy, you cannot enforce this inheritance.
- Security Intelligence: This feature acts as the first line of defense against malicious internet content. It enables blocking of connections based on IP address, URL, and domain name reputation intelligence. To ensure continual access to vital resources, you can override block list entries with custom do not block list entries.
- HTTP responses: When a user's website request is blocked, you can display either a generic response page or a custom page. Alternatively, you can warn users but allow them to proceed to the requested site.
- Logging: Logging settings enable configuration of default syslog destinations for the policy. These settings apply to the policy and all included decryption, prefilter, and intrusion policies unless overridden in individual rules or policies.
- Advanced access control options: Advanced policy settings typically require little modification, and defaults are often suitable. Modifiable options include traffic preprocessing, decryption, identity, and various performance configurations.

# Access control policy default actions

The default action is the access control policy configuration that determines how the system handles and logs traffic that is not handled by any other access control configuration. The default action can block or trust all traffic without further inspection, or inspect traffic for intrusions and discovery data.

In a simple access control policy, the default action specifies how a device handles all traffic. In a more complex policy, the default action handles traffic that:

- is not fast pathed by the prefilter policy
- is not on a Security Intelligence block list

- is not blocked by the decryption policy (encrypted traffic only)
- matches none of the rules in the policy (except monitor rules, which match and log—but do not handle or inspect—traffic)

**Note** You cannot perform file or malware inspection on traffic handled by the default action. Logging of connections handled by the default action is initially disabled but can be enabled.

If you are using policy inheritance, the default action for the lowest-level descendant determines how traffic is handled. An access control policy can inherit its default action from the base policy, but this inheritance is not enforced.

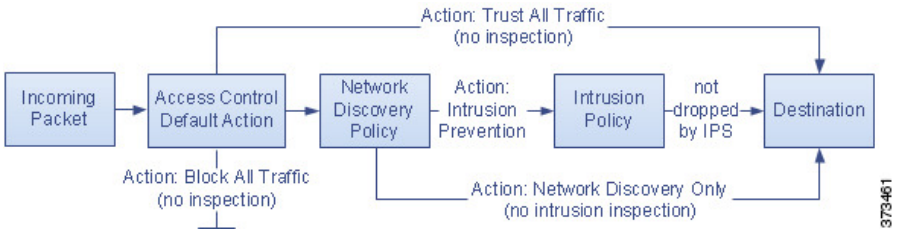
This table describes the types of inspection you can perform on traffic handled by each default action.

***Table 1: Types of default actions and inspection options***

| Default action | Effect on traffic | Inspection type and policy |
|---|---|---|
| Access Control: Block All Traffic | Block without further inspection. | None. |
| Access Control: Trust All Traffic | Trust (allow to its final destination without further inspection). | None. |
| Intrusion Prevention | Allow, as long as it is passed by the intrusion policy you specify. | Intrusion, using the specified intrusion policy and associated variable set.<br>Discovery, using the network discovery policy. |
| Network Discovery Only | Allow. | Discovery only, using the network discovery policy. |
| Inherit from base policy | Defined in base policy. | Defined in base policy. |

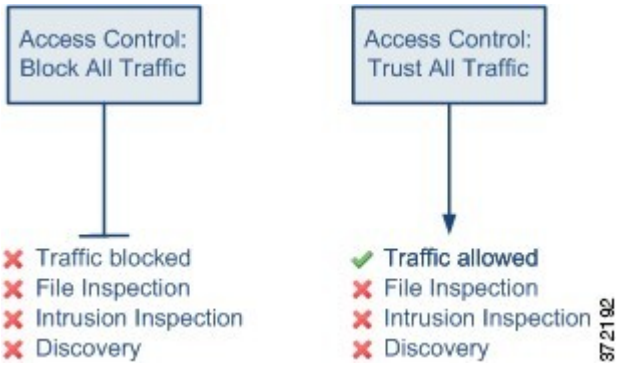
This diagram illustrates the table.

***Figure 1: Default action configuration options***




This diagram illustrates the **Block All Traffic** and **Trust All Traffic** default actions.

*Figure 2: Block all traffic and trust all traffic default actions*




This diagram illustrates the **Intrusion Prevention** and **Network Discovery Only** default actions.

*Figure 3: Intrusion prevention and network discovery only default actions*




**Note** The purpose of **Network Discovery Only** is to improve performance in a discovery-only deployment. Different configurations can disable discovery if you are only interested in intrusion detection and prevention.

# Access control policy inheritance

Access control policy inheritance is a hierarchical policy-based implementation that allows child policies to inherit rules and settings from ancestor or base policies. This structure ensures that mandatory rules defined in parent policies are consistently enforced, while providing the flexibility for child policies to selectively override default configurations.

### Access control policy hierarchy structure

You can nest access control policies, where each policy inherits the rules and settings from an ancestor policy. You can enforce this inheritance, or allow lower-level policies to override their ancestors.

Just as you create a domain hierarchy, you can create a corresponding hierarchy of access control policies. A descendant, or child, access control policy inherits rules and settings from its direct parent, or base, policy. That base policy may have its own parent policy from which it inherits rules and settings, and so on.

An access control policy's rules are nested between its parent policy's Mandatory and Default rule sections. This implementation enforces Mandatory rules from ancestor policies, but allows the current policy to write rules that preempt Default rules from ancestor policies.

You can lock these settings to enforce them in all descendant policies:

- Security Intelligence—Connections that are allowed or blocked based on the latest reputation intelligence for IP addresses, URLs, and domain names.
- HTTP response pages—Displaying a custom or system-provided response page when you block a user's website request.
- Advanced settings—Specifying associated subpolicies, network analysis settings, performance settings, and other general options.

When using policy inheritance, the default action for the lowest-level descendant determines final traffic handling. Although an access control policy can inherit its default action from an ancestor policy, you cannot enforce this inheritance.

### Policy inheritance and multitenancy

Access control's hierarchical policy-based implementation complements multitenancy.

In a typical multidomain deployment, access control policy hierarchy corresponds to domain structure, and you apply the lowest-level access control policy to managed devices. This implementation allows selective access control enforcement at a higher domain level, while lower-level domain administrators can tailor deployment-specific settings. (You must use roles, not policy inheritance and enforcement alone, to restrict administrators in descendant domains.)

For example, as a Global domain administrator for your organization, you can create an access control policy at the Global level. You can then require that all your devices, which are divided into subdomain by function, use that Global-level policy as a base policy.

When subdomain administrators log into the Secure Firewall Management Center to configure access control, they can deploy the Global-level policy as-is. Or, they can create and deploy a descendant access control policy within the boundaries of the Global-level policy.

**Note** Although the most useful implementation of access control inheritance and enforcement complements multitenancy, you can create a hierarchy of access control policies within a single domain. You can also assign and deploy access control policies at any level.

# Rule and other policy warnings

Policy and rule editors use icons to mark configurations that could adversely affect traffic analysis and flow. Depending on the issue, the system may warn you when you deploy or prevent you from deploying entirely.

**Note** Move your pointer over an icon to read the warning, error, or informational text.

*Table 2: Policy error icons*

| Icon | Description | Example |
|---|---|---|
| **Errors** ( ) | If a rule or configuration has an error, you cannot deploy until you correct the issue, even if you disable any affected rules. | A rule that performs category and reputation-based URL filtering is valid until you assign a device that does not have a URL Filtering license. At that point, an error icon appears next to the rule, and you cannot deploy until you edit or delete the rule, unassign the device, or enable the license. |
| **Warning** ( ) | You can deploy a policy that displays rule or other warnings. However, misconfigurations marked with warnings have no effect.<br>If you disable a rule with a warning, the warning icon disappears. It reappears if you enable the rule without correcting the underlying issue. | Preempted rules or rules that cannot match traffic due to misconfiguration have no effect. This includes conditions using empty object groups, application filters that match no applications, excluded LDAP users, invalid ports, and so on.<br>However, if a warning icon marks a licensing error or model mismatch, you cannot deploy until you correct the issue. |
| **Information** ( ) | Information icons convey helpful information about configurations that may affect the flow of traffic. These issues do not prevent you from deploying. | The system might skip matching the first few packets of a connection against some rules, until the system identifies the application or web traffic in that connection. This allows connections to be established so that applications and HTTP requests can be identified. |
| **Rule Conflict** ( ) | When you enable rule conflict analysis, this icon appears in the rule table for rules that have conflicts. | Conflicts include redundant rules, redundant objects, and shadowed rules. Redundant and shadowed rules do not match traffic because previous rules would already match the criteria. Redundant objects make your rules unnecessarily complex. |

# Requirements and prerequisites for access control policies

Ensure that your environment meets these requirements before configuring access control policies.

- You must have **Administrator**, **Network Admin**, or **Access Admin** privileges to configure access control policies.
-

# Manage access control policies

Configure and manage access control policies by editing system-provided policies and creating custom access control policies.

**Procedure**

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

At the top of the page, there are convenient links to some related features: Object management, Intrusion policies, Network Analysis policies, DNS policies, and policy Import/Export.

**Step 2** These actions help you to manage your policy:

- Create—Click **New Policy**; see Creating a basic access control policy, on page 7.
- Edit—Click **Edit** ( ); see Editing an access control policy, on page 8
- Delete—Click **Delete** ( ). You must remove any device assignments before deleting a policy.
- Copy—Click **Copy** ( ). Device assignments are not retained in the copy.
- Lock or unlock a policy—See Lock an access control policy, on page 10.

**Step 3** These options help you to analyze, monitor, and customize your policies:

- Inheritance—Click **Plus** next to a policy with descendants to expand your view of the policy's hierarchy.
- Report—Click **Report** ( ).. The report is generated as a background process. Go to the message/notifications center and look in the tasks list. When the report is finished, you can download it from the notification.

# Creating a basic access control policy

When you create a new access control policy, it contains default actions and settings. After creating the policy, you are immediately placed in an edit session so that you can adjust the policy to suit your requirements.

**Procedure**

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

**Step 2** Click **New Policy**.

**Step 3** Enter a unique **Name** and, optionally, a **Description**.

**Step 4** Optionally, choose a base policy.

If an access control policy is enforced on your domain, this step is not optional. You must choose the enforced policy or one of its descendants as the base policy.

If you select a base policy, the base policy defines the default action and you cannot select a new one in this dialog box. Logging for connections handled by the default action depends on the base policy.

**Step 5** When you do not select a base policy, specify the initial **Default Action**:

- **Block all traffic** creates a policy with the **Access Control: Block All Traffic** default action.
- **Intrusion Prevention** creates a policy with the **Intrusion Prevention: Balanced Security and Connectivity** default action, associated with the default intrusion variable set.
- **Network Discovery** creates a policy with the **Network Discovery Only** default action.

When you select a default action, logging of connections handled by the default action is initially disabled. You can enable it later when you edit the policy.

**Note**
If you want to trust all traffic by default, or if you chose a base policy and do not want to inherit the default action, you can change the default action later.

**Step 6** Optionally, choose the devices to assign to the policy. To narrow the devices that appear, enter a search string.

If you want to deploy this policy immediately, you must perform this step.

**Step 7** Click **Save**.

The new policy opens for edit. You can add rules to it and make other changes as needed. See Editing an access control policy, on page 8.

# Editing an access control policy

When you edit an access control policy, you should lock it to ensure that your changes do not get overridden by another person who might edit it simultaneously.

You can only edit access control policies that were created in the current domain. Also, you cannot edit settings that are locked by an ancestor access control policy.

To protect the privacy of your session, a warning appears after 30 minutes of inactivity on the policy editor. After 60 minutes, the system discards your changes.



**Note** If you do not lock the policy, consider the following: Only one person should edit a policy at a time, using a single browser window. Although more than one user can edit an access control policy at a time, when one user saves changes, all changes made by the other users are immediately deleted and these users are moved out of edit mode and back to the access control policy listing page. These other users must start again.

**Procedure**

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

**Step 2** Click **Edit** ( ) next to the access control 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** Edit your access control policy.

**Note**
You can operate on multiple rules at one time by selecting their checkboxes in the left column, then selecting the action you want to perform from the **Select Bulk Rule Actions** drop-down list next to the search box. Bulk editing is available for enabling and disabling, copying, moving, deleting, and editing rules, or viewing hit counts or related events.

You can change the following settings or perform these actions:

- Name and description—Click **Edit** ( ) next to the name, make your changes, and click **Save** .
- Default action and settings—Choose a value from the **Default Action** drop-down list, then click **Cog** ( ) , make your changes to the settings, and click **OK** . For detailed information, see Configure the access control default action, on page 11 .
- Associated policies—To edit or change policies in the packet flow, click the policy type in the packet flow representation below the policy name. You can select the **Prefilter Rules** , **Decryption** , **Security Intelligence** , and **Identity** policies. When necessary, click **Access Control** to return to the access control rules.
- Policy assignment—To identify the managed devices targeted by this policy, or enforce this policy in a subdomain, click the **Targeted: x devices** link.
- Rules—To manage access control rules, and to inspect and block malicious traffic using intrusion and file policies, click **Add Rule** , or right-click an existing rule and select **Edit** or another appropriate action. The actions are also available from the **More** ( ) button for each rule. See Create and edit access control rules .
- Layout—Use the **Grid/Table View** icon above the list of rules to change the layout. Grid view provides color-coded objects in an easy-to-see layout. Table view provides a summary list so that you can see more rules at once. You can freely switch views without impacting the rules.
- Columns (table view only)—Click the **Show/Hide Columns** icon above the list of rules to select which information to show in the table. Click **Show/Hide Empty Columns** to quickly add, or remove, all columns that have no information, that is, you are not using those conditions in any rule. Click **Revert to Default** to undo all of your customizations.
- Analyze rule logic—You can select the following options from the **Analyze** menu to examine the logic of your rules:
  - **Manage Rule Hit Counts** —To view statistics on how many connections matched each rule. See View rule hit counts, on page 17 .
  - **Enable/Disable Rule Conflicts** —Select whether you want to see information on whether rules interfere with each other. You can then view the results using the following commands. See Analyzing rule conflicts and warnings, on page 19
    - **Show Warnings and Errors** —See whether there are rules with configuration issues that you need to address.
    - **Show Policy Warnings** —See whether there are configuration issues with the policy.
    - **Show Rule Conflicts** —See whether you have redundant or shadowed rules. These conflicts could prevent certain rules from ever being matched by connections, meaning either that you need to fix the match criteria, move the rule, or simply delete the rule.

**Step 4** To change additional settings for the policy, select one of these options from the **More** drop-down arrow at the end of the packet flow line.

- **Advanced Settings** —To set preprocessing, decryption, identity, performance, and other advanced options. See Configure advanced settings for the access control policy .
- **HTTP Responses** —To specify what the user sees in a browser when the system blocks a website request. See Choose HTTP response pages .
- **Inheritance Settings** —To change the base access control policy for this policy, and to enforce this policy's settings in its descendant policies. See Choose a base access control policy, on page 12 and Lock settings in descendant access control policies, on page 13 .
- **Logging** —To set the default logging options for the policy.

**Step 5** Click **Save** .

# Lock an access control policy

Lock an access control policy to prevent other administrators from editing it, ensuring that your changes will not be invalidated if another administrator edits the policy and saves changes before you save your changes.

Without locking, if multiple administrators edit the policy simultaneously, the first person who saves changes wins, and all other users have their changes erased.

The lock is for the access control policy itself. The lock does not apply to objects used in the policy. For example, another user can edit a network object that is used in a locked access control policy. Your lock remains in place until you explicitly unlock the policy, so you can log out and come back to your edits later.

When locked, other administrators have read-only access to the policy. However, other administrators can assign a locked policy to a managed device.

**Before you begin**

Any user role that has permission to modify the access control policy has permission to lock it, and to unlock a policy that was locked by another user.

However, the ability to unlock a policy that was locked by another administrator is controlled by the following permission: **Policies** > **Access Control heading** > **Access Control** and choose the **Access Control Policy** > **Modify Access Control Policy** > **Override Access Control Policy Lock**.

If you are using custom roles, your organization might have limited your unlocking abilities by not assigning this permission. Without this permission, only the administrator who locks a policy can unlock it.

**Procedure**

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

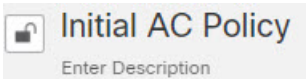
**Step 2** Click **Edit** ( ) next to the access control policy you want to lock or unlock.

In the access control policy list:

- The **Lock Status** column shows whether a policy is already locked, and if so, who locked it. An empty cell indicates that the policy is not locked.

- If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration. Or, it is locked by another user.

**Step 3** Click the lock icon next to the policy name to lock or unlock the policy.




If the policy inherits settings from a parent policy, you must choose one of these options when you click the lock icon.

- **Lock/Unlock this policy**—The locking or unlocking is for this policy only.
- **Lock/Unlock all policies in hierarchy**—This policy and all parent policies are locked or unlocked. If a parent policy is already locked by another administrator, you will see a message and you will not be able to lock that parent policy. When unlocking policies, if you have the Override Access Control Policy Lock permission, all parent policies are unlocked even if they were locked by other users.

# Configure the access control default action

Configure the default action for an access control policy when connections do not match specific access control rules.

The default action for an access control policy is applied to any connection that:

- is not fast pathed by the prefilter policy
- is not on a Security Intelligence block list
- is not blocked by the decryption policy (encrypted traffic only)
- matches none of the rules in the policy (except monitor rules, which match and log—but do not handle or inspect—traffic)

## Procedure

**Step 1** Edit the access control policy whose inheritance settings you want to change; see Editing an access control policy, on page 8.

**Step 2** Select the **Default Action** at the bottom of the rules list.

For detailed information on what each option does, see Access control policy default actions, on page 2.

**Step 3** Click **Cog** ( ) to configure the default action settings.

You can configure the following options. Click **OK** when finished.

- Logging options—Whether to log the connection. You can **Log at beginning of connection**, **Log at end of connection**, or both. If you select block as the default action, you can log at the beginning of the connection only.

- **Send connection events to**—If you select one of the logging options, select whether to send events to any combination of the following:
  - **Firewall Management Center**—Send events to the manager.
  - **Syslog server**—Send events to the default syslog server configured for the policy. You can configure overrides to specify a different severity level or syslog server destination.
- **SNMP trap**—If you enable logging, you can send SNMP traps to an SNMP server. Select an SNMP configuration, or click + to configure a new one.
- **Default action variable set**—If you selected one of the intrusion prevention default actions, select the variable set that should be used with the intrusion policy you selected.

**Step 4** Click **Save**.

# Manage access control policy inheritance

Inheritance relates to using another policy as a base policy for an access control policy. This allows you to use one policy to define some baseline characteristics that can be applied to multiple policies. To understand how inheritance works, refer to Access control policy inheritance, on page 4.

**Procedure**

**Step 1** Edit the access control policy whose inheritance settings you want to change; see Editing an access control policy, on page 8.

**Step 2** Manage policy inheritance:

- Change base policy — To change the base access control policy for this policy, select **Inheritance Settings** from the **More** drop-down arrow at the end of the packet flow line and proceed as described in Choose a base access control policy, on page 12.
- Lock settings in descendants — To enforce this policy's settings in its descendant policies, select **Inheritance Settings** from the **More** drop-down arrow at the end of the packet flow line and proceed as described in Lock settings in descendant access control policies, on page 13.
- Inherit settings from base policy — To inherit settings from a base access control policy, proceed as directed in Inherit access control policy settings from the base policy, on page 13.
- Required in domains — To enforce this policy in a subdomain, click the **Targeted: x devices** link and proceed as described in Enforce an access control policy in a domain, on page 14.

## Choose a base access control policy

Choose a base access control policy to establish policy inheritance, allowing a child policy to inherit rules and settings from its parent while maintaining the ability to modify unlocked configurations.

When you change the base policy for the current access control policy, the system updates the current policy with any locked settings from the new base policy.

### Procedure

**Step 1** In the access control policy editor, select **Inheritance Settings** from the **More** drop-down arrow at the end of the packet flow line.

**Step 2** Choose a policy from the **Select Base Policy** drop-down list.

To remove inheritance, select **None**.

**Step 3** Click **OK**.

**Step 4** Click **Save** to save the access control policy.

### What to do next

- Deploy configuration changes; see Deploy configuration changes.

## Lock settings in descendant access control policies

Lock a setting in an access control policy to enforce the setting in all descendant policies.

Descendant policies can override unlocked settings. When you lock settings, the system saves overrides already made in descendant polices so that the overrides can be restored if you unlock settings again.

### Procedure

**Step 1** In the access control policy editor, select **Inheritance Settings** from the **More** drop-down arrow at the end of the packet flow line.

**Step 2** In the Child Policy Inheritance Settings area, check the settings you want to lock.

If the controls are dimmed, settings are inherited from an ancestor policy, or you do not have permission to modify the configuration.

**Step 3** Click **OK** to save the inheritance settings.

**Step 4** Click **Save** to save the access control policy.

### What to do next

- Deploy configuration changes; see Deploy configuration changes.

## Inherit access control policy settings from the base policy

Configure a child access control policy to inherit settings from a base policy or override them with custom configurations.

Child access control policies inherit many settings from their base policy by default. If these settings are unlocked in the base policy, you can override them.

If you later reinherit the settings from the base policy, the Firewall Management Center displays the base policy's settings and dims the controls. However, it saves the overrides you made and restores them if you disable inheritance again.

### Procedure

**Step 1** In the access control policy editor, click **Security Intelligence**, or select **HTTP Responses**, **Logging**, **Encrypted Visibility Engine**, or **Advanced Settings** from the **More** drop-down arrow at the end of the packet flow line

**Step 2** Check the **Inherit from (*base policy*)** check box for each setting you want to inherit.

If the controls are dimmed, settings are inherited from an ancestor policy, or you do not have permission to modify the configuration.

**Step 3** Click **Save**.

### What to do next

- Deploy configuration changes; see Deploy configuration changes.

## Enforce an access control policy in a domain

Require that every device in a domain use the same base access control policy or one of its descendant policies.

This procedure is relevant in a multi-domain deployment only.

### Procedure

**Step 1** In the access control policy editor, click the **Targeted: x devices** link.

**Step 2** Click **Required on Domains**.

**Step 3** Build your domain list:

- Add — Select the domains where you want to enforce the current access control policy, then click **Add** or drag and drop into the list of selected domains.
- Delete — Click **Delete** ( ) next to a leaf domain, or right-click an ancestor domain and choose **Delete Selected**.
- Search — Type a search string in the search field. Click **Clear** ( ) to clear the search.

**Step 4** Click **OK** to save the domain enforcement settings.

**Step 5** Click **Save** to save the access control policy.

### What to do next

- Deploy configuration changes; see Deploy configuration changes.

# Assign access control policy to devices

An access control policy specifies the devices that use it. Each device can be assigned to only one access control policy.

### Procedure

**Step 1** In the access control policy editor, click the **Targeted: x devices** link.

**Step 2** Build your target list:

- Add — Select one or more **Available Devices**, then click **Add to Policy** or drag and drop into the list of **Selected Devices**.
- Delete — Click **Delete** ( ) next to a single device, or select multiple devices, right-click, then choose **Delete Selection**.
- Search — Type a search string in the search field. Click **Clear** ( ) to clear the search.

Under **Impacted Devices**, the system lists the devices whose assigned access control policies are children of the current policy. Any change to the current policy affects these devices.

**Step 3** Click **OK** to save your assigned device settings.

**Step 4** Click **Save** to save the access control policy.

### What to do next

- Deploy configuration changes; see Deploy configuration changes.

# Logging settings for access control policies

Logging settings for access control policies are configuration option that allow you to configure default syslog destinations for the current access control policy. The settings are applicable to the access control policy and all the included decryption, prefilter, and intrusion policies unless the syslog destination settings are explicitly overridden with custom settings in included rules and policies.

To configure logging settings for an access control policy, select **Logging** from the **More** drop-down arrow at the end of the packet flow line.

You can configure default syslog destinations and syslog ALERT for the access control policy. The settings are applicable to the access control policy and all the included decryption, prefilter, and intrusion policies unless the syslog destination settings are explicitly overridden with custom settings in included rules and policies.

Logging for connections handled by the default action is initially disabled.

Intrusion settings and file and malware settings are effective only after you have selected an option at the top of the page for sending syslog messages generally.

### Default syslog settings

Default syslog settings provide two destination options:

- **Send using specific syslog ALERT**: Events are sent based on the selected syslog ALERT as configured using the instructions in *Creating a Syslog ALERT Response* in the *Cisco Secure Firewall Management Center Administration Guide*. You can select the syslog ALERT from the list or add one by specifying the name, logging host, port, facility, and severity.

  When using this option, the system sends syslog messages to the server using the Management interface. Ensure there is a route from the Management interface to the syslog server, or messages will not arrive at the server.

- **Use the syslog settings configured in the Threat Defense Platform Settings policy deployed on the device**: Connection or intrusion events are sent with the selected severity to syslog collectors configured in Platform Settings. Using this option, you can unify the syslog configuration by configuring it in platform settings and reusing the settings in access control policy. Severity selected in this section is applied to all connection and intrusion events. The default severity is ALERT.

### Intrusion settings

Intrusion settings include these options:

- **Send syslog messages for intrusion events**: Send intrusion events as syslog messages. The defaults set above are used unless you override them.
- **Show overrides** or **Hide overrides**: If you want to use the default syslog destination and severity, leave these options empty. Otherwise, you can set a different syslog server destination for intrusion events, and change the severity of the events.

### File and malware settings

File and malware settings include these options:

- **Send syslog messages for file and malware events**: Send file and malware events as syslog messages. The defaults set above are used unless you override them.
- **Show overrides** or **Hide overrides**: If you want to use the default syslog destination and severity, leave these options empty. Otherwise, you can set a different syslog server destination for file and malware events, and change the severity of the events.

### AI defense settings

# Associate prefilter, decryption, and identity policies with an access control

Associate subpolicies with an access control policy to create a comprehensive security framework that handles different aspects of network traffic processing.

The easiest way to associate subpolicies to an access control policy is by clicking the policy's link in the packet flow shown at the top of an access control policy. Alternatively, you can use the policy's advanced settings to associate the subpolicy, as described in this topic. The subpolicies that you can associate with access control policy are:

- Prefilter policy—Performs early traffic handling using limited network (layer 4) outer-header criteria.
- Decryption policy—Monitors, decrypts, blocks, or allows application layer protocol traffic encrypted with Secure Socket Layer (SSL) or Transport Layer Security (TLS).

- Identity policy—Performs user identification based on the realm and authentication method associated with the traffic.

### Before you begin

Before associating a decryption policy with an access control policy, review the information about TLS server identity discovery in TLS server identity discovery.

### Procedure

**Step 1** In the access control policy editor, select **Advanced Settings** from the **More** drop-down arrow at the end of the packet flow line.

**Step 2** Locate the policy settings area for prefilter, decryption, or identity policies.

**Step 3** Click **Edit** ( ) in the corresponding policy settings area.

If **View** ( ) appears instead, settings are inherited from an ancestor policy, or you do not have permission to modify the settings. If the configuration is unlocked, uncheck **Inherit from base policy** to enable editing.

**Step 4** Choose a policy from the drop-down list.

If you choose a user-created policy, you can click the edit icon that appears to edit the policy.

**Step 5** Click **OK**.

**Step 6** Click **Save** to save the access control policy.

### What to do next

- Deploy configuration changes; see Deploy configuration changes.

# View rule hit counts

Hit count indicates the number of times a policy rule or default action has been matched to a connection. The hit count is incremented only for the first packet of a connection that matches a rule. You can use this information to identify the efficacy of your rules. Hit count information is available only for access control and prefilter rules applied to Firewall Threat Defense devices.

**Note**

- The count persists through reboots and upgrades.
- Counts are maintained by each unit in an HA pair or cluster separately.
- You will not be able to derive the hit count information from a device when deployment or a task is in progress on the device.
- You can also see rule hit count information in the device CLI using the **show rule hits** command.
- If you have accessed the Hit Count page from the Access Control Policy page, you will not be able to view or edit prefilter rules and vice-versa.
- Hit counts are not available for rules that use the Monitor action.

**Before you begin**

If you use custom user roles, ensure that the roles include these privileges:

- View Device, to see the hit counts.
- Modify Device, to refresh the hit counts.

Follow these steps to view rule hit counts:

## Procedure

**Step 1** In the access control policy or prefilter policy editor, click **Analyze** > **Rule Hit Counts** on the top-right of the page.

**Step 2** On the Hit Count page, select the device from the **Select a device** drop-down list.

If it is not the first time that you are generating hit counts for this device, the last fetched hit count information appears next to the drop-down box. Also, verify the **Last Deployed** time to confirm recent policy changes.

**Step 3** If necessary, click **Refresh** ( ) to obtain current hit count data from the selected device.

In the prefilter policy, you might need to click **Fetch Current Hit Count** to get initial hit count data.

You cannot refresh the hit count while deployment to the device is in progress.

**Step 4** View and analyze the data.

You can do these actions:

- Click **Prefilter** or **Access Control** to switch between the hit counts for these policies.
- Search for a specific rule by entering a search string in the **Filter** box.
- Broadly limit the list to **Hit Rules** or **Never Hit Rules** by selecting these options in the **Filter by** field. When viewing hit rules, you can further limit the list by selecting a time range in the **In Last** field (for example, in the last 1 day).
- (When viewed from the access control policy.) You can do these actions with individual rules:

- Edit the rule by clicking **Edit** ( ).
- Delete the rule from the policy by clicking **Delete** ( ).
- Enable or disable the rule by clicking the **Slider** ( ).
- Clear the hit count (reset it to zero) for the rule by clicking the **X** for the rule. You cannot undo this action.

- Change the displayed columns by clicking and selecting the columns to show. When viewed from the prefilter policy, use the **Cog** ( ) to select the columns to display.
- (When viewed from the prefilter policy.) Click on a rule name to edit it, or click **View** ( ) in the last column to view the rule details. Clicking on the rule name highlights it in the policy page where you can edit it.
- (When viewed from the prefilter policy.) Clear the hit count information (reset it to zero) for a rule by right-clicking the rule and selecting **Clear Hit Count**. You can select multiple rules by using Ctrl+click. You cannot undo this action.

**Step 5** Optionally, you can generate a comma-separated values report of the details on the page by clicking **Generate CSV** on the bottom-left of the page.

**Step 6** Click **Close** to return to the policy page.

# Analyzing rule conflicts and warnings

View warnings and information about rule conflicts to examine the logic of your access control policy and to identify rules that need changes.

When rules overlap, you can end up with unnecessary rules in the policy, and these rules will never be matched to traffic. The analysis can help you delete unnecessary rules, or identify rules that should be moved or modified so they enforce the desired policy.

Policy warnings and errors point out things that you should understand and perhaps address to ensure that your rules provide the desired services.

Rule conflict analysis identifies these problems:

- Object overlap—An element in a field of a rule is a subset of one or more elements in the same field of the rule. For example, the source field might include a network object for 10.1.1.0/24, and another object for the host 10.1.1.1. Because 10.1.1.1 is within the network covered by 10.1.1.0/24, the object for 10.1.1.1 is redundant and can be deleted, simplifying the rule and saving device memory.
- Redundant rule—Two rules apply the same action to the same type of traffic and removing the base rule would not change the ultimate result. For example, if a rule permitting FTP traffic for a particular network were followed by a rule allowing IP traffic for that same network, and there were no rules in between denying access, then the first rule is redundant and you can delete it.
- Shadowed rule—This is the reverse of a redundant rule. In this case, one rule will match the same traffic as another rule such that the second rule will never be applied to any traffic because it comes later in the access list. If the action for both rules is the same, you can delete the shadowed rule. If the two rules

specify different actions for traffic, you might need to move the shadowed rule or edit one of the two rules to implement your desired policy. For example, the base rule might deny IP traffic, and the shadowed rule might permit FTP traffic, for a given source or destination.

### Before you begin

When doing the analysis:

- Only the first conflict for a given rule is identified. Once you fix the problem, the rule might be identified as having a conflict with another rule in the table. However, a rule might have multiple warnings or errors.
- Rule conflict analysis considers source/destination security zone, network, VLAN, and service/port match conditions and action only. It does not consider other match criteria, so an apparently redundant rule might not be completely redundant.
- FQDN network objects cannot be analyzed for conflict, because the IP address of an FQDN cannot be known prior to DNS lookup.
- Disabled rules are ignored.
- Time range attributes are ignored. Rules for different time periods might be marked as redundant when they are not actually redundant for the time ranges.
- Icons for warnings and errors, and rule conflicts when you enable the feature, are shown in the rules table. For a reference to the icons, see Rule and other policy warnings, on page 5.

### Procedure

**Step 1** Choose **Policies** > **Access Control heading** > **Access Control** and edit an access control policy.

**Step 2** Do one of these to open the rule conflicts and warnings dialog box:

- To view rule conflicts, click the **Analyze** drop-down and click **Enable Rule Conflicts**. Then, click **Show Rule Conflicts** from the same menu to see the specific results.
- To view rule warnings and errors, click **Analyze** > **Show Rule Warnings and Errors**.
- If you are finished viewing rule conflicts, click **Analyze** > **Disable Rule Conflicts**.

**Step 3** In the rule conflicts and warnings dialog box:

- Rule warnings and errors are shown on a separate tab from rule conflicts.
- Each tab contains sub-tabs to let you examine individual types of problems, such as redundant vs. shadowed or warnings vs. errors. You can also search for an item.
- **More** (⋮) next to each rule name provides shortcuts to edit, disable, or delete the rule.

**Step 4** Click **Close** when finished.

# Search for rules

Use search to help you find rules, especially when you have a large number of them.

When you search for an IP address in the source or destination network (but not as a simple text search), the system returns rules that will match the address. This includes not only exact matches, but also subnet matches. For example, searching for 10.1.1.1 will include rules for 10.1.1.0/24.

**Procedure**

**Step 1** When editing an access control policy, build the search string by clicking in the **Search** box.

- For a simple text string search, type the string. The search returns rules that have that string in any column.
- To search on a specific column, start typing the name of the column until the system prompts you with the full name, such as Source Networks. When you select the search tag, you can then enter the search string for that tag. For example, **Source Networks 10.1.1.1**.
- Searches in the port field return exact match only.
- When entering mulitple values, separate the values with commas.
- When building a search string with multiple tags, do not include spaces between the tags.
- When you select a tag, you are prompted with values that appear in those columns. Select the values you want to search for.
- You can quickly filter based on some common features by clicking the filter icon to the left of the search box and selecting to show rules with any combination of the following: Allow, Block, Monitor, Intrusion Policy, Time Range.

**Step 2** With your cursor at the end of the search string in the search box, press Enter.

The rules that satisfy the search string are highlighted and non-matching rules are hidden. You can deselect **Show Only Matching Rules** to see the entire table, with the highlighted rules within the table. This lets you see the surrounding rules.

Beside the Show Only Matching Rules checkbox is a summary of the total number of rules in the policy compared to the number that match the search string.

**Step 3** After your first search, clicking in the search box prompts you with recent searches and tags. You can quickly repeat a search by selecting it, or build similar searches by selecting previous searches or tags and building on them.

**Step 4** To close the search and return to the unfiltered and unhighlighted table, click the **X** at the right of the search box. You can also put your cursor at the end of the search string and press the Esc key.

# Filter rules by device

Some policy editors allow you to filter your rule view by affected devices.

The Firewall Management Center uses a rule's interface constraints to determine if the rule affects a device. If you constrain a rule by interface (security zone or interface group condition), the device where that interface

is located is affected by that rule. Rules with no interface constraint apply to any interface, and therefore every device.

QoS rules are always constrained by interface.

**Procedure**

**Step 1** In the access control policy editor, click **Rules**, then click **Filter by Device**.
A list of targeted devices and device groups appears.

**Step 2** Check one or more check boxes to display only the rules that apply to those devices or groups. Alternatively, check **All** to reset and display all of the rules.

**Note**
Move your pointer over a rule criterion to see its value. If the criterion represents an object with device-specific overrides, the Firewall Management Center displays the override value when you filter the rules list by only that device. If the criterion represents an object with domain-specific overrides, the Firewall Management Center displays the override value when you filter the rules list by devices in that domain.

**Step 3** Click **OK**.

# History for access control policies

This history table provides a comprehensive history of feature changes for access control policies in Cisco Secure Firewall Management Center across releases.

***Table 3: Access control policy feature history by release***

| Feature | Minimum Firewall Management Center | Minimum Firewall Threat Defense | Details |
|---|---|---|---|
| New access control policy user interface and rule conflict analysis. | 7.3.0 | Any | The access control policy user interface introduced in 7.2 is now the default interface. You can also enable rule conflict analysis to help identify redundant rules and objects, and shadowed rules that cannot be matched due to previous rules in the policy. |
| Access control policy locking. | 7.2.0 | Any | You can lock an access control policy to prevent other administrators from editing it. Locking the policy ensures that your changes will not be invalidated if another administrator edits the policy and saves changes before you save your changes. Any user who has permission to modify the access control policy has permission to lock it.<br>We added an icon to lock or unlock a policy next to the policy name while editing the policy. In addition, there is a new permission to allow users to unlock policies locked by other administrators: Override Access Control Policy Lock. This permission is enabled by default in the Administrator, Access Admin, and Network Admin roles. |

| Feature | Minimum Firewall Management Center | Minimum Firewall Threat Defense | Details |
|---|---|---|---|
| Rule hit counts persist over reboot. | 7.2.0 | Any | Rebooting a managed device no longer resets access control rule hit counts to zero. Hit counts are reset only if you actively clear the counters. In addition, counts are maintained by each unit in an HA pair or cluster separately. You can use the **show rule hits** command to see cumulative counters across the HA pair or cluster, or see the counts per node.<br>We modified the device CLI command: **show rule hits**. |
| Usability improvements for the access control policy. | 7.2.0 | Any | There is a new user interface available for the access control policy. You can continue to use the legacy user interface, or you can try out the new user interface. The new interface has both a table and a grid view for the rules list, the ability to show or hide columns, enhanced search, infinite scroll, a clearer view of the packet flow related to policies associated with the access control policy, and a simplified add/edit dialog box for creating rules. You can freely switch back and forth between the legacy and new user interfaces while editing an access control policy. |
| DNS filtering | 7.0.0<br>6.7.0 (experimental) | Any | If URL filtering is enabled and configured, a new option to enhance category and reputation filtering efficacy is enabled by default for each new access control policy.<br>For more information, see DNS filtering and subtopics.<br>The Advanced tab of access control policy has a new option under General Settings: **Enable reputation enforcement on DNS traffic**. |
| TLS server identity discovery | 6.7.0 | Any | Enable access control policies to evaluate URL and application conditions when a client connects to a TLS 1.3-enabled server. TLS server identity discovery enables these conditions to be evaluated without decrypting traffic.<br>Enabling this feature can impact device performance, depending on model.<br>The Advanced tab page of access control policy has new options:<br>• Warning is displayed on the Advanced tab; moving the slider to the right enables TLS server identity discovery.<br>• New option on the Advanced tab page: **TLS Server Identity Discovery**. |

| Feature | Minimum Firewall Management Center | Minimum Firewall Threat Defense | Details |
|---|---|---|---|
| New Security Intelligence categories | — | Any | These categories were introduced at about the time of the 6.6 release, but are not specific to 6.6:<br>• banking_fraud<br>• high_risk<br>• ioc<br>• link_sharing<br>• malicious<br>• newly_seen<br>• spyware |