



Access Control Policies

The following topics describe how to work with access control policies:

- [Access Control Policy Components, on page 1](#)
- [System-Created Access Control Policies, on page 2](#)
- [Requirements and Prerequisites for Access Control Policies, on page 2](#)
- [Managing Access Control Policies, on page 3](#)
- [History for Access Control Policies, on page 25](#)

Access Control Policy Components

Following are the main elements 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

Each access control policy identifies the devices that use it. Each device can be targeted by only one access control policy. You can also assign the policy to device templates.

Rules

Access control rules provide a granular method of handling network traffic. Rules in an access control policy are numbered, starting at 1, including rules inherited from ancestor policies. The system matches traffic to access control rules in top-down order by ascending rule number.

Usually, the system handles network traffic according to the *first* access control rule where *all* the rule's conditions match the traffic. Conditions can be simple or complex, and their use often depends on certain licenses.

Default Action

The default action 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.

Although an access control policy can inherit its default action from an ancestor policy, you cannot enforce this inheritance.

Security Intelligence

Security Intelligence is a first line of defense against malicious internet content. This feature allows you to block connections based on the latest 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 the system blocks a user's website request, you can either display a generic system-provided response page, or a custom page. You can also display a page that warns users, but also allows them to continue to the originally requested site.

Logging

Settings for access control policy logging 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 SSL, prefilter, and intrusion policies unless the syslog destination settings are explicitly overridden with custom settings in included rules and policies.

Advanced Access Control Options

Advanced access control policy settings typically require little or no modification. Often, the default settings are appropriate. Advanced settings you can modify include traffic preprocessing, SSL inspection, identity, and various performance options.

System-Created Access Control Policies

Depending on your devices' initial configurations, system-provided policies can include:

- Default Access Control—Blocks all traffic without further inspection.
- Default Intrusion Prevention—Allows all traffic, but also inspects with the Balanced Security and Connectivity intrusion policy and default intrusion variable set.
- Default Network Discovery—Allows all traffic while inspecting it for discovery data but not intrusions or exploits.

Requirements and Prerequisites for Access Control Policies

Model Support

Any

Supported Domains

Any

User Roles

- Admin
- Access Admin
- Network Admin
- You can define custom user roles to differentiate between the intrusion configuration in access control policy and rules and the rest of the access control policy and rules. Using these permissions, you can separate the responsibilities of your network administration team and your intrusion administration teams. The existing pre-defined user roles that included the Modify Access Control Policy permission support all sub-permissions; you need to create your own custom roles if you want to apply granular permissions. The granular permissions are:
 - **Policies > Access Control > Access Control Policy > Modify Access Control Policy > Modify Threat Configuration** allows the selection of intrusion policy, variable set, and file policy in a rule, the configuration of the advanced options for Network Analysis and Intrusion Policies, the configuration of the Security Intelligence policy for the access control policy, and intrusion actions in the policy default action. If a user has this option only, the user can modify no other part of the policy or rule.
 - **Modify Remaining Access Control Policy Configuration** controls the ability to edit all other aspects of the policy.

Managing Access Control Policies

You can edit system-provided access control policies and create 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 Manage access control policies:

- **Analyze Policy**—Select one or more policy and then click **Analyze Policies** to evaluate access control policies for anomalies such as redundant or shadowed rules, and take action to fix discovered anomalies. The analysis job is sent to the cloud and takes time to complete. See [Identifying and Fixing Anomalies with Policy Analyzer & Optimizer, on page 19](#).

The Anomalies column shows the results of the analysis. Click the **% Optimizable** link to see the anomalies, or **Re-Analyze** to run the analysis again. The Last Analyzed column shows when Policy Analyzer & Optimizer was last run.

After completing analysis and optimization, you can download the reports by selecting the following options from the **More** (⋮) menu: **Download Last Policy Analysis**, **Remediation History**.

Note

To use the Policy Analysis feature, you must be using a Cloud-Delivered Firewall Management Center or an connected to Cisco Security Cloud Control (Security Cloud Control). If your setup does not meet requirements, the explanatory dialog that opens when you click this button includes an **Integrate** button to help get you started. Policy Analyzer & Optimizer operates in the cloud only.

- **Create**—Click **New Policy**; see [Creating a Basic Access Control Policy, on page 4](#).
- **Columns**—Click the **Show/Hide Columns** icon above the list of rules to select which information to show in the table. Click **Show All/Hide All** to quickly add or remove all listed columns, excepting name and actions. Click **Default** to undo all of your customizations.
- **Inheritance**—Click **Plus** next to a policy with descendants to expand your view of the policy's hierarchy.
- **Edit**—Click **Edit** (✎); see [Editing an Access Control Policy, on page 5](#)
- **Delete**—Click **Delete** (🗑). You must remove any device assignments before deleting a policy.
To delete more than one policy at a time, select the check boxes for the policies, then select **Delete Policies** above the table.
- **Copy**—Select **Clone** from the **More** (⋮) menu. Device assignments are not retained in the copy.
- **Report**—Select **Generate Report** from the **More** (⋮) menu.. 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.
- **View the Audit Log**—Click **Go to Audit Log** from the **More** (⋮).
- **Lock or unlock a policy**—See [Locking an Access Control Policy, on page 7](#).

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.

Tip

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. The list includes both devices and device templates.

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 5](#).

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.

Tip

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 Action** drop-down list next to the search box. Bulk editing is available for enabling and disabling, copying, cloning, moving, deleting, and editing rules, or viewing hit counts or related events. You can also remove object overlaps in the selected rules.

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**—Choose a value from the **Default Action** drop-down list.
- **Default Action Settings**—Click **Cog** (⚙), make your changes, and click **OK**. You can configure settings for logging, the location of an external syslog server or SNMP trap server, and the variable set associated with an intrusion prevention default action.
- **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. You can assign the policy to devices or device templates.
- **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 **Hide Empty Columns** to quickly 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:
 - **Hit Count**—To view statistics on how many connections matched each rule.
 - **Enable/Disable Rule Conflicts**—Select whether you want to see information on whether rules interfere with each other.
 - **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.

- **Show Warnings**—See whether there are rules with configuration issues that you need to address.
- **Additional Settings**—To change additional settings for the policy, select one of the following options from the **More** drop-down arrow at the end of the packet flow line.
 - **Advanced Settings**—To set preprocessing, SSL inspection, identity, performance, and other advanced options. See [Access Control Policy Advanced Settings, on page 12](#).
 - **HTTP Responses**—To specify what the user sees in a browser when the system blocks a website request. See [Choosing 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 [Choosing a Base Access Control Policy, on page 9](#) and [Locking Settings in Descendant Access Control Policies, on page 10](#).
- **Logging**—To set the default logging options for the policy.

Step 4 Click **Save**.

What to do next

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

Locking an Access Control Policy

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. 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 > 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.

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 the following options when you click the lock icon.

- **Lock/Unlock This Policy**—The locking or unlocking is for this policy only.
- **Lock/Unlock This Policy and Parents in the 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.

Managing 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, see [Access Control Policy Inheritance](#).

Procedure

Step 1 Edit the access control policy whose inheritance settings you want to change; see [Editing an Access Control Policy, on page 5](#).

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 [Choosing a Base Access Control Policy, on page 9](#).
- **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 [Locking Settings in Descendant Access Control Policies, on page 10](#).
- **Required in Domains** — To enforce this policy in a subdomain, click the **Targeted: x devices** link and proceed as described in [Requiring an Access Control Policy in a Domain, on page 10](#).

- **Inherit Settings from Base Policy** — To inherit settings from a base access control policy, click **Security Intelligence**, or select **HTTP Responses** or **Advanced Settings** from the drop-down arrow at the end of the packet flow line, and proceed as directed in [Inheriting Access Control Policy Settings from the Base Policy, on page 9](#).

Choosing a Base Access Control Policy

You can use one access control policy as the base (parent) for another. By default, a child policy inherits its settings from its base policy, though you can change unlocked settings.

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. |
| Step 3 | Click Save . |
-

What to do next

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

Inheriting Access Control Policy Settings from the Base Policy

A new child policy inherits many settings from its base policy. If these settings are unlocked in the base policy, you can override them.

If you later reinherit the settings from the base policy, the system displays the base policy's settings and dims the controls. However, the system 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 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](#).

Locking 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 policies 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](#).

Requiring an Access Control Policy in a Domain

You can 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](#).

Assigning Devices to an Access Control Policy

An access control policy specifies the devices that use it. Each device can be assigned to only one access control policy. You can also assign the policy to device templates. Templates are included in the list of available and selected devices.

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 Selected**.
- 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 targeted 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

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 SSL/TLS 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.

IPS 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

- **Send using specific syslog alert**—If you select this option, the 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. For more information, see *Facilities and Severities for Intrusion Syslog Alerts* in the [Cisco Secure Firewall Management Center Administration Guide](#). This option is applicable to all devices.

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**—If you select this option and select the severity, 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.

This option is applicable only to Secure Firewall Threat Defense devices 6.3 and later.

IPS Settings

- **Send Syslog messages for IPS events**—Send IPS events as syslog messages. The defaults set above are used unless you override them.
- **Show/Hide Overrides**—If you want to use the default syslog destination and severity, leaves these options empty. Otherwise, you can set a different syslog server destination for IPS events, and change the severity of the events.

File and Malware Settings

- **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/Hide Overrides**—If you want to use the default syslog destination and severity, leaves these options empty. Otherwise, you can set a different syslog server destination for file and malware events, and change the severity of the events.

Access Control Policy Advanced Settings

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

Advanced access control policy settings typically require little or no modification. The default settings are appropriate for most deployments. Note that many of the advanced preprocessing and performance options in access control policies may be modified by rule updates as described in *Update Intrusion Rules* in the [Cisco Secure Firewall Management Center Administration Guide](#).

If **View** (👁) appears instead, settings are inherited from an ancestor policy, or you do not have permission to modify the settings.


Caution

See [Configurations that Restart the Snort Process When Deployed or Activated](#) for a list of advanced setting modifications that restart the Snort process, temporarily interrupting traffic inspection. Whether traffic drops during this interruption or passes without further inspection depends on how the assigned device handles traffic. See [Snort Restart Traffic Behavior](#) for more information.

Inheriting Settings from a Parent Policy

If the access control policy has a base policy, you can elect to inherit settings from the base policy. Select **Inherit from base policy** for each setting group where you want to use the parent policy's settings. If inheritance has been configured so that these settings are locked, you cannot configure unique settings for the policy, these settings are read-only.

If you are allowed to configure unique settings for the policy, you must deselect **Inherit from base policy** to make your edits.

General Settings

Option	Description
Maximum URL characters to store in connection events	To customize the number of characters you store for each URL requested by your users. See <i>Limiting Logging of Long URLs</i> in the Cisco Secure Firewall Management Center Administration Guide for more information. To customize the length of time before you re-block a website after a user bypasses an initial block, see Setting the User Bypass Timeout for a Blocked Website.
Allow an Interactive Block to bypass blocking for (seconds)	See Setting the User Bypass Timeout for a Blocked Website .
Retry URL cache miss lookup	The first time the system encounters a URL that does not have a locally stored category and reputation, it looks up that URL in the cloud and adds the result to the local data store, for faster processing of that URL in the future. This setting determines what the system does when it needs to look up a URL's category and reputation in the cloud. By default, this setting is enabled: The system momentarily delays the traffic while it checks the cloud for the URL's reputation and category, and uses the cloud verdict to handle the traffic. If you disable this setting: When the system encounters a URL that is not in its local cache, the traffic is immediately passed and handled according to the rules configured for Uncategorized and reputationless traffic. In passive deployments, the system does not retry the lookup, as it cannot hold packets.

Option	Description
Enable Threat Intelligence Director	Disable this option to stop publishing TID data to your configured devices.
Enable reputation enforcement on DNS traffic	This option is enabled by default, for improved URL filtering performance and efficacy. For details and additional instructions, see DNS Filtering: Identify URL Reputation and Category During DNS Lookup and subtopics.
Inspect traffic during policy apply	To inspect traffic when you deploy configuration changes unless specific configurations require restarting the Snort process, ensure that Inspect traffic during policy apply is set to its default value (enabled). When this option is enabled, resource demands could result in a small number of packets dropping without inspection. Additionally, deploying some configurations restarts the Snort process, which interrupts traffic inspection. Whether traffic drops during this interruption or passes without further inspection depends on how the assigned device handles traffic. See Snort Restart Scenarios for more information.

Associated Policies

Use advanced settings to associate subpolicies (decryption, identity, prefilter) with access control; see [Associating Other Policies with Access Control, on page 18](#).

TLS Server Identity Discovery

The latest version of the Transport Layer Security (TLS) protocol 1.3, defined by [RFC 8446](#), is the preferred protocol for many web servers to provide secure communications. Because the TLS 1.3 protocol encrypts the server's certificate for additional security, and the certificate is needed to match application and URL filtering criteria in access control rules, the Firepower System provides a way to extract the server certificate *without* decrypting the entire packet.

You can enable this feature, referred to as *TLS server identity discovery*, when you configure advanced settings for an access control policy. Certain features are not supported, such as STARTTLS traffic, the HTTP CONNECT method, and in a network where another device is already decrypting traffic.

If you enable this option, we recommend you also enable the decryption policy's advanced TLS adaptive server identity probe option as well. Together, these options enable more efficient decryption of TLS 1.3 traffic. For more information, see [TLS 1.3 Decryption Best Practices](#).

When a new connection starts that will be affected by TLS server identity discovery, the Firewall Threat Defense holds the original ClientHello packet to determine the identity of the server to which it connects before continuing. The Firewall Threat Defense device sends a specialized connection from the Firewall Threat Defense to the server. The server's response includes the server certificate, the specialized connection is terminated, and the original connection is evaluated as required by the access control policy.

TLS server identity discovery prioritizes the certificate's Common Name (CN) over the [Server Name Indication \(SNI\)](#).

To enable TLS server identity discovery, click the **Advanced** tab, click **Edit** (✎) for the setting, and select **Early application detection and URL categorization**.

TLS Server Identity Discovery

☒ Early application detection and URL categorization

We recommend that you enable early application detection and server identity. Since TLS 1.3 certificates are encrypted, for traffic encrypted with TLS to match access rules that use application or URL filtering, the system must decrypt it. The setting decrypts the certificate only; the connection remains encrypted. Enabling this option is sufficient to decrypt TLS 1.3 certificates; you do not need to create a corresponding SSL decryption rule.

[Revert to Defaults](#)[Cancel](#)[OK](#)

We strongly recommend enabling it for any traffic you want to match on application or URL criteria, especially if you want to perform deep inspection of that traffic. A decryption policy is not required because *traffic is not decrypted* in the process of extracting the server certificate.

**Note**

- TLS server identity discovery *cannot* be used with any of the following:
 - STARTTLS traffic
 - The HTTP CONNECT method
 - Traffic that is already being decrypted by another device on the network
- Because the certificate is decrypted, TLS server identity discovery can reduce performance depending on the hardware platform.
- TLS server identity discovery is not supported in inline tap mode or passive mode deployments.
- Enabling TLS server identity discovery is not supported on any Secure Firewall Threat Defense Virtual deployed to AWS. If you have any such managed devices managed by the Secure Firewall Management Center, the connection event **PROBE_FLOW_DROP_BYPASS_PROXY** increments every time the device attempts to extract the server certificate.
- TLS Server Identity Discovery also operates on TLS 1.2 sessions.

Network Analysis and Intrusion Policies

Advanced network analysis and intrusion policy settings allow you to:

- Specify the intrusion policy and associated variable set that are used to inspect packets that must pass before the system can determine exactly how to inspect that traffic.
- Change the access control policy's default network analysis policy, which governs many preprocessing options.
- Use custom network analysis rules and network analysis policies to tailor preprocessing options to specific security zones, networks, and VLANs.

For more information, see [Advanced Access Control Settings for Network Analysis and Intrusion Policies](#).

Threat Defense Service Policy

You can use the Threat Defense Service Policy to apply services to specific traffic classes. For example, you can use a service policy to create a timeout configuration that is specific to a particular TCP application, as opposed to one that applies to all TCP applications. This policy applies to Firewall Threat Defense devices only, and will be ignored for any other device type. The service policy rules are applied after the access control rules. For more information, see [Service Policies](#).

File and Malware Settings

[Tuning File and Malware Inspection Performance and Storage](#) provides information on performance options for file control and malware defense.

Portscan Threat Detection

Portscan detector is a threat detection mechanism designed to help you detect and prevent portscan activity in all types of traffic to protect networks from eventual attacks. Portscan traffic can be detected efficiently in both allowed and denied traffic. For more information, see [Threat Detection](#).

Elephant Flow Settings

Elephant flows are large, long duration, and fast flows that can cause duress for Snort cores. There are two actions that can be applied on elephant flows to reduce system stress, CPU hogging, packet drops, and so on. These actions are:

- **Bypass any or all applications**—This action bypasses flow from Snort inspection.
- **Throttle**—This action applies dynamic rate limit policy (10% reduction) on elephant flows.

For more information, see the Elephant Flow Detection chapter in the [Cisco Secure Firewall Management Center Snort 3 Configuration Guide](#).

Intelligent Application Bypass Settings

Intelligent Application Bypass (IAB) is an expert-level configuration that specifies applications to bypass or test for bypass if traffic exceeds a combination of inspection performance and flow thresholds. For more information, see [Intelligent Application Bypass](#).

Transport/Network Layer Preprocessor Settings

Advanced transport and network preprocessor settings apply globally to all networks, zones, and VLANs where you deploy your access control policy.

- **Ignore the VLAN header when tracking connections** —Different VLAN tags in traffic traveling in different directions for the same connection can affect traffic reassembly and rule processing. For example, traffic for the same connection could be transmitted over VLAN A and be received over VLAN B. Select this option to configure the system to ignore the VLAN header so packets can be correctly processed for your deployment.
- **Maximum Active Responses**—For a TCP connection that triggers a preprocessor/intrusion rule that is configured to provide an active response, the maximum number of active responses per TCP connection. When additional traffic occurs on a connection where an active response has been initiated, and the traffic occurs more than **Minimum Response Seconds** after a previous active response, the system sends another

active response unless the specified maximum has been reached. A setting of 0 disables additional active responses triggered active response rules.

- **Minimum Response Seconds**—Until **Maximum Active Responses** occur, specifies the number of seconds to wait before any additional traffic on a connection where the system has initiated an active response results in a subsequent active response.
- **Session Termination Logging Threshold**—Do not modify this option unless instructed to do so by Cisco Technical Support. This option specifies for the number of bytes that result in a logged message when the session terminates and the specified number was exceeded. Changing the option can affect system performance.

Detection Enhancement Settings

Detection Enhancement Settings determine whether adaptive profiles are used for application detection and intrusion rules in the access control policy. Typically, the system uses the static settings in your network analysis policy to preprocess and analyze traffic. With adaptive profile updates, the system can adapt processing behavior using host information either detected by network discovery or imported from a third party.

To enable Adaptive Profiles in Snort 3, you must select both the **Enable** and **Enable Profile Updates** options.

- **Enable**—You must enable adaptive profiling (its default state) for access control rules to perform application and file control, including malware protection (AMP), and for intrusion rules to use service metadata.
- **Enable Profile Updates**—Profile updates, like the target-based profiles you can configure manually in a network analysis policy, help to defragment IP packets and reassemble streams in the same way as the operating system on the target host. The intrusion rules engine then analyzes the data in the same format as that used by the destination host. Profile updates also compare metadata in an intrusion rule to host information to determine whether a rule should apply for a particular host.
- **Adaptive Profiles – Attribute Update Interval**—When profile updates are enabled, you can control how frequently in minutes network map data is synced from the management center to its managed devices. The system uses the data to determine what profiles should be used when processing traffic. Increasing the value for this option can improve performance in a large network.
- **Adaptive Profiles – Networks**—Optionally, when profile updates are enabled, you can improve performance by constraining profile updates to a comma-separated list of IP addresses, address blocks, and network variables. If you use a network variable, the system uses the variable's value in the variable set linked to the default intrusion policy for your access control policy. For example, you could enter: 192.168.1.101, 192.168.4.0/24, \$HOME_NET. IPv4 and IPv6 are supported.

The default value (0.0.0.0/0) applies adaptive profile updates to all networks.

Performance Settings and Latency-Based Performance Settings

[About Intrusion Prevention Performance Tuning](#) provides information on improving the performance of your system as it analyzes traffic for attempted intrusions.

For information specific to latency-based performance settings, see [Packet and Intrusion Rule Latency Threshold Configuration](#).

Encrypted Visibility Engine

For details about this feature, see the Encrypted Visibility Engine chapter in the [Cisco Secure Firewall Management Center Snort 3 Configuration Guide](#).

Associating Other Policies with Access Control

The easiest way to associate the major policy's to an access control policy is by clicking the policy's link in the packet flow shown at the topic of the access control policy. You can quickly select the associated policy. Alternatively, you can use the policy's advanced settings to associate the policy, as described in this topic. These policies include the following:

- 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).



Caution

Snort 2 only. Adding or removing an SSL policy restarts the Snort process when you deploy configuration changes, temporarily interrupting traffic inspection. Whether traffic drops during this interruption or passes without further inspection depends on how the assigned device handles traffic. See [Snort Restart Traffic Behavior](#) for more information.

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

Before you begin

Before associating an SSL policy with an access control policy, review the information about TLS server identity discovery in [Access Control Policy Advanced Settings, on page 12](#).

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** Click **Edit** (✎) in the appropriate 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 3** Choose a policy from the drop-down list.
If you choose a user-created policy, you can click edit that appears to edit the policy.
- Step 4** Click **OK**.
- Step 5** Click **Save** to save the access control policy.

What to do next

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

Identifying and Fixing Anomalies with Policy Analyzer & Optimizer

You can use the Policy Analyzer & Optimizer to evaluate access control policies for anomalies such as redundant or shadowed rules, and take action to fix discovered anomalies. The Policy Analyzer & Optimizer is hosted in the cloud and is different from the rule analysis available when you are not integrated with the cloud. The non-cloud policy analysis is not available once you integrate with the cloud.

The system automatically performs policy analysis on a daily basis (every 24 hours). You can also manually start an analysis. When you initially enable the service, the system starts an analysis of all existing access control policies.



Note Before optimizing a policy, create a copy of the policy. If you are then dissatisfied with the results of optimization, you can easily reassign the managed devices to the copy and return the system to its starting state.

Before you begin

- You can use Policy Analyzer & Optimizer directly from Security Cloud Control with Firewall Management Center 7.2+, but you can cross-launch the feature starting with 7.6 only. The software version running on the managed devices assigned to the access control policies does not matter. Policy Analyzer & Optimizer is hosted in the cloud only, so whether you run the analysis from within the Firewall Management Center or directly from Security Cloud Control does not make a difference.
- To use the feature, you must select **Enable Policy Analysis & Optimization** when you integrate with the Cisco Security cloud, under **Integration > Cisco Security Cloud**.
- If you have enabled Change Management, Policy Analyzer & Optimizer automatically creates a ticket for the changes, and submits the ticket. The approver must approve the ticket before the changes can be deployed.
- If you are using domains, you cannot click the link in the Anomaly column to see the report. Instead, log into Security Cloud Control and use the feature from that application.
- Policy Analyzer & Optimizer adds rule comments on rules that are updated, disabled, or merged. You can later search on these comments to find optimized rules.
- Changes implemented by Policy Analyzer & Optimizer are reflected in the audit log as API calls under the default name internaladmin.

Procedure

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

If you have already run an analysis, the Anomaly column shows the number of issues with the policy and the percentage the policy can be optimized, and the state of the policy analysis, such as Error or Completed. Last Analyzed shows the date/time when the analysis was run.

Step 2 Select one or more policy, then click **Analyze Policy**.

The analysis runs as a background process in the cloud. When the analysis is complete, the results appear in the Anomaly column.

Notes:

- You can also start an analysis when editing the policy by selecting **Analyze > Policy**. Other options from that menu allow you to show hit counts and warnings.
- If you have not connected to the cloud yet, the explanatory dialog that opens when you click this button includes an **Integrate** button to help get you started. Policy Analyzer & Optimizer operates in the cloud only.

Step 3 When the analysis is complete, click the % **Optimizable** link in the Anomaly column to launch Policy Analyzer & Optimizer in the cloud.

When you have done all the actions you want to take, click **Apply Remediations** (in the cloud). You are shown a confirmation of what will be done. Click **Proceed** to implement the changes.

If the initial analysis ended in an error, you could instead click **Re-analyze** to restart the process.

Step 4 Deploy the policy to complete the changes.

If you have Change Management enabled, the approver must first approved the ticket that contains the remediations before you can deploy them.

Viewing 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 the following privileges:

- View Device, to see the hit counts.
- Modify Device, to refresh the 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 the following:
- 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 **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.) Clear the hit counts for one or more rule by selecting the checkbox for the rule, then clicking **Clear Hit Counts**. When confirming the action, choose **Clear and Reload** to refresh the hit count data. You can clear hit counts for up to 500 rules at a time. You cannot undo clearing hit counts.

Note

Click the checkbox in the table header to select all rules visible in the list. To select a range of rules, select the first rule's checkbox, then Shift+click the last rule's checkbox; all rules in between are also selected.

- (When viewed from the access control policy.) You can do the following with individual rules:
 - Edit the rule by clicking **Edit** from the **More** (⋮) menu.
 - Delete the rule from the policy by clicking **Delete** from the **More** (⋮) menu.
 - Enable or disable the rule by clicking **Enable/Disable Rule** from the **More** (⋮) menu.
 - Clear the hit count (reset it to zero) for the rule by clicking **Clear Hit Count** from the **More** (⋮) menu. You cannot undo this action.

- (When viewed from the prefilter policy.) Change the displayed columns by clicking **Cog** (⚙️) and selecting the columns to show.
- (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.
- Generate a comma-separated values report of the details on the page by clicking **Generate CSV** on the bottom-left of the page.

Step 5 Click **Close** to return to the policy page.

Analyzing Rule Conflicts and Warnings



Note The feature described in this topic is available only if you do not integrate with the Cisco Security Cloud. When you integrate with the cloud, the more powerful Policy Analyzer & Optimizer replaces this feature. See [Identifying and Fixing Anomalies with Policy Analyzer & Optimizer, on page 19](#).

You can 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 the following types of problem:

- **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](#).

Procedure

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

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

- To view rule conflicts, click the **Analyze** drop-down and click **Enable Rule Conflicts**. When the analysis completes, you see a summary of the conflicts at the top of the page. Then, click **Show Rule Conflicts** from the same menu to see the specific results.

You must re-enable rule conflict detection each time you open the policy or make a change and save the policy.

- To view rule warnings and errors, click **Analyze > Show Rule Warnings**.

After you make a change to the policy, you can refresh the results by clicking the reload icon next to the Analyze button.

- To view policy warnings, click **Analyze > Show Policy Warnings**.
- If you are finished viewing rule conflicts, click **Analyze > Disable Rule Conflicts**.

Step 3 In the rule conflicts and warnings dialog box:

- 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.

Searching for Rules

You can 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. You cannot do a string search in conjunction with a tag search, such as combining a string search with a Source Networks search.
- 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, or select the name from the list of searchable fields. When you select the search tag, you can then enter the search string for that tag. For example, **Source Networks 10.1.1.1**.
- When entering multiple values, separate the values with commas.
- 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.
- 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, Conflicts, Warnings, Errors, Disabled rules, expired rules, rules that include objects whose definitions overlap..
- To see the rules that apply to a specific device or set of devices, click the filter icon and select the devices. Rules apply to a device if they use security zones that contain at least one interface on the device, or if they do not include security zones.

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** 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.

History for Access Control Policies

Table 1:

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Object group search performance enhancements.	7.6.0	7.6.0	Object group search is now faster and uses fewer CPU resources. New CLI commands: clear asp table network-object , show asp table network-object , debug acl ogs Modified CLI comments (enhanced output): , packet-tracer , show access-list , show object-group
Policy Analyzer & Optimizer for access control.	From mgmt. center: 7.6.0 From Security Cloud Control: 7.2.0	Any	The Policy Analyzer & Optimizer evaluates access control policies for anomalies such as redundant or shadowed rules, and can take action to fix discovered anomalies. You can launch the access control Policy Analyzer & Optimizer directly from a Version 7.6+ Firewall Management Center; this requires Cisco Security Cloud. For Versions 7.2–7.4 Firewall Management Centers, use Security Cloud Control. New/modified screens: • To enable: Integration > Cisco Security Cloud > Enable Policy Analyzer & Optimizer • To analyze policies: Policies > Access Control, select policies, click Analyze Policies.

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Granular permissions for modifying access control policies and rules.	7.4.0	Any	You can define custom user roles to differentiate between the intrusion configuration in access control policies and rules and the rest of the access control policy and rules. Using these permissions, you can separate the responsibilities of your network administration team and your intrusion administration teams. When defining user roles, you can select the Policies > Access Control > Access Control Policy > Modify Access Control Policy > Modify Threat Configuration option to allow the selection of intrusion policy, variable set, and file policy in a rule, the configuration of the advanced options for Network Analysis and Intrusion Policies, the configuration of the Security Intelligence policy for the access control policy, and intrusion actions in the policy default action. You can use the Modify Remaining Access Control Policy Configuration to control the ability to edit all other aspects of the policy. The existing pre-defined user roles that included the Modify Access Control Policy permission continue to support all sub-permissions; you need to create your own custom roles if you want to apply granular permissions.
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. 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.
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. We modified the following device CLI command: show rule hits.

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
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 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. For more information, see DNS Filtering: Identify URL Reputation and Category During DNS Lookup and subtopics. 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. Enabling this feature can impact device performance, depending on model. The Advanced tab page of access control policy has new options: - Warning is displayed on the Advanced tab; moving the slider to the right enables TLS server identity discovery. - New option on the Advanced tab page: TLS Server Identity Discovery.
New Security Intelligence categories	—	Any	The following categories were introduced at about the time of the 6.6 release, but are not specific to 6.6: - banking_fraud - high_risk - ioc - link_sharing - malicious - newly_seen - spyware

