Access Control Policies

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

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, which is especially useful in multidomain deployments.

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. In a multidomain deployment, you can require that all the devices in a domain use the same base policy.

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

Managing Access Control Policies

You can edit system-provided access control policies and create custom access control policies.

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

Procedure

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

Step 2 Manage access control policies:

- Copy—Click **Copy** (⧉)..
- Create—Click **New Policy**; see Creating a Basic Access Control Policy, on page 3.
- Delete—Click **Delete** (🗑).
- Edit—Click **Edit** (✎); see Editing an Access Control Policy, on page 4
- Lock or unlock a policy—See Locking an Access Control Policy, on page 6.
- Inheritance—Click **Plus** next to a policy with descendants to expand your view of the policy's hierarchy.
- Import/Export—Click **Import/Export**; see *Import/Export* in the Cisco Secure Firewall Management Center Administration Guide.
- Report—Click **Report** (📋); see Generating Current Policy Reports.

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**.

Step 2 Click **New Policy**.

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

Step 4 Optionally, choose a base policy from the **Select Base Policy** drop-down list.

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 **Available Devices** where you want to deploy the policy, then click **Add to Policy** (or drag and drop) to add the selected devices. To narrow the devices that appear, type a search string in the **Search** field.

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 4 .

Related Topics

Access Control Policy Default Action

Setting Target Devices for an Access Control Policy, on page 10

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.

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. If multiple users save the same policy, the last saved changes are retained. For your convenience, the system displays information on who (if anyone) is currently editing each 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.

Procedure

Step 1 Choose **Policies** > **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 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.
- 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. Ssee 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 8 and Locking Settings in Descendant Access Control Policies, on page 9 .
    - **Logging**—To set the default logging options for the policy.

Step 4 Click **Save**.

What to do next

- Deploy configuration changes.

Related Topics

Rule and Other Policy Warnings

Deep Inspection Using File and Intrusion Policies

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**.

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.

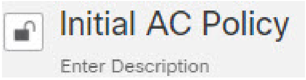

Initial AC Policy

Enter Description

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 4.

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 8.
- 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 9 .
- 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 9.
- 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 8.

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.

In a multidomain deployment, an access control policy might be required in the current domain. You can choose only the enforced policy or one of its descendants as the base policy.

Step 3 Click **Save**.

What to do next

- 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.

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 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.

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.

Setting Target Devices for an Access Control Policy

An access control policy specifies the devices that use it. Each device can be targeted by only one access control policy. In multidomain deployments, you can require that all the devices in a domain use the same base policy.

Procedure

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

Step 2 On **Targeted Devices**, 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 (Multi-domain deployments only.) Optionally, click **Required on Domains** to require that all the devices in the subdomains you choose use the same base policy.

Step 4 Click **OK** to save your targeted device settings.

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

What to do next

- 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.
- **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. If the configuration is unlocked, uncheck **Inherit from base policy** to enable editing.

⚠

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 target device handles traffic. See Snort Restart Traffic Behavior for more information.

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. 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 target 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.

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.

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

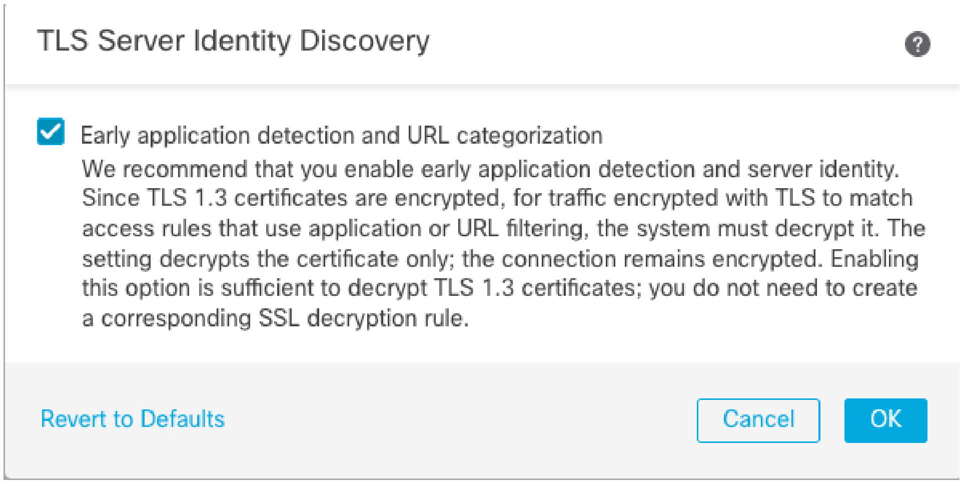

TLS Server Identity Discovery

☑ 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.

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

- 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.

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 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..

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.

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. You configure these advanced settings in an access control policy rather than in a network analysis policy. For more information, see Advanced Transport/Network Preprocessor Settings.

Detection Enhancement Settings

Advanced detection enhancement settings allow you to configure adaptive profiles so you can:

- Use file policies and applications in access control rules.
- Use service metadata in intrusion rules.
- In passive deployments, improve reassembly of packet fragments and TCP streams based on your network’s host operating systems.

For more information, see Adaptive Profiles.

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 Encrypted Visibility Engine.

Encrypted Visibility Engine

The encrypted visibility engine (EVE) is used to provide more visibility into the encrypted sessions without the need to decrypt them. These insights into encrypted sessions are obtained by Cisco's open-source library that is packaged in Cisco's vulnerability database (VDB). The library fingerprints and analyzes incoming encrypted sessions and matches it against a set of known fingerprints. This database of known fingerprints is also available in the Cisco VDB.

Use the **Encrypted Visibility Engine (EVE)** toggle button, available under the **Advanced** tab of the access control policy, to enable or disable EVE. In management center 7.1, encrypted visibility engine is used to only provide more visibility into the encrypted traffic. It does not enforce any actions on that traffic.

In management center 7.2, encrypted visibility engine (EVE) has the following enhanced features:

- To use EVE on management center 7.2, you must have a valid IPS license on your device. In the absence of a IPS license, the policy displays a warning and deployment is not allowed.
- You can take access control policy actions on the traffic using information derived from EVE.
- The VDB included in Cisco Secure Firewall 7.2 has the ability to assign applications to some processes detected by EVE with a high confidence value. Alternatively, you can create custom application detectors to:
  - Map EVE-detected processes to new user-defined applications.
  - Override the built-in value of process confidence that is used to assign applications to EVE-detected processes.

    See Configuring Custom Application Detectors and **Specifying EVE Process Assignments** section in the **Application Detection** chapter of the *Cisco Secure Firewall Device Management Configuration Guide*.
- EVE can detect the operating system type and version of the client that created a Client Hello packet in the encrypted traffic.
- EVE supports fingerprinting and analysis of Quick UDP Internet Connections (QUIC) traffic too. The server name from the Client Hello packet is displayed in the URL field of the **Connection Events** page.

Note The encrypted visibility engine feature is supported only on management center-managed devices running Snort 3. This feature is not supported on Snort 2 devices, device manager-managed devices, or CDO.

After the **Encrypted Visibility Engine** toggle button is enabled and your access control policy is deployed, you can start sending live traffic through your system. You can view the logged connection events in the **Connection Events** page. To access the connection events, in management center, go to **Analysis** > **Connections** > **Events**, and click the **Table View of Connection Events** tab. You can also view the connection event fields in the **Unified Events** viewer, which is under the **Analysis** menu. Encrypted Visibility Engine can identify the client process that initiated a connection, the OS on the client, and if the process contains malware or not.

In the **Connection Events** page, the following columns are added for Encrypted Visibility Engine. Note that you must explicitly enable the mentioned columns.

- Encrypted Visibility Process Name
- Encrypted Visibility Process Confidence Score
- Encrypted Visibility Threat Confidence
- Encrypted Visibility Threat Confidence Score
- Detection Type

For information about these fields, see the section *Connection and Security Intelligence Event Fields* in the Cisco Firepower Management Center Administration Guide.

Note In the **Connection Events** page, if processes are assigned applications, the **Detection Type** column displays **Encrypted Visibility Engine** indicating that the client application was identified by EVE. Without application assignments to process names, the **Detection Type** column displays **AppID** indicating that the engine that identified the client application was AppID.

You can view the analysis information in two dashboards. Under **Overview** > **Dashboards**, click **Dashboard**. In the **Summary Dashboard** window, click the **switch dashboard** link and select **Application Statistics** from the dropdown box. Select the **Encrypted Visibility Engine** tab to view the following two dashboards:

- **Top Encrypted Visibility Engine Discovered Processes**—Displays the top TLS process names being used in your network and the connection count. You can click the process name in the table to see the filtered view of the **Connection Events** page, which is filtered by the process name.
- **Connections by Encrypted Visibility Engine Threat Confidence**—Displays the connections by the confidence levels (Very High, Very Low, and so on). You can click the Threat confidence level in the table to see the filtered view of the **Connection Events** page, which is filtered by the confidence level.

In management center 7.2, EVE can detect the operating system type and version of SSL sessions. Normal usage of the operating system, such as running applications, package management software, and so on, can trigger OS detection. To view client OS detection, in addition to enabling the EVE toggle, you must enable **Hosts** under **Policies** > **Network Discovery**. To view a list of possible Operating Systems on the host IP address, click **Analysis** > **Hosts** > **Network Map**, and then select the required host.

From management center 7.3.0, the host's Indications of Compromise (IoC) events for encrypted visibility engine detection allows you to check connection events with a very high malware confidence level, as reported by EVE. IoC events are triggered for encrypted sessions generated from a host using a malicious client. You can view information, such as IP address, MAC address, and OS information of the malicious host, and the timestamp of the suspicious activity.

A session with Encrypted Visibility Threat Confidence score 'Very High' as seen in connection events genreates an IoC event. You must enable **Hosts** from **Policies** > **Network Discovery**. In management center, you can view the IoC event existence from:

- **Analysis** > **Indications of Compromise**.
- **Analysis** > **Network Map** > **Indications of Compromise** > Choose the host that must be checked.

You can view the process information of the session that generated the IoC from:

Analysis > **Connection Events** > **Table View of Connection Events** > **IoC** column. Note that you must manually select the Encrypted Visbility fields and IoC field.

Snort can identify client applications in QUIC sessions based on EVE. QUIC fingerprinting can:

- Detect applications over QUIC without enabling decryption.
- Identify malware without enabling decryption.
- Detect service applications. You can assign access control rules based on the service detected over QUIC protocol.

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 target 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.

Related Topics

Snort Restart Scenarios

Viewing Policy Hit Counts

Hit count indicates the number of times a policy rule or default action has been matched to a connection. The policy hit count is incremented only for the first packet of a connection that matches a policy. 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 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 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.) You can do the following 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.
- (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

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:

- Redundant Object—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**. Then, click **Show Rule Conflicts** from the same menu to see the specific results.
- To view rule warnings and errors, click **Analyze** > **Show 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, 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. Separate multiple strings with a semi-colon (;).
- 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**.
- 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.

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.