

Rule Management: Common Characteristics

The following topics describe how to manage common characteristics of rules in various policies on the Firepower Management Center:

- Requirements and Prerequisites for Rule Management, on page 1
- Introduction to Rules, on page 1
- Rule Condition Types, on page 3
- Searching for Rules, on page 31
- Filtering Rules by Device, on page 31
- Identify Rules with Issues, on page 32
- Rule and Other Policy Warnings, on page 33

Requirements and Prerequisites for Rule Management

Model Support

Any.

Supported Domains

Any

User Roles

- Admin
- Access Admin
- Network Admin

Introduction to Rules

Rules in various policies exert granular control over network traffic. The system evaluates traffic against rules in the order that you specify, using a first-match algorithm.

Although these rules may include other configurations that are not consistent across policies, they share many basic characteristics and configuration mechanics, including:

- Conditions: Rule conditions specify the traffic that each rule handles. You can configure each rule with multiple conditions. Traffic must match all conditions to match the rule.
- Action: A rule's action determines how the system handles matching traffic. Note that even if a rule does not have an **Action** list you can choose from, the rule still has an associated action. For example, a custom network analysis rule uses a network analysis policy as its "action." As another example, QoS rules do not have an explicit action because all QoS rules do the same thing: rate limit traffic.
- Position: A rule's position determines its evaluation order. When using a policy to evaluate traffic, the system matches traffic to rules in the order you specify. Usually, the system handles traffic according to the first rule where all the rule's conditions match the traffic. (Monitor rules, which are designed to track and log, are an exception.) Proper rule order reduces the resources required to process network traffic, and prevents rule preemption.
- Category: To organize some rule types, you can create custom rule categories in each parent policy.
- Logging: For many rules, logging settings govern whether and how the system logs connections handled by the rule. Some rules (such as identity and network analysis rules) do not include logging settings because the rules neither determine the final disposition of connections, nor are they specifically designed to log connections. As another example, QoS rules do not include logging settings; you cannot log a connection simply because it was rate limited.
- Comments: For some rule types, each time you save changes, you can add comments. For example, you might summarize the overall configuration for the benefit of other users, or note when you change a rule and the reason for the change.



Tip A right-click menu in many policy editors provides shortcuts to many rule management options, including editing, deleting, moving, enabling, and disabling.

Rules with Shared Characteristics

This chapter documents many common aspects of the following rules and configurations. For information on non-shared configurations, see:

- Access control rules: Access Control Rules
- Prefilter rules: Tunnel and Prefilter Rule Components
- Tunnel rules: Tunnel and Prefilter Rule Components
- SSL rules: Creating and Modifying TLS/SSL Rules
- DNS rules: Creating and Editing DNS Rules
- Identity rules: Create an Identity Rule
- Network analysis rules: Configuring Network Analysis Rules
- QoS rules: Configuring QoS Rules
- Intelligent Application Bypass (IAB): Intelligent Application Bypass
- Application filters: Application Filters

Rules without Shared Characteristics

Rules whose configurations are not documented in this chapter include:

- Intrusion rules: Tuning Intrusion Policies Using Rules
- File and malware rules: File Rules
- Correlation rules: Configuring Correlation Rules
- NAT rules (Classic): NAT for 7000 and 8000 Series Devices
- NAT rules (Firepower Threat Defense): Network Address Translation (NAT) for Firepower Threat Defense
- 8000 Series fastpath rules: Configure Fastpath Rules (8000 Series)

Rule Condition Types

The following table describes the common rule conditions documented in this chapter, and lists the configurations where they are used.

Condition	Controls Traffic By...	Supported Rules/Configurations
Interface Conditions, on page 6	Source and destination interfaces, and where supported, tunnel zones	Access control rules Tunnel rules Prefilter rules SSL rules DNS rules Identity rules Network analysis rules QoS rules
Network Conditions, on page 8	Source and destination IP address, and where supported, geographical location or originating client	Access control rules Prefilter rules SSL rules DNS rules Identity rules Network analysis rules QoS rules
Tunnel Endpoint Conditions, on page 10	Source and destination tunnel endpoints for plaintext, passthrough tunnels	Tunnel rules

Condition	Controls Traffic By...	Supported Rules/Configurations
VLAN Conditions, on page 12	VLAN tag	Access control rules **Note** For FTD, VLAN tags in access rules only apply to inline sets; they cannot be used in access rules applied to firewall interfaces. Tunnel rules Prefilter rules SSL rules DNS rules Identity rules Network analysis rules
Port and ICMP Code Conditions, on page 13	Source and destination ports, protocols, and ICMP codes	Access control rules Prefilter rules SSL rules Identity rules QoS rules
Encapsulation Conditions, on page 15	Encapsulation protocol (nonencrypted)	Tunnel rules
Application Conditions (Application Control), on page 15	Application or application characteristic (type, risk, business relevance, category, and tags)	Access control rules SSL rules Identity rules QoS rules Application filters Intelligent Application Bypass (IAB)
URL Conditions (URL Filtering), on page 24	URL, and where supported, URL characteristic (category and reputation)	Access control rules SSL rules QoS rules
User, Realm, and ISE Attribute Conditions (User Control), on page 24	Logged-in authoritative user of a host, or that user's realm, group, or ISE attributes	Access control rules SSL rules (no ISE attributes) QoS rules

Condition	Controls Traffic By...	Supported Rules/Configurations
Custom SGT Conditions, on page 29	Custom Security Group Tag (SGT)	Access control rules QoS rules

Rule Condition Mechanics

Rule conditions specify the traffic that each rule handles. You can configure each rule with multiple conditions, and traffic must match all conditions to match the rule. The available condition types depend on the rule type.

In rule editors, each condition type has its own tab page. Build conditions by choosing the traffic characteristics you want to match. In general, choose criteria from one or two lists of available items on the left, then add or combine those criteria into one or two lists of selected items on the right. For example, in URL conditions in access control rules, you can combine URL category and reputation criteria to create a single group of websites to block.

To help you build conditions, you can match traffic using various system-provided and custom configurations, including realms, ISE attributes, and various types of objects and object groups. Often, you can manually specify rule criteria.

Leave matching criteria empty whenever possible, especially those for security zones, network objects, and port objects. When you specify multiple criteria, the system must match against every combination of the contents of the criteria you specify.



Caution Failure to set up your access control rules properly can have unexpected results, including traffic being allowed that should be blocked. In general, application control rules should be lower in your access control list because it takes longer for those rules to match than rules based on IP address, for example.

Access control rules that use *specific* conditions (such as networks and IP addresses) should be ordered *before* rules that use general conditions (such as applications). If you're familiar with the Open Systems Interconnect (OSI) model, use similar numbering in concept. Rules with conditions for layers 1, 2, and 3 (physical, data link, and network) should be ordered first in your access control rules. Conditions for layers 5, 6, and 7 (session, presentation, and application) should be ordered later in your access control rules. For more information about the OSI model, see this Wikipedia article.

Source and Destination Criteria

Where a rule involves source and destination criteria (zones, networks, ports), usually you can use either or both criteria as constraints. If you use both, matching traffic must originate from one of the specified source zones, networks, or ports and leave through one of the destination zones, networks, or ports.

Items per Condition

You can add up to 50 items to each condition. For rules with source and destination criteria, you can use up to 50 of each. Traffic that matches any of the selected items matches the condition.

Simple Rule Mechanics

In rule editors, you have the following general choices. For detailed instructions on building conditions, see the topics for each condition type.

- Choose Item—Click an item or check its check box. Often you can use Ctrl or Shift to choose multiple items, or right-click to **Select All**.
- Search—Enter criteria in the search field. The list updates as you type. The system searches item names and, for objects and object groups, their values. Click **Reload** ( ) or **Clear** ( ) to clear the search.
- Add Predefined Item—After you choose one or more available items, click an **Add** button or drag and drop. The system prevents you from adding invalid items: duplicates, invalid combinations, and so on.
- Add Manual Item—Click the field under the **Selected** items list, enter a valid value, and click **Add**. When you add ports, you may also choose a protocol from the drop-down list.
- Create Object—Click **Add** ( ) to create a new, reusable object that you can immediately use in the condition you are building, then manage in the object manager. When using this method to add application filters on the fly, you cannot save a filter that includes another user-created filter.
- Delete—Click the **Delete** ( ) for an item, or choose one or more items and right-click to **Delete Selected**.

Interface Conditions

Interface rule conditions control traffic by its source and destination interfaces.

Depending on the rule type and the devices in your deployment, you can use predefined *interface objects* called *security zones* or *interface groups* to build interface conditions. Interface objects segment your network to help you manage and classify traffic flow by grouping interfaces across multiple devices; see Interface Objects: Interface Groups and Security Zones.

Tip Constraining rules by interface is one of the best ways to improve system performance. If a rule excludes all of a device's interfaces, that rule does not affect that device's performance.

Just as all interfaces in an interface object must be of the same type (all inline, passive, switched, routed, or ASA FirePOWER), all interface objects used in an interface condition must be of the same type. Because devices deployed passively do not transmit traffic, in passive deployments you cannot constrain rules by destination interface.

Tunnel Zones vs Security Zones

In some configurations, you can use tunnel zones instead of security zones to constrain interface conditions. Tunnel zones allow you to use prefiltering to tailor subsequent traffic handling to certain types of encapsulated connections.

Note If a configuration supports tunnel zone constraints, a rezoned connection—a connection with an assigned tunnel zone—does **not** match security zone constraints. For more information, see Tunnel Zones and Prefiltering.

Rules with Interface Conditions

Rule Type	Supports Security Zones?	Supports Tunnel Zones?	Supports Interface Groups?
Access control	yes	yes	no
Tunnel and prefilter	yes	n/a; you assign tunnel zones in the prefilter policy	yes
SSL	yes	no	no
DNS (source only)	yes	no	no
Identity	yes	no	no
Network analysis	yes	no	no
QoS (routed only, required)	yes	no	yes

Example: Access Control Using Security Zones

Consider a deployment where you want hosts to have unrestricted access to the internet, but you nevertheless want to protect them by inspecting incoming traffic for intrusions and malware.

First, create two security zones: Internal and External. Then, assign interface pairs on one or more devices to those zones, with one interface in each pair in the Internal zone and one in the External zone. Hosts connected to the network on the Internal side represent your protected assets.

Note You are not required to group all internal (or external) interfaces into a single zone. Choose the grouping that makes sense for your deployment and security policies.

Then, configure an access control rule with a destination zone condition set to Internal. This simple rule matches traffic that leaves the device from any interface in the Internal zone. To inspect matching traffic for intrusions and malware, choose a rule action of **Allow**, then associate the rule with an intrusion and a file policy.

Configuring Interface Conditions

Before you begin

- (Access control only) If you want to constrain traffic by tunnel zones instead of security zones, make sure the associated prefilter policy assigns those zones; see Associating Other Policies with Access Control.

Procedure

Step 1 In the rule editor, click the following for interface conditions:

- Interface groups and security zones (tunnel, prefilter, QoS)—Click **Interface Objects**.
- Security zones (access control, SSL, DNS, identity, network analysis)—Click **Zones**.
- Tunnel zones (access control)—Click **Zones**.

Step 2 Find and choose the interfaces you want to add from the **Available Interface Objects** or **Available Zones** list.

(Access control only) To match connections in rezoned tunnels, choose tunnel zones instead of security zones. You cannot use tunnel and security zones in the same rule. For more information, see Tunnel Zones and Prefiltering.

Step 3 Click **Add to Source** or **Add to Destination**, or drag and drop.

Step 4 Save or continue editing the rule.

What to do next

- (Access control only) If you rezoned tunnels during prefiltering, configure additional rules if necessary to ensure complete coverage. Connections in rezoned tunnels do **not** match rules with security zone constraints. For more information, see Using Tunnel Zones.
- Deploy configuration changes; see Deploy Configuration Changes.

Network Conditions

Network rule conditions control traffic by its source and destination IP address, using inner headers. Tunnel rules, which use outer headers, have tunnel endpoint conditions instead of network conditions.

You can use predefined objects to build network conditions, or manually specify individual IP addresses or address blocks.

Note The system builds a separate network map for each leaf domain. In a multidomain deployment, using literal IP addresses to constrain this configuration can have unexpected results. Using override-enabled objects allows descendant domain administrators to tailor Global configurations to their local environments.

Leave matching criteria empty whenever possible, especially those for security zones, network objects, and port objects. When you specify multiple criteria, the system must match against every combination of the contents of the criteria you specify.

Geolocation in Network Conditions

Some rules can match traffic using the geographical location of the source or destination. If a rule type supports geolocation, you can mix network and geolocation criteria. To ensure you are using up-to-date geolocation data to filter your traffic, Cisco strongly recommends you regularly update the geolocation database (GeoDB).

Original Client in Network Conditions (Filtering Proxied Traffic)

For some rules, you can handle proxied traffic based on the originating client. Use a source network condition to specify proxy servers, then add an original client constraint to specify original client IP addresses. The system uses a packet's X-Forwarded-For (XFF), True-Client-IP, or custom-defined HTTP header field to determine original client IP.

Traffic matches the rule if the proxy's IP address matches the rule's source network constraint, **and** the original client's IP address matches the rule's original client constraint. For example, to allow traffic from a specific original client address, but only if it uses a specific proxy, create three access control rules:

Access Control Rule 1: Blocks non-proxied traffic from a specific IP address (209.165.201.1)

Source Networks: 209.165.201.1
Original Client Networks: none/any
Action: Block

Access Control Rule 2: Allows proxied traffic from the same IP address, but only if the proxy server for that traffic is one you choose (209.165.200.225 or 209.165.200.238)

Source Networks: 209.165.200.225 and 209.165.200.238
Original Client Networks: 209.165.201.1
Action: Allow

Access Control Rule 3: Blocks proxied traffic from the same IP address if it uses any other proxy server.

Source Networks: any
Original Client Networks: 209.165.201.1
Action: Block

Rules with Network Conditions

Rule Type	Supports Geolocation Constrains?	Supports Original Client Constraints?
Access control	yes	yes
Prefilter	no	no
SSL	yes	no
DNS (source networks only)	no	no
Identity	yes	no
Network analysis	no	no
QoS	yes	yes

Configuring Network Conditions

Procedure

Step 1 In the rule editor, click **Networks**.

Step 2 Find and choose the predefined networks you want to add from the **Available Networks** list.

If the rule supports geolocation, you can mix network and geolocation criteria in the same rule:

- Networks—Click **Networks** to choose networks.
- Geolocation—Click **Geolocation** to choose geolocation objects.

Step 3 (Optional) If the rule supports original client constraints, under **Source Networks**, configure the rule to handle proxied traffic based on its original client:

- Source/Proxy—Click **Source** to specify proxy servers.
- Original Client—Click **Original Client** to add a network as an original client constraint. In proxied connections, the original client's IP address must match one of these networks to match the rule.

Step 4 Click **Add to Source**, **Add to Original Client**, or **Add to Destination**, or drag and drop.

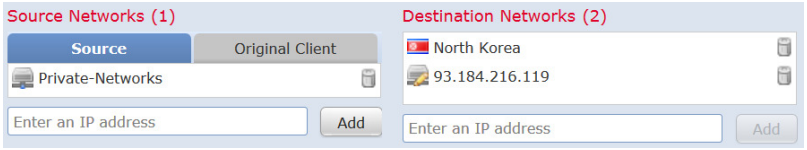
Step 5 Add networks that you want to specify manually. Enter a source, original client, or destination IP address or address block, then click **Add**.

Note The system builds a separate network map for each leaf domain. In a multidomain deployment, using literal IP addresses to constrain this configuration can have unexpected results. Using override-enabled objects allows descendant domain administrators to tailor Global configurations to their local environments.

Step 6 Save or continue editing the rule.

Example: Network Condition in an Access Control Rule

The following graphic shows the network condition for an access control rule that blocks connections originating from your internal network and attempting to access resources either in North Korea or on 93.184.216.119 (example.com).



In this example, a network object group called Private Networks (that comprises the IPv4 and IPv6 Private Networks network objects, not shown) represents your internal networks. The example also manually specifies the example.com IP address, and uses a system-provided North Korea geolocation object to represent North Korea IP addresses.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes.

Tunnel Endpoint Conditions

Tunnel endpoint conditions are specific to tunnel rules. They are similar to the network conditions for other rule types.

Tunnel endpoint conditions control certain types of plaintext, passthrough tunnels (see Encapsulation Conditions, on page 15) by their source and destination IP address, using outer encapsulation headers. These are the IP addresses of the tunnel endpoints—the routed interfaces of the network devices on either side of the tunnel.

Tunnel rules are bidirectional by default, and handle all matching tunnels between any of the source endpoints and any of the destination endpoints. However, you can configure unidirectional tunnel rules that match source-to-destination traffic only; see Tunnel and Prefilter Rule Components.

You can use predefined network objects to build tunnel endpoint conditions, or manually specify individual IP addresses or address blocks.



Note The system builds a separate network map for each leaf domain. In a multidomain deployment, using literal IP addresses to constrain this configuration can have unexpected results. Using override-enabled objects allows descendant domain administrators to tailor Global configurations to their local environments.

Configuring Tunnel Endpoint Conditions

Tunnel endpoint conditions apply to Firepower Threat Defense devices only.

Procedure

Step 1 In the rule editor, click **Tunnel Endpoints**.

Step 2 Find and choose the predefined networks you want to add from the **Available Tunnel Endpoints** list.

Because tunnel endpoints are simply the IP addresses of the routed interfaces of the network devices on either side of the tunnel, you can use network objects to build tunnel endpoint conditions.

Step 3 Click **Add to Source** or **Add to Destination**, or drag and drop.

Tunnel rules are bidirectional by default so they can handle all traffic between the two endpoints. However, if you choose to **Match tunnels only from source**, the tunnel rule matches source-to-destination traffic only.

Step 4 Add endpoints that you want to specify manually. Enter a source or destination IP address or address block, then click **Add**.

Note The system builds a separate network map for each leaf domain. In a multidomain deployment, using literal IP addresses to constrain this configuration can have unexpected results. Using override-enabled objects allows descendant domain administrators to tailor Global configurations to their local environments.

Step 5 Save or continue editing the rule.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes.

VLAN Conditions

VLAN rule conditions control VLAN-tagged traffic, including Q-in-Q (stacked VLAN) traffic. The system uses the innermost VLAN tag to filter VLAN traffic, with the exception of the prefilter policy, which uses the outermost VLAN tag in its rules.

Note the following Q-in-Q support:

- NGIPSv, Firepower 7000, Firepower 8000—Supports Q-in-Q for all interface types.
- ASA FirePOWER module—Does not support Q-in-Q (supports only one VLAN tag).
- FTD on Firepower 4100/9300—Does not support Q-in-Q (supports only one VLAN tag).
- FTD on all other models:
  - Inline sets and passive interfaces—Supports Q-in-Q, up to 2 VLAN tags.
  - Firewall interfaces—Does not support Q-in-Q (supports only one VLAN tag).

You can use predefined objects to build VLAN conditions, or manually enter any VLAN tag from **1** to **4094**. Use a hyphen to specify a range of VLAN tags.

You can specify a maximum of 50 VLAN conditions.



Note The system builds a separate network map for each leaf domain. In a multidomain deployment, using literal VLAN tags to constrain this configuration can have unexpected results. Using override-enabled objects allows descendant domain administrators to tailor Global configurations to their local environments.

Rules with VLAN Conditions

The folllowing rule types support VLAN conditions:

- Access control



Note For FTD, VLAN tags in access rules only apply to inline sets; they cannot be used in access rules applied to firewall interfaces.

- Tunnel and prefilter (uses outermost VLAN tag)
- SSL
- DNS
- Identity
- Network analysis

Port and ICMP Code Conditions

Port conditions allow you to control traffic by its source and destination ports. Depending on the rule type, "port" can represent any of the following:

- TCP and UDP—You can control TCP and UDP traffic based on the transport layer protocol. The system represents this configuration using the protocol number in parentheses, plus an optional associated port or port range. For example: TCP(6)/22.
- ICMP—You can control ICMP and ICMPv6 (IPv6-ICMP) traffic based on its internet layer protocol plus an optional type and code. For example: ICMP(1):3:3.
- No port—You can control traffic using other protocols that do not use ports.

Leave matching criteria empty whenever possible, especially those for security zones, network objects, and port objects. When you specify multiple criteria, the system must match against every combination of the contents of the criteria you specify.

Best Practices for Port-Based Rules

Specifying ports is the traditional way to target applications. However, applications can be configured to use unique ports to bypass access control blocks. Thus, whenever possible, use application filtering criteria rather than port criteria to target traffic.

Application filtering is also recommended for applications, like FTP, that open separate channels dynamically for control vs. data flow. Using port-based access control rules can prevent these kinds of applications from performing correctly, and could result in blocking desirable connections.

Using Source and Destination Port Constraints

If you add both source and destination port constraints, you can only add ports that share a single transport protocol (TCP or UDP). For example, if you add DNS over TCP as a source port, you can add Yahoo Messenger Voice Chat (TCP) as a destination port but not Yahoo Messenger Voice Chat (UDP).

If you add only source ports or only destination ports, you can add ports that use different transport protocols. For example, you can add both DNS over TCP and DNS over UDP as source port conditions in a single access control rule.

Matching Non-TCP Traffic with Port Conditions

Although you can configure port conditions to match non-TCP traffic, there are some restrictions:

- Access control rules—For Classic devices, you can match GRE-encapsulated traffic with an access control rule by using the GRE (47) protocol as a destination port condition. To a GRE-constrained rule, you can add only network-based conditions: zone, IP address, port, and VLAN tag. Also, the system uses outer headers to match **all** traffic in access control policies with GRE-constrained rules. For Firepower Threat Defense devices, use tunnel rules in the prefilter policy to control GRE-encapsulated traffic.
- SSL rules—SSL rules support TCP port conditions only.

Caution Adding the first or removing the last active authentication rule when SSL decryption is disabled (that is, when the access control policy does not include 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.

Note that an active authentication rule has either an **Active Authentication** rule action, or a **Passive Authentication** rule action with **Use active authentication if passive or VPN identity cannot be established** selected.

- IMCP echo—A destination ICMP port with the type set to 0 or a destination ICMPv6 port with the type set to 129 only matches unsolicited echo replies. ICMP echo replies sent in response to ICMP echo requests are ignored. For a rule to match on any ICMP echo, use ICMP type 8 or ICMPv6 type 128.

Rules with Port Conditions

The following rules support port conditions:

- Access control
- Prefilter
- SSL (supports TCP traffic only)
- Identity (active authentication supports TCP traffic only)
- QoS

Configuring Port Conditions

Procedure

Step 1 In the rule editor, click **Ports**.

Step 2 Find and choose the predefined ports you want to add from the **Available Ports** list.

Step 3 Click **Add to Source** or **Add to Destination**, or drag and drop.

Step 4 Add any source or destination ports that you want to specify manually:

- Source—Choose a **Protocol**, enter a single **Port** from 0 to 65535, and click **Add**.
- Destination (non-ICMP)—Choose or enter a **Protocol**. If you do not want to specify a protocol, or if you choose **TCP** or **UDP**, enter a single **Port** from 0 to 65535. Click **Add**.
- Destination (ICMP)—Choose **ICMP** or **IPv6-ICMP** from the **Protocol** drop down list, then choose a **Type** and related **Code** in the pop-up window that appears. For more information on ICMP types and codes, see the Internet Assigned Numbers Authority (IANA) website.

Step 5 Save or continue editing the rule.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes.

Encapsulation Conditions

Encapsulation conditions are specific to tunnel rules.

These conditions control certain types of plaintext, passthrough tunnels by their encapsulation protocol. You must choose at least one protocol to match before you can save the rule. You can choose:

- GRE (47)
- IP-in-IP (4)
- IPv6-in-IP (41)
- Teredo (UDP (17)/3455)

Application Conditions (Application Control)

When the system analyzes IP traffic, it can identify and classify the commonly used applications on your network. This discovery-based *application awareness* is the basis for *application control*—the ability to control application traffic.

System-provided *application filters* help you perform application control by organizing applications according to basic characteristics: type, risk, business relevance, category, and tags. You can create reuseable user-defined filters based on combinations of the system-provided filters, or on custom combinations of applications.

At least one detector must be enabled for each application rule condition in the policy. If no detector is enabled for an application, the system automatically enables all system-provided detectors for the application; if none exist, the system enables the most recently modified user-defined detector for the application. For more information about application detectors, see Application Detector Fundamentals.

You can use both application filters and individually specified applications to ensure complete coverage. However, understand the following note before you order your access control rules.

As part of application control, you can also use access control rules to enforce content restriction (such as Safe Search and YouTube EDU).



Caution Failure to set up your access control rules properly can have unexpected results, including traffic being allowed that should be blocked. In general, application control rules should be lower in your access control list because it takes longer for those rules to match than rules based on IP address, for example.

Access control rules that use *specific* conditions (such as networks and IP addresses) should be ordered *before* rules that use general conditions (such as applications). If you're familiar with the Open Systems Interconnect (OSI) model, use similar numbering in concept. Rules with conditions for layers 1, 2, and 3 (physical, data link, and network) should be ordered first in your access control rules. Conditions for layers 5, 6, and 7 (session, presentation, and application) should be ordered later in your access control rules. For more information about the OSI model, see this Wikipedia article.

Benefits of Application Filters

Application filters help you quickly configure application control. For example, you can easily use system-provided filters to create an access control rule that identifies and blocks all high risk, low business relevance applications. If a user attempts to use one of those applications, the system blocks the session.

Using application filters simplifies policy creation and administration. It assures you that the system controls application traffic as expected. Because Cisco frequently updates and adds application detectors via system and vulnerability database (VDB) updates, you can ensure that the system uses up-to-date detectors to monitor application traffic. You can also create your own detectors and assign characteristics to the applications they detect, automatically adding them to existing filters.

Configurations with Application Conditions

The configurations in the following table help you perform application control. The table also shows how you can constrain application control, depending on the configuration.

Configuration	Type, Risk, Relevance, Category	Tags	User-Defined Filters	Content Restriction
Access control rules	yes	yes	yes	yes
SSL rules	yes	no; automatically constrained to encrypted application traffic by the SSL Protocol tag	no	no
Identity rules (to exempt applications from active authentication)	yes	no; automatically constrained by the User-Agent Exclusion tag	no	no
QoS rules	yes	yes	yes	no
User-defined application filter in the object manager	yes	yes	no; you cannot nest user-defined filters	no
Intelligent Application Bypass (IAB)	yes	yes	yes	no

Related Topics

Overview: Application Detection

Configuring Application Conditions and Filters

To build an application condition or filter, choose the applications whose traffic you want to control from a list of available applications. Optionally (and recommended), constrain the available applications using filters. You can use filters and individually specified applications in the same condition.

Before you begin

- Adaptive profiling **must** be enabled (its default state) as described in Configuring Adaptive Profiles for access control rules to perform application control.
- For Classic device models, you must have the Control license to configure these conditions.

Procedure

Step 1 Invoke the rule or configuration editor:

- Access control, SSL, QoS rule condition—In the rule editor, click **Applications**.
- Identity rule condition—In the rule editor, click **Realms & Settings** and enable active authentication; see Create an Identity Rule.
- Application filter—On the Application Filters page of the object manager, add or edit an application filter. Provide a unique **Name** for the filter.
- Intelligent Application Bypass (IAB)—In the access control policy editor, click **Advanced**, edit IAB settings, then click **Bypassable Applications and Filters**.

Step 2 (Optional) For an access control rule, enable content restriction features by clicking dimmed for **Safe search** ( ) or **YouTube EDU** ( ) and setting related options.

For additional configuration requirements, see Using Access Control Rules to Enforce Content Restriction.

In most cases, enabling content restriction populates the condition's **Selected Applications and Filters** list with the appropriate values. The system does not automatically populate the list if applications or filters related to content restriction are already present in the list when you enable content restriction.

Continue with the procedure to refine your application and filter selections, or skip to saving the rule.

Step 3 Find and choose the applications you want to add from the **Available Applications** list.

To constrain the applications displayed in **Available Applications**, choose one or more **Application Filters** or search for individual applications.

Tip Click **Information** ( ) next to an application to display summary information and internet search links. **Unlock** marks applications that the system can identify only in decrypted traffic.

When you choose filters, singly or in combination, the Available Applications list updates to display only the applications that meet your criteria. You can choose system-provided filters in combination, but not user-defined filters.

- Multiple filters for the same characteristic (risk, business relevance, and so on)—Application traffic must match only one of the filters. For example, if you choose both the medium and high-risk filters, the Available Applications list displays all medium and high-risk applications.
- Filters for different application characteristics—Application traffic must match both filter types. For example, if you choose both the high-risk and low business relevance filters, the Available Applications list displays only applications that meet both criteria.

Step 4 Click **Add to Rule**, or drag and drop.

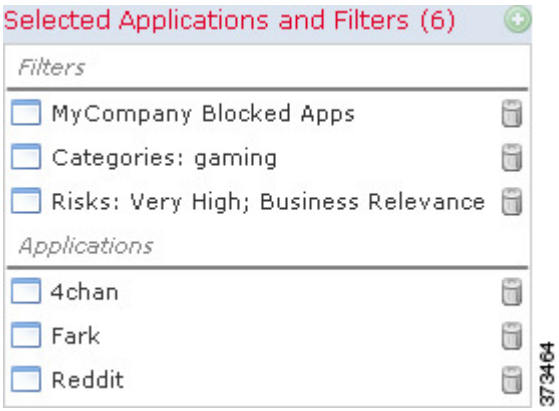
Tip Before you add more filters and applications, click **Clear Filters** to clear your current choices.

The web interface lists filters added to a condition above and separately from individually added applications.

Step 5 Save or continue editing the rule or configuration.

Example: Application Condition in an Access Control Rule

The following graphic shows the application condition for an access control rule that blocks a user-defined application filter for MyCompany, all applications with high risk and low business relevance, gaming applications, and some individually selected applications.



What to do next

- Deploy configuration changes; see Deploy Configuration Changes.

Application Characteristics

The system characterizes each application that it detects using the criteria described in the following table. Use these characteristics as application filters.

Table 1: Application Characteristics

Characteristic	Description	Example
Type	Application protocols represent communications between hosts. Clients represent software running on a host. Web applications represent the content or requested URL for HTTP traffic.	HTTP and SSH are application protocols. Web browsers and email clients are clients. MPEG video and Facebook are web applications.
Risk	The likelihood that the application is being used for purposes that might be against your organization's security policy.	Peer-to-peer applications tend to have a very high risk.

Characteristic	Description	Example
Business Relevance	The likelihood that the application is being used within the context of your organization's business operations, as opposed to recreationally.	Gaming applications tend to have a very low business relevance.
Category	A general classification for the application that describes its most essential function. Each application belongs to at least one category.	Facebook is in the social networking category.
Tag	Additional information about the application. Applications can have any number of tags, including none.	Video streaming web applications often are tagged `high bandwidth` and `displays ads`.

Best Practices for Application Control

Keep in mind the following guidelines and limitations for application control:

Ensuring that Adaptive Profiling is Enabled

If adaptive profiling is not enabled (its default state), access control rules cannot perform application control.

Automatically Enabling Application Detectors

If no detector is enabled for an application you want to detect, the system automatically enables all system-provided detectors for the application. If none exist, the system enables the most recently modified user-defined detector for the application.

Configure Your Policy to Examine the Packets That Must Pass Before an Application Is Identified

The system cannot perform application control, including Intelligent Application Bypass (IAB) and rate limiting, before *both* of the following occur:

- A monitored connection is established between a client and server
- The system identifies the application in the session

This identification should occur in 3 to 5 packets, or after the server certificate exchange in the SSL handshake if the traffic is encrypted.

Important! To ensure that your system examines these initial packets, see Specify a Policy to Handle Packets That Pass Before Traffic Identification.

If early traffic matches all other criteria but application identification is incomplete, the system allows the packet to pass and the connection to be established (or the SSL handshake to complete). After the system completes its identification, the system applies the appropriate action to the remaining session traffic.

Create Separate Rules for URL and Application Filtering

Create separate rules for URL and application filtering whenever possible, because combining application and URL criteria can lead to unexpected results, especially for encrypted traffic.

Rules that include both application and URL criteria should come after application-only or URL-only rules, unless the application+URL rule is acting as an exception to a more general application-only or URL-only rule.

URL Rules Before Application and Other Rules

For the most effective URL matching, place rules that include URL conditions before other rules, particularly if the URL rules are block rules and the other rules meet both of the following criteria:

- They include application conditions.
- The traffic to be inspected is encrypted.

Application Control for Encrypted and Decrypted Traffic

The system can identify and filter encrypted and decrypted traffic:

- Encrypted traffic—The system can detect application traffic encrypted with StartTLS, including SMTPS, POPS, FTPS, TelnetS, and IMAPS. In addition, it can identify certain encrypted applications based on the Server Name Indication in the TLS ClientHello message, or the subject distinguished name value from the server certificate. These applications are tagged `SSL Protocol`; in an SSL rule, you can choose only these applications. Applications without this tag can only be detected in unencrypted or decrypted traffic.
- Decrypted traffic—The system assigns the `decrypted traffic` tag to applications that the system can detect in decrypted traffic only, not encrypted or unencrypted.

Exempting Applications from Active Authorization

In an identity policy, you can exempt certain applications from active authentication, allowing traffic to continue to access control. These applications are tagged `User-Agent Exclusion`. In an identity rule, you can choose only these applications.

Handling Application Traffic Packets Without Payloads

When performing access control, the system applies the default policy action to packets that do not have a payload in a connection where an application is identified.

Handling Referred Application Traffic

To handle traffic referred by a web server, such as advertisement traffic, match the referred application rather than the referring application.

Controlling Application Traffic That Uses Multiple Protocols (Skype, Zoho)

Some applications use multiple protocols. To control their traffic, make sure your access control policy covers all relevant options. For example:

- Skype—To control Skype traffic, choose the **Skype** tag from the **Application Filters** list rather than selecting individual applications. This ensures that the system can detect and control all Skype traffic the same way.
- Zoho—To control Zoho mail, choose *both* **Zoho** and **Zoho mail** from the Available Application list.

Search Engines Supported for Content Restriction Features

The system supports Safe Search filtering for specific search engines only. The system assigns the `safesearch supported` tag to application traffic from these search engines.

Controlling Evasive Application Traffic

See Application-Specific Notes and Limitations, on page 22.

Additional Guidelines for Rule Ordering for Application Control

For guidelines about rule ordering for application control, see Best Practices for Configuring Application Control, on page 21.

Related Topics

Inspection of Packets That Pass Before Traffic Is Identified
Special Considerations for Application Detection

Best Practices for Configuring Application Control

We recommend controlling applications' access to the network as follows:

- To allow or block application access from a less secure network to a more secure network: Use **Port** (Selected Destination Port) conditions on the access control rule

  For example, allow ICMP traffic from the internet (less secure) to an internal network (more secure.)

- To allow or block applications being accessed by user groups: Use **Application** conditions on the access control rule

  For example, block Facebook from being accessed by members of the Contractors group



Caution Failure to set up your access control rules properly can have unexpected results, including traffic being allowed that should be blocked. In general, application control rules should be lower in your access control list because it takes longer for those rules to match than rules based on IP address, for example.

Access control rules that use *specific* conditions (such as networks and IP addresses) should be ordered *before* rules that use general conditions (such as applications). If you're familiar with the Open Systems Interconnect (OSI) model, use similar numbering in concept. Rules with conditions for layers 1, 2, and 3 (physical, data link, and network) should be ordered first in your access control rules. Conditions for layers 5, 6, and 7 (session, presentation, and application) should be ordered later in your access control rules. For more information about the OSI model, see this Wikipedia article.

The following table provides an example of how to set up your access control rules:

Type of control	Action	Zones, Networks, VLAN Tags	Users	Applications	Ports	URLs	SGT/ISE Attributes	Inspection, Logging, Comments
Application from more secure to less secure network when application uses a port (for example, SSH)	Your choice (**Allow** in this example)	Destination zones or networks using the outside interface	Any	Do not set	Available Ports : **SSH** Add to **Selected Destination Ports**	Any	Use only with ISE/ISE-PIC.	Any
Application from more secure to less secure network when application does not use a port (for example, ICMP)	Your choice (**Allow** in this example)	Destination zones or networks using the outside interface	Any	Do not set	Selected Destination Ports **Protocol**: **ICMP** **Type**: **Any**	Do not set	Use only with ISE/ISE-PIC.	Any
Application access by a user group	Your choice (**Block** in this example)	Your choice	Choose a user group (Contractors group in this example)	Choose the name of the application (**Facebook** in this example)	Do not set	Do not set	Use only with ISE/ISE-PIC.	Your choice

Application-Specific Notes and Limitations

- Office 365 Admin Portal:

  Limitation: If the access policy has logging enabled at the beginning as well as at the end, the first packet will be detected as Office 365 and the end of connection will be detected as Office 365 Admin Portal. This should not affect blocking.

- Skype:

  See Best Practices for Application Control, on page 19

- GoToMeeting

  In order to fully detect GoToMeeting, your rule must include all of the following applications:

  - GoToMeeting
  - Citrix Online
  - Citrix GoToMeeting Platform
  - LogMeIn
  - STUN

- Zoho:

  See Best Practices for Application Control, on page 19

- Evasive applications such as Bittorrent, Tor, Psiphon, and Ultrasurf:

  For evasive applications, only the highest-confidence scenarios are detected by default. If you need to take action on this traffic (such as block or implement QoS), it may be necessary to configure more aggressive detection with better effectiveness. To do this, contact TAC to review your configurations as these changes may result in false positives.

- WeChat:

  It is not possible to selectively block WeChat Media if you allow WeChat.

Troubleshoot Application Control Rules

If your application control rules don't function as you expect, use the guidelines discussed in this section.

We recommend controlling applications' access to the network as follows:

- To allow or block application access from a less secure network to a more secure network: Use **Port** (Selected Destination Port) conditions on the access control rule

  For example, allow ICMP traffic from the internet (less secure) to an internal network (more secure.)

- To allow or block applications being accessed by user groups: Use **Application** conditions on the access control rule

  For example, block Facebook from being accessed by members of the Contractors group



Caution Failure to set up your access control rules properly can have unexpected results, including traffic being allowed that should be blocked. In general, application control rules should be lower in your access control list because it takes longer for those rules to match than rules based on IP address, for example.

Access control rules that use *specific* conditions (such as networks and IP addresses) should be ordered *before* rules that use general conditions (such as applications). If you're familiar with the Open Systems Interconnect (OSI) model, use similar numbering in concept. Rules with conditions for layers 1, 2, and 3 (physical, data link, and network) should be ordered first in your access control rules. Conditions for layers 5, 6, and 7 (session, presentation, and application) should be ordered later in your access control rules. For more information about the OSI model, see this Wikipedia article.

The following table provides an example of how to set up your access control rules:

Type of control	Action	Zones, Networks, VLAN Tags	Users	Applications	Ports	URLs	SGT/ISE Attributes	Inspection, Logging, Comments
Application from more secure to less secure network when application uses a port (for example, SSH)	Your choice (**Allow** in this example)	Destination zones or networks using the outside interface	Any	Do not set	Available Ports : **SSH** Add to **Selected Destination Ports**	Any	Use only with ISE/ISE-PIC.	Any
Application from more secure to less secure network when application does not use a port (for example, ICMP)	Your choice (**Allow** in this example)	Destination zones or networks using the outside interface	Any	Do not set	Selected Destination Ports **Protocol**: **ICMP** **Type**: **Any**	Do not set	Use only with ISE/ISE-PIC.	Any
Application access by a user group	Your choice (**Block** in this example)	Your choice	Choose a user group (Contractors group in this example)	Choose the name of the application (**Facebook** in this example)	Do not set	Do not set	Use only with ISE/ISE-PIC.	Your choice

Initial Packets Are Passing Uninspected

See Inspection of Packets That Pass Before Traffic Is Identified and subtopics.

Related Topics

Best Practices for Ordering Rules

URL Conditions (URL Filtering)

Use URL conditions to control the websites that users on your network can access.

For complete information, see URL Filtering.

User, Realm, and ISE Attribute Conditions (User Control)

You can perform *user control* with the *authoritative user identity data* collected by the Firepower System.

Identity sources monitor users as they log in and out, or as they authenticate using Microsoft Active Directory (AD) or LDAP credentials. You can then configure rules that use this collected identity data to handle traffic based on the logged-in authoritative user associated with a monitored host. A user remains associated with a

host until the user logs off (as reported by an identity source), a realm times out the session, or you delete the user data from the system's database.

For information on the authoritative user identity sources supported in your version of the Firepower System, see About User Identity Sources.

You can use the following rule conditions to perform user control:

- User and realm conditions—Match traffic based on the logged-in authoritative user of a host. You can control traffic based on realms, individual users, or the groups those users belong to.
- ISE attribute conditions—Match traffic based on a user's ISE-assigned Security Group Tag (SGT), Device Type (also referred to as Endpoint Profile), or Location IP (also referred to as Endpoint Location). Requires that you configure ISE as an identity source.

Note The ISE-PIC identity source does not provide ISE attribute data.

Note In some rules, custom SGT conditions can match traffic tagged with SGT attributes that were **not** assigned by ISE. This is not considered user control, and only works if you are not using ISE as an identity source; see Custom SGT Conditions, on page 29.

Rules with User Conditions

Rule Type	Supports User and Realm Conditions?	Supports ISE Attribute Conditions?
Access control	yes	yes
SSL	yes	no
QoS	yes	yes

Related Topics

The User Agent Identity Source
The ISE/ISE-PIC Identity Source
The Terminal Services (TS) Agent Identity Source
The Captive Portal Identity Source

User Control Prerequisites

Configure Identity Sources/Authentication Methods

Configure identity sources for the types of authentication you want to perform. For more information, see About User Identity Sources.

If you configure an ISE/ISE/PIC , user agent, or TS Agent device to monitor a large number of user groups, or if you have a very large number of users mapped to hosts on your network, the system may drop user mappings based on groups, due to your Firepower Management Center user limit. As a result, rules with realm, user, or user group conditions may not match traffic as expected.

Configure Realms

Configure a realm for each AD or LDAP server you want to monitor, including your ISE/ISE-PIC, User Agent, and TS Agent servers, and perform a user download. For more information, see Create a Realm.

Note If you are configuring an ISE SGT attribute rule condition, configuring a realm is optional. The ISE SGT attribute rule condition can be configured in policies with or without an associated identity policy (where realms are invoked).

When you configure a realm, you specify the users and user groups whose activity you want to monitor. Including a user group automatically includes all of that group's members, including members of any secondary groups. However, if you want to use the secondary group as a rule criterion, you must explicitly include the secondary group in the realm configuration.

For each realm, you can enable automatic download of user data to refresh authoritative data for users and user groups.

Create Identity Policies

Create an identity policy to associate the realm with an authentication method, and associate that policy with access control. For more information, see Create an Identity Policy.

Policies that perform user control on a device (access control, SSL, QoS) share an identity policy. That identity policy determines the realms, users, and groups that you can use in rules affecting traffic on those devices.

Before you configure a user condition in a QoS rule, you **must** make sure the devices targeted by the QoS policy are using the correct identity policy, as defined in the access control policy deployed to the devices. Because the QoS policy and access control policy deployed to the same device are not explicitly linked, the QoS rule editor can allow you to choose invalid realms, users, and groups. These invalid elements are those from identity policies that exist on the Firepower Management Center, but that are not applied to the QoS-targeted devices. If you use these elements, the system cannot determine that you made an invalid choice until deploy-time.

Configuring User and Realm Conditions

You can constrain a rule by realm, or by users and user groups within that realm.

Before you begin

- Fulfill the user control prerequsities described in User, Realm, and ISE Attribute Conditions (User Control), on page 24.
- For Classic device models, you must have the Control license to configure these conditions.

Procedure

Step 1 In the rule editor, click **Users**.

Step 2 (Optional) Find and choose the realm you want to use from the **Available Realms**.

Step 3 (Optional) Further constrain the rule by choosing users and groups from the **Available Users** list.

Step 4 Click **Add to Rule**, or drag and drop.

Step 5 Save or continue editing the rule.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes.

Configuring ISE Attribute Conditions

Before you begin

- Fulfill the user control prerequisites described in User, Realm, and ISE Attribute Conditions (User Control), on page 24.
- For Classic device models, you must have the Control license to configure these conditions.

Procedure

Step 1 In the rule editor, click the following for ISE attribute conditions:

- Access control—Click **SGT/ISE Attributes**.

You can use ISE-assigned Security Group Tags (SGTs) to constrain ISE attribute conditions. To use custom SGTs in access control rules, see Custom SGT Conditions, on page 29.

Step 2 Find and choose the ISE attributes you want to use from the **Available Attributes** list:

- Security Group Tag (SGT)
- Device Type (also referred to as Endpoint Profile)
- QoS—Click **ISE Attributes**.
- Location IP (also referred to as Endpoint Location)

Step 3 Further constrain the rule by choosing attribute metadata from the **Available Metadata** list. Or, keep the default: **any**.

Step 4 Click **Add to Rule**, or drag and drop.

Step 5 (Optional) Constrain the rule with an IP address in the **Add a Location IP Address** field, then click **Add**.

The system builds a separate network map for each leaf domain. In a multidomain deployment, using literal IP addresses to constrain this configuration can have unexpected results.

Step 6 Save or continue editing the rule.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes.

Troubleshoot User Control

If you notice unexpected user rule behavior, consider tuning your rule, identity source, or realm configurations. For other related troubleshooting information, see:

- Troubleshoot the User Agent Identity Source
- Troubleshoot the ISE/ISE-PIC or Cisco TrustSec Issues
- Troubleshoot the TS Agent Identity Source
- Troubleshoot the Captive Portal Identity Source
- Troubleshoot Realms and User Downloads

Rules targeting realms, users, or user groups are not matching traffic

If you configure a TS Agent, user agent, or ISE/ISE-PIC device to monitor a large number of user groups, or if you have a very large number of users mapped to hosts on your network, the system may drop user records due to your Firepower Management Center user limit. As a result, rules with user conditions may not match traffic as expected.

Rules targeting user groups or users within user groups are not matching traffic as expected

If you configure a rule with a user group condition, your LDAP or Active Directory server must have user groups configured. The system cannot perform user group control if the server organizes the users in basic object hierarchy.

Rules targeting users in secondary groups are not matching traffic as expected

If you configure a rule with a user group condition that includes or excludes users who are members of a secondary group on your Active Directory server, your server may be limiting the number of users it reports.

By default, Active Directory servers limit the number of users they report from secondary groups. You must customize this limit so that all of the users in your secondary groups are reported to the Firepower Management Center and eligible for use in rules with user conditions.

Rules are not matching users when seen for the first time

After the system detects activity from a previously-unseen user, the system retrieves information about them from the server. Until the system successfully retrieves this information, activity seen by this user is *not* handled by matching rules. Instead, the user session is handled by the next rule it matches (or the policy's default action, if applicable).

For example, this might explain when:

- Users who are members of user groups are not matching rules with user group conditions.
- Users who were reported by a TS Agent, user agent, or ISE/ISE-PIC device are not matching rules, when the server used for user data retrieval is an Active Directory server.

Note that this might also cause the system to delay the display of user data in event views and analysis tools.

Rules are not matching all ISE users

This is expected behavior. You can perform user control on ISE users who were authenticated by an Active Directory domain controller. You cannot perform user control on ISE users who were authenticated by an LDAP, RADIUS, or RSA domain controller.

Rules are not matching all ISE/ISE-PIC users

This is expected behavior. You can perform user control on ISE/ISE-PIC users who were authenticated by an Active Directory domain controller. You cannot perform user control on ISE/ISE-PIC users who were authenticated by an LDAP, RADIUS, or RSA domain controller.

Custom SGT Conditions

If you do not configure ISE/ISE-PIC as an identity source, you can control traffic using Security Group Tags (SGTs) that were **not** assigned by ISE. SGTs specify the privileges of traffic sources within a trusted network.

Custom SGT rule conditions use manually created SGT objects to filter traffic, rather than ISE SGTs obtained from the system's connection to an ISE server. These manually created SGT objects correspond to the SGT attributes on the traffic you want to control. Controlling traffic using custom SGTs is not considered user control.

Rules with Custom SGT Conditions

The following rules support custom SGT conditions:

- Access control
- QoS

ISE SGT vs Custom SGT Rule Conditions

Some rules allow you to control traffic based on assigned SGT. Depending on the rule type and your identity source configuration, you can use either ISE-assigned SGTs or custom SGTs to match traffic with assigned SGT attributes.



Note If you use ISE SGTs to match traffic, even if a packet does not have an assigned SGT attribute, the packet still matches an ISE SGT rule if the SGT associated with the packet's source IP address is known in ISE.

Condition Type	Requires	SGTs Listed in Rule Editor
ISE SGT	ISE identity source	SGTs obtained by querying the ISE server, with automatically updated metadata
Custom SGT	No ISE/ISE-PIC identity source	Static SGT objects you create

Related Topics

User, Realm, and ISE Attribute Conditions (User Control), on page 24

Autotransition from Custom SGTs to ISE SGTs

If you create rules that match custom SGTs, then configure ISE/ISE-PIC as an identity source, the system:

- Disables **Security Group Tag** options in the object manager. Although the system retains existing SGT objects, you cannot modify them or add new ones.
- Retains existing rules with custom SGT conditions. However, these rules do not match traffic. You also cannot add additional custom SGT criteria to existing rules, or create new rules with custom SGT conditions.

If you configure ISE, Cisco recommends that you delete or disable existing rules with custom SGT conditions. Instead, use ISE attribute conditions to match traffic with SGT attributes.

Related Topics

Configure ISE/ISE-PIC for User Control

Configuring Custom SGT Conditions

The following procedure describes how to filter traffic tagged with SGT attributes that were **not** assigned by ISE. This is not considered user control, and only works if you are not using ISE/ISE-PIC as an identity source; see ISE SGT vs Custom SGT Rule Conditions, on page 29.

Before you begin

- Disable ISE/ISE-PIC connections. Custom SGT matching does not work if you use ISE/ISE-PIC as an identity source.
- Configure Security Group Tag objects that correspond with the SGTs you want to match; see Creating Security Group Tag Objects.
- For Classic device models, you must have the Control license to configure these conditions.

Procedure

Step 1 In the rule editor, click **SGT/ISE Attributes**.

Step 2 Choose **Security Group Tag** from the **Available Attributes** list.

Step 3 In the **Available Metadata** list, find and choose a custom SGT object.

If you choose **Any**, the rule matches all traffic with an SGT attribute. For example, you might choose this value if you want an access control rule to block traffic from hosts that are not configured for TrustSec.

Step 4 Click **Add to Rule**, or drag and drop.

Step 5 Save or continue editing the rule.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes.

Troubleshooting Custom SGT Conditions

If you notice unexpected rule behavior, consider tuning your custom SGT object configuration.

Security Group Tag objects unavailable

Custom SGT objects are only available if you do not configure ISE/ISE-PIC as an identity source. For more information, see Autotransition from Custom SGTs to ISE SGTs, on page 30.

Searching for Rules

In many policies, you can search for and within rules. The system matches your input to rule names and condition values, including objects and object groups.

You cannot search for values in a Security Intelligence or URL list or feed.

Procedure

Step 1 In the policy editor, click **Rules**.

Step 2 Click **Search Rules**, enter a complete or partial search string, then press Enter.
The matching value is highlighted for each matching rule. A status message displays the current match and the total number of matches.

Step 3 View the rules you are interested in.

To navigate between matching rules, click **Next-Match** or **Previous-Match** .

What to do next

- Before you begin a new search, click **Clear** (×) to clear the search and any highlighting.

Filtering Rules by Device

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

Filter by device only works for rules that use zones or interface groups. (Otherwise a rule applies to all devices.)

The system uses a rule's interface constraints to determine if the rule affects a device. If you constrain a rule by interface (security zone or interface group condition), the device where that interface is located is affected by that rule. Rules with no interface constraint apply to any interface, and therefore every device.

QoS rules are always constrained by interface.

Procedure

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

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

Tip Hover your pointer over a rule criterion to see its value. If the criterion represents an object with device-specific overrides, the system displays the override value when you filter the rules list by only that device. If the criterion represents an object with domain-specific overrides, the system displays the override value when you filter the rules list by devices in that domain.

Step 3 Click **OK**.

Related Topics

Create and Edit Access Control Rules
Configure Prefiltering
Configuring QoS Rules
Configure NAT for Threat Defense

Identify Rules with Issues

The system will flag each rule that will prevent deploy (these are marked with a red icon) or that will never match traffic because another rule above it in the rule order will match instead (these are marked with a yellow icon).



Important The system does not flag rules that partially match other rules, which may also prevent some subsequent rules from matching.

Procedure

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

Step 2 Click a policy name.

Step 3 Do one or both of the following:

- Look for **Show Warnings** near the top of the window.

  If the system has not identified issues, this button will not appear.

  If there are issues, click this button to open a list of all rules with issues.

  To see all issues, click both tabs (Rule Errors and Rule Warnings).

  To locate a rule in the table of rules below, click the rule name in the error or warnings list.

- Select the **Show Rule Conflicts** check box.

  This will indicate each problem rule in the list with an Error (red) or Warning (yellow).

  If necessary, scroll down to see all rules in the policy.

Step 4 To view issue details, hover your pointer over the icon.

Step 5 Look for duplications that are not flagged because they are only partial matches and address them.

Step 6 If you make changes, you must click **Save** or deselect and reselect **Show Rule Conflicts** to evaluate the changed rules for conflicts.

What to do next

- Address any issues you see by removing or modifying the problematic rule.
- Examine your SSL and QoS policies for similar errors and warnings and address those issues.

Rule and Other Policy Warnings

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

Tip Hover your pointer over an icon to read the warning, error, or informational text.

Table 2: Policy Error Icons

Icon	Description	Example
Errors ( ) error	If a rule or configuration has an error, you cannot deploy until you correct the issue, even if you disable any affected rules.	A rule that performs category and reputation-based URL filtering is valid until you target a device that does not have a URL Filtering license. At that point, an error icon appears next to the rule, and you cannot deploy until you edit or delete the rule, retarget the policy, or enable the license.
Warning ( ) warning	You can deploy a policy that displays rule or other warnings. However, misconfigurations marked with warnings have no effect. If you disable a rule with a warning, the warning icon disappears. It reappears if you enable the rule without correcting the underlying issue.	Preempted rules or rules that cannot match traffic due to misconfiguration have no effect. This includes conditions using empty object groups, application filters that match no applications, excluded LDAP users, invalid ports, and so on. However, if a warning icon marks a licensing error or model mismatch, you cannot deploy until you correct the issue.
Information ( ) information	Information icons convey helpful information about configurations that may affect the flow of traffic. These issues do not prevent you from deploying.	With application control, the system might skip matching the first few packets of a connection against some rules, until the system identifies the application or web traffic in that connection. This allows connections to be established so that applications and HTTP requests can be identified.

Related Topics

Best Practices for Application Control, on page 19
Best Practices for URL Filtering