



Quality of Service

The following topics describe how to use the Quality of Service (QoS) feature to police network traffic using threat defense devices:

- [Introduction to QoS, on page 1](#)
- [About QoS Policies, on page 1](#)
- [Requirements and Prerequisites for QoS, on page 2](#)
- [Rate Limiting with QoS Policies, on page 3](#)

Introduction to QoS

Quality of Service, or QoS, rate limits (polices) network traffic that is allowed or trusted by access control. The system does not rate limit traffic that was fastpathed.

Though QoS is supported only on the routed interfaces of threat defense devices, it is not supported on site-to-site VPN and VTI interfaces.

Logging Rate-Limited Connections

There are no logging configurations for QoS. A connection can be rate limited without being logged, and you cannot log a connection simply because it was rate limited. To view QoS information in connection events, you must independently log the ends of the appropriate connections to the management center database.

Connection events for rate-limited connections contain information on how much traffic was dropped, and which QoS configurations limited the traffic. You can view this information in event views (workflows), dashboards, and reports.

About QoS Policies

QoS policies deployed to managed devices govern rate limiting. Each QoS policy can target multiple devices; each device can have one deployed QoS policy at a time.

The system matches traffic to QoS rules in the order you specify. The system rate limits traffic according to the first rule where all rule conditions match the traffic. Traffic that does not match any of the rules is not rate limited.



Note The total number of rules including QoS rules on the device cannot exceed 255. When this threshold is reached, a deployment warning message is displayed. You need to reduce the number of rules for a successful deployment.

You must constrain QoS rules by source or destination (routed) interfaces. The system enforces rate limiting *independently* on *each* of those interfaces; you cannot specify an aggregate rate limit for a set of interfaces.

QoS rules can also rate limit traffic by other network characteristics, as well as contextual information such as application, URL, user identity, and custom Security Group Tags (SGTs).

You can rate limit download and upload traffic independently. The system determines download and upload directions based on the connection initiator.



Note QoS is not subordinate to a main access control configuration; you configure QoS independently. However, the access control and QoS policies deployed to the same device share identity configurations; see [Associating Other Policies with Access Control](#).

QoS Policies and Multitenancy

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.

Administrators in ancestor domains can deploy the same QoS policy to devices in different descendant domains. Administrators in those descendant domains can use this read-only ancestor-deployed QoS policy, or replace it with a local policy.

Requirements and Prerequisites for QoS

Model Support

Threat Defense

Supported Domains

Any

User Roles

Admin

Access Admin

Network Admin

Rate Limiting with QoS Policies

To perform policy-based rate limiting, configure and deploy QoS policies to managed devices. Each QoS policy can target multiple devices; each device can have one deployed QoS policy at a time.

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 **Devices > QoS**.
- Step 2** Click **New Policy** to create a new QoS policy and, optionally, assign target devices; see [Creating a QoS Policy, on page 3](#).
- You can also **Copy** (📄) or **Edit** (✎) an existing policy.
- Step 3** Configure QoS rules; see [Configuring QoS Rules, on page 4](#) and [QoS Rule Conditions, on page 6](#).
- The Rules in the QoS policy editor lists each rule in evaluation order, and displays a summary of the rule conditions and rate limiting configurations. A right-click menu provides rule management options, including moving, enabling, and disabling.
- Helpful in larger deployments, you can **Filter by Device** to display only the rules that affect a specific device or group of devices. You can also search for and within rules; the system matches text you enter in the **Search Rules** field to rule names and condition values, including objects and object groups.
- Note** Properly creating and ordering rules is a complex task, but one that is essential to building an effective deployment. If you do not plan carefully, rules can preempt other rules, require additional licenses, or contain invalid configurations. Icons represent comments, warnings, and errors. If issues exist, click **Show Warnings** to display a list. For more information, see [Best Practices for Access Control Rules](#).
- Step 4** Click **Policy Assignments** to identify the managed devices targeted by the policy; see [Setting Target Devices for a QoS Policy, on page 4](#).
- If you identified target devices during policy creation, verify your choices.
- Step 5** Save the QoS policy.
- Step 6** Because this feature must allow some packets to pass, you must configure your system to examine those packets. See [Best Practices for Handling Packets That Pass Before Traffic Identification](#) and [Specify a Policy to Handle Packets That Pass Before Traffic Identification](#).
- Step 7** Deploy configuration changes.
-

Creating a QoS Policy

A new QoS policy with no rules performs no rate limiting.

Procedure

- Step 1** Choose **Devices** > **QoS**.
- Step 2** Click **New Policy**.
- Step 3** Enter a **Name** and, optionally, a **Description**.
- Step 4** (Optional) Choose the **Available Devices** where you want to deploy the policy, then click **Add to Policy**, or drag and drop to the **Selected Devices**. To narrow the devices that appear, type a search string in the **Search** field.
- You must assign devices before you deploy the policy.
- Step 5** Click **Save**.
-



What to do next

- Configure and deploy the QoS policy; see [Rate Limiting with QoS Policies, on page 3](#).

Setting Target Devices for a QoS Policy

Each QoS policy can target multiple devices; each device can have one deployed QoS policy at a time.

Procedure

- Step 1** In the QoS policy editor, click **Policy Assignments**.
- Step 2** Build your target list:
- Add—Choose 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 choose multiple devices, right-click, then choose **Delete Selected**.
 - Search—Enter a search string in the search field. Click **Clear** () to clear the search.
- Step 3** Click **OK** to save policy assignments.
- Step 4** Click **Save** to save the policy.
-

What to do next

- Deploy configuration changes.

Configuring QoS Rules

When you create or edit a rule, use the upper portion of the rule editor to configure general rule properties. Use the lower portion of the rule editor to configure rule conditions and comments.

Procedure

- Step 1** On Rules of the QoS policy editor:
- Add Rule—Click **Add Rule**.
 - Edit Rule—Click **Edit** (✎).
- Step 2** Enter a **Name**.
- Step 3** Configure rule components:
- Enabled—Specify whether the rule is **Enabled**.
 - Apply QoS On—Choose the interfaces you want to rate limit, either **Interfaces in Destination Interface Objects** or **Interfaces in Source Interface Objects**. Your choice must correspond with a populated interface constraint (not **any**).
 - Traffic Limit Per Interface—Enter a **Download Limit** and an **Upload Limit** in Mbits/sec. The default value of **Unlimited** prevent matching traffic from being rate limited in that direction.
 - Conditions—Click the corresponding condition you want to add. You must configure a source or destination interface condition, corresponding to your choice for **Apply QoS On**.
 - Comments—Click **Comments**. To add a comment click **New Comment**, enter a comment, and click **OK**. You can edit or delete this comment until you save the rule.
- For detailed information on rule components, see [QoS Rule Components, on page 5](#).
- Step 4** Save the rule.
- Step 5** In the policy editor, set the rule position. Click and drag or use the right-click menu to cut and paste.
- Rules are numbered starting at 1. The system matches traffic to rules in top-down order by ascending rule number. The first rule that traffic matches is the rule that handles that traffic. Proper rule order reduces the resources required to process network traffic and prevents rule preemption.
- Step 6** Click **Save** to save the policy.
-

What to do next

- Deploy configuration changes.

Related Topics

[Best Practices for Access Control Rules](#)

QoS Rule Components

State (Enabled/Disabled)

By default, rules are enabled. If you disable a rule, the system does not use it and stops generating warnings and errors for that rule.

Interfaces (Apply QoS On)

You cannot save a QoS rule that rate limits all traffic. For each QoS rule, you must apply QoS on either:

- Interfaces in Source Interface Objects—Rate limits traffic through the rule's source interfaces. If you choose this option, you must add at least one source interface constraint (cannot be **any**).
- Interfaces in Destination Interface Objects—Rate limits traffic through the rule's destination interfaces. If you choose this option, you must add at least one destination interface constraint (cannot be **any**).

Traffic Limit Per Interface

A QoS rule enforces rate limiting *independently* on *each* of the interfaces you specify with the Apply QoS On option. You cannot specify an aggregate rate limit for a set of interfaces.

You can rate limit traffic by Mbits per second. The default value of **Unlimited** prevents matching traffic from being rate limited.

You can rate limit download and upload traffic independently. The system determines download and upload directions based on the connection initiator.

If you specify a limit greater than the maximum throughput of an interface, the system does not rate limit matching traffic. Maximum throughput may be affected by an interface's hardware configuration, which you specify in each device's properties (**Devices > Device Management**).

Conditions

Conditions specify the specific traffic the rule handles. You can configure each rule with multiple conditions. Traffic must match all conditions to match the rule. Each condition type has its own tab in the rule editor. For more information, see [QoS Rule Conditions, on page 6](#).

Comments

Each time you save changes to a rule 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.

In the policy editor, the system displays how many comments a rule has. In the rule editor, use the Comments tab to view existing comments and add new ones.

QoS Rule Conditions

Conditions specify the specific traffic the rule handles. You can configure each rule with multiple conditions. Traffic must match all conditions to match the rule. Each condition type has its own tab in the rule editor. You can rate limit traffic using:

See one of the following sections for more information.

Related Topics

- [Interface Rule Conditions, on page 7](#)
- [Network Rule Conditions, on page 7](#)
- [User Rule Conditions, on page 7](#)
- [Application Rule Conditions, on page 8](#)
- [Port Rule Conditions, on page 9](#)
- [URL Rule Conditions, on page 10](#)
- [Custom SGT Rule Conditions, on page 11](#)

Interface Rule 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](#).



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.

Network Rule 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 You *cannot* use FDQN network objects in identity rules.



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.

User Rule Conditions

User rule conditions match traffic based the user who initiates the connection, or the group to which the user belongs. For example, you could configure a Block rule to prohibit anyone in the Finance group from accessing a network resource.

For access control rules only, you must first associate an identity policy with the access control policy as discussed in [Associating Other Policies with Access Control](#).

In addition to configuring users and groups for configured realms, you can set policies for the following Special Identities users:

- Failed Authentication: User that failed authentication with the captive portal.
- Guest: Users configured as guest users in the captive portal.

- No Authentication Required: Users that match an identity **No Authentication Required** rule action.
- Unknown: Users that cannot be identified; for example, users that are not downloaded by a configured realm.

Application Rule Conditions

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 reusable 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.

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.

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.

Related Topics

[Best Practices for Configuring Application Control](#)

Port Rule Conditions

Port conditions allow you to control traffic by its source and destination 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 threat defense, 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.

Port, Protocol, and ICMP Code Rule Conditions

Port conditions match traffic based on the 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 port. 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.
- Protocol—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. Note that application filtering is not available in prefilter rules.

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 destination port conditions in a single access control rule.

Matching Non-TCP Traffic with Port Conditions

You can match non-port-based protocols. By default, if you do not specify a port condition, you are matching IP traffic. 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 threat defense devices, use tunnel rules in the prefilter policy to control GRE-encapsulated traffic.
- Decryption rules—These rules support TCP port conditions only.
- ICMP 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.

URL Rule Conditions

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

For complete information, see [URL Filtering](#).

Custom SGT Rule 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.

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

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.

