Quality of Service

Quality of Service is a network management feature that:

- enables the policing of network traffic on Firewall Threat Defense devices.
- assists administrators in managing network traffic,
- provides guidelines for rate limiting.

QoS

QoS is a traffic management mechanism that

- rate limits (polices) network traffic that is allowed or trusted by access control
- does not rate limit traffic that was fastpathed and
- is supported only on the routed interfaces of Firewall Threat Defense devices.

QoS interface limitations

QoS is supported only on the routed interfaces of Firewall Threat Defense devices. However, it does not support site-to-site VPN and VTI interfaces.

Logging Rate-Limited Connections

You cannot configure logging for QoS. A connection can be rate limited without being logged. Conversely, a logged connection may not be rate-limited. To view QoS details, you must separately log connection ends in your Cloud-Delivered Firewall Management Center database.

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

QoS policies

A QoS policy is a configuration mechanism that

- governs rate limiting on managed devices
- can target multiple devices with each device having one deployed QoS policy at a time, and
- contains up to 63 QoS rules that handle network traffic in specified order.

QoS policy characteristics

QoS policies have these key characteristics:

-
- The system matches traffic to QoS rules in the order specified. The system rate limits traffic according to the first rule where all rule conditions match the traffic. Traffic not matching any of the rules is not rate limited.
- QoS rules can be limited by source or destination (routed) interfaces. The system enforces rate limiting *independently* on *each* of those interfaces and does not support an aggregate rate limit for a set of interfaces.
-
- You can independently rate limit download and upload traffic based on the connection initiator.

Note A device can have no more than 255 rules, including QoS rules. Exceeding this number causes a deployment warning necessitating a reduction in the number of rules..

Note QoS operates independently and is not subordinate to a main access control configuration. However, the access control and QoS policies deployed to the same device share identity configurations. For more infomration, see Associating other policies with access control.

Requirements and prerequisites for QoS

This reference provides the requirements and prerequisites for implementing QoS, including supported models, domains, and user roles.

Model support

Firewall Threat Defense

Supported domains

Any

User roles

Admin

Access Admin

Network Admin

Guidelines and limitations for QoS policies

QoS rules and service policy rules (in the access control policy) are translated into firewall commands. There can be only one policy map per interface as the system combines QoS and service policy configurations automatically.

However, you can also create policy maps manually using FlexConfig. For lower-end firewalls, such as the Secure Firewall 3100, you can create up to 64 policy maps using FlexConfig. For more powerful firewalls, such as the Firepower 4100, you can create up to 128 policy maps. Ensure not to exceed this limit to prevent deployment failures. You can check the device configuration using the CLI or the deployment transcript to see how rules are translated into commands.

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 **Policies** > **Network policies** > **QoS**.

Step 2 Click **New Policy** to create a new QoS policy and, optionally assign target devices; see Create a QoS policy, on page 4.

You can also **Copy** ( ) or **Edit** ( ) an existing policy.

Step 3 Configure QoS rules; see Configuring QoS rules, on page 6 and QoS rule conditions, on page 8.

The QoS policy editor lists each rule in evaluation order and displays a summary of the rule conditions and rate limiting configurations. Use the right-click menu to manage rules, including moving, enabling, and disabling them.

This feature is 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 rules, and the system matches text entered 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 but essential to effective deployment. Careful planning is necessary to prevent rules from preempting each other, requiring additional licenses, or containing invalid

configurations. Icons in the interface 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 which managed devices the policy targets; see Set target devices for a QoS policy , on page 4.

Verify your choices, if you identified target devices during policy creation.

Step 5 Save the QoS policy.

Step 6 Since this feature must allow some packets to pass, configure your system to examine those packets.

Step 7 Deploy configuration changes.

The QoS policy is configured and deployed to the target devices, enabling policy-based rate limiting for network traffic.

Create a QoS policy

Create a QoS policy to establish the foundation for implementing rate limiting and traffic management rules on network devices.

The new QoS policy you create will not perform any rate limiting unless rules are added.

Procedure

Step 1 Choose **Policies** > **Network policies** > **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.

Devices must be assigned before policy deployment.

Step 5 Click **Save**.

A new QoS policy is created and ready for configuration with rate limiting rules.

What to do next

- Configure and deploy the QoS policy; see Rate limiting with QoS policies, on page 3.

Set target devices for a QoS policy

Before you begin

Ensure you are logged into the QoS policy editor and have the necessary access rights to configure system policies.

A QoS policy can target multiple devices, but 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 either 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.

The QoS policy is successfully configured with assigned target devices and is ready for deployment.

What to do next

- Deploy configuration changes.

Filter QoS rules by device

You can filter the list of QoS policy rules based on the device to which the rules apply. A rule's interface constraints determine if the rule affects a device.

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 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.

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.

Step 3 Click **OK**.

A message appears next to the button showing the filter used. To remove the filter, click the X in the message. Alternatively, select **All** in the Filter by device dialog box.

The QoS rules list displays only the rules that apply to the selected devices or device groups. It simplifies the process of reviewing and managing device-specific QoS policies.

Configuring QoS rules

QoS rules enable you to control and rate limit network traffic on specific interfaces by configuring rule properties, conditions, and traffic limits.

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—Select 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 7.

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.

The QoS rule you created or modified is saved to the policy and configured with your set traffic limits and conditions.

What to do next

- Deploy configuration changes.

Related Topics

Best practices for access control rules

QoS rule components

QoS rules consist of several key components that control traffic rate limiting and filtering behavior. Understanding these components helps you configure effective quality of service policies.

QoS rules include these main components:

- State (Enabled/Disabled): Controls whether the rule is active and enforced by the system.
- Interfaces (Apply QoS On): Specifies which interfaces the rate limiting applies to.
- Traffic Limit Per Interface: Sets the rate limiting values for matching traffic.
- Conditions: Defines the traffic criteria that must be matched for the rule to apply.
- Comments: Provides documentation and change tracking for the rule.

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, apply QoS on one of the options:

- 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* interface specified with the Apply QoS On option. Aggregate rate limits cannot be applied to a group 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 independently rate limit download and upload traffic. The system determines download and upload directions based on the connection initiator.

If you specify a limit greater than the interface's maximum throughput, 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 that the rule handles. You can configure each rule with multiple conditions. Traffic must meet 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 8.

Comments

Each time you save changes to a rule, add comments. You might summarize the overall configuration for users, or note any changes made and the reasons.

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 traffic the rule handles. You can configure each rule with multiple conditions. All conditions must be met for traffic to match the rule. Each condition type has its own tab in the rule editor. You can rate limit traffic using these conditions.

Related Topics

Interface rule conditions, on page 8
Network rule conditions, on page 9
User rule conditions, on page 9
Application rule conditions, on page 9
Best practice: use application-based rules over port-based rules, on page 11
Recommendation: use URL conditions for website access control, on page 12
Custom SGT rule conditions, on page 12

Interface rule conditions

An interface rule condition is a network traffic control that

- controls traffic by its source and destination interfaces,
- uses predefined interface objects called security zones or interface groups to build interface conditions, and
- improves system performance by constraining rules to specific interfaces.

Interface rule condition characteristics

Interface objects segment your network by grouping interfaces across multiple devices to manage and classify traffic flow. For more information, refer to Interfaces.

QoS rules can be applied only on routed interfaces.

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.

Network rule conditions

A network condition is a traffic matching criterion that

- controls or decrypts traffic based on source and destination IP addresses using inner packet headers,
- enables administrators to specify exact IP addresses or blocks for granular policy enforcement, and
- can be built using predefined network objects or by manually entering IP addresses or address blocks.

Network rule condition restrictions

Minimize the number of matching criteria 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.

Note You *cannot* use FDQN network objects in identity rules.

User rule conditions

A user rule condition is a traffic matching mechanism that

- matches traffic based on whether the user initiating the connection or the group to which they belong
- can be configured for users in Microsoft Active Directory realms only, and
- enables policy enforcement based on user identity and group membership.

Special identities users

In addition to configuring users and groups for configured realms, you can set policies for the special identities users:

- Failed Authentication: User who 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 such as users who are not downloaded by a configured realm.

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.

User rule condition example

Configure a Block rule to prohibit anyone in the Finance group from accessing a network resource.

Application rule conditions

Application rule conditions are application awareness and control mechanisms that

- identify and classify applications through IP traffic analysis

- enable control with organized filters
- require one enabled detector for each condition.

Application filters and control benefits

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

At least one detector must be enabled for each application rule condition in the policy. If no detector is enabled, the system automatically enables all available system-provided detectors, or the most recently modified user-defined detector. For more information about application detectors, see Application detectors.

You can utilize both application filters and individually specified applications to ensure complete coverage. Understand the application's requirements to effectively order access control rules.

Application filters help you to quickly configure application control effectively. 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.

The system characterizes detected applications using several criteria which form effective 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 applications 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.
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 practice: use application-based rules over port-based rules

Use application filtering criteria instead of port filtering to target traffic. Applications can be configured to use unique ports to bypass access control blocks.

Port conditions allow you to control traffic by its source and destination ports.

Minimize the number of matching criteria 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.

The traditional method targets applications by using port specifications, but applications can manipulate unique ports to bypass these controls

Application filtering is recommended for applications to avoid blocking legitimate applications, like Firewall Threat Defense, that dynamically open separate channels for control and data flow. Port-based access control rules could block such applications, impacting desirable operations.

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 (concept)

A port, protocol, and ICMP code rule condition is a traffic matching criterion that

- matches traffic based on source and destination ports
- supports TCP, UDP, ICMP, and other protocols that do not use ports, and
- controls traffic flow in access control and other policy rules.

Port condition types

Port conditions match traffic based on the source and destination ports. Depending on the rule type, "port" can represent any of these types:

- TCP and UDP—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—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—Control traffic using other protocols that do not use ports.

Minimize the number of matching criteria 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 method used to target applications. However, applications can be configured to use unique ports to bypass access control blocks. Use application filtering criteria when possible; however, note that application filtering is unavailable in prefilter rules.

Use application filtering for FTP and similar applications that open separate channels for control and 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 configuring port conditions can match non-TCP traffic, restrictions apply:

- Access control rules—For Classic devices, 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, add only network-based conditions: zone, IP address, port, and VLAN tag. The system uses outer headers to match **all** traffic in access control policies with GRE-constrained rules. For Firewall Threat Defense devices, use tunnel rules in the prefilter policy to control GRE-encapsulated traffic.
- Decryption rules—These rules support only TCP port conditions.
- 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. To match any ICMP echo, use ICMP type 8 or ICMPv6 type 128.

Recommendation: use URL conditions for website access control

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

For complete information, see About URL Filtering with Category and Reputation.

Custom SGT rule conditions

A custom SGT rule condition is a traffic filtering mechanism that

- Use manually created SGT objects to filter traffic. Do not rely on ISE SGTs from a server.
- Correspond to the SGT attributes on the traffic to be controlled.
- Exclude user control considerations when using custom SGT traffic.

SGT traffic control without ISE configuration

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

ISE SGT versus custom SGT rule conditions

Certain rules allow you to control traffic based on assigned SGT. The rule type and identity source configuration determine whether ISE-assigned SGTs or custom SGTs are used to match traffic with assigned SGT attributes.

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

Note If using ISE SGTs to match traffic, any packet will match an ISE SGT rule if the SGT associated with the packet source IP address is recognized in ISE, even without an assigned SGT attribute.

Autotransition from custom Security Group Tags to ISE Security Group Tags

An autotransition from custom Security Group Tags (SGTs) to ISE Security Group Tags is a system process that automatically adjusts rule configuration when you configure ISE/ISE-PIC as an identity source after creating rules with custom Security Group Tags.

System behavior during autotransition

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

- Disable **Security Group Tag** options: The system retains existing SGT objects, but you cannot modify them or add new ones.
- Retain existing rules: Existing rules with custom SGT conditions remain, but they do not match traffic, and you cannot add more custom SGT conditions.

If you configure ISE, it is recommended to delete or disable existing rules with custom Security Group Tag conditions and use ISE attribute conditions to match traffic with Security Group Tag attributes.