Configuring Quality of Service

Cisco Security Manager supports the management and configuration of security services and other platform-specific services on Cisco Catalyst switches and Cisco 7600 Series routers.

You can manage Catalyst switches and 7600 devices configured in VTP transparent or VTP client/server mode. Security Manager manages switches configured in client/server mode by bypassing VLAN database management on the device (including VLAN creation, deletion, and monitoring VLANs in the VLAN database on switches).

This chapter contains the following topics

Quality of Service on Cisco IOS Routers

Note From version 4.17, though Cisco Security Manager continues to support IOS features/functionality, it does not support any bug fixes or enhancements.

Quality of service (QoS) refers to the ability of a network to provide priority service to selected network traffic over various underlying technologies, including Frame Relay, ATM, Ethernet and 802.1 networks, SONET, and IP-routed networks. QoS features enhance the predictability of network service by:

- Supporting dedicated bandwidth.
- Improving loss characteristics.
- Avoiding and managing network congestion.
- Shaping network traffic.
- Setting traffic priorities across the network.

QoS is generally used at entry points to service providers, as well as at consolidation points where multiple lines converge. QoS is also useful where speed mismatches occur (for example, at the boundary between a WAN and a LAN), as these places are often traffic congestion points.

QoS policies in Security Manager are based on the Cisco Systems Modular QoS CLI (MQC). MQC standardizes the CLI and semantics for QoS features across all platforms supported by Cisco IOS software, which provides

a modular and highly extensible framework for deploying QoS. Security Manager provides an easy-to-use interface for MQC that concentrates key QoS features inside a single dialog box, streamlining the creation of QoS policies for selected traffic entering and leaving the router.

For a description of the procedure for defining a QoS policy in Security Manager, see Defining QoS Policies , on page 11.

Related Topics

- Quality of Service and CEF , on page 2
- Understanding Marking Parameters , on page 3
- Understanding Queuing Parameters , on page 4
- Understanding Policing and Shaping Parameters , on page 7

Quality of Service and CEF

Cisco Express Forwarding (CEF) is an advanced Layer 3 IP switching technology that optimizes network performance and scalability for all kinds of networks. It defines the fastest method by which a Cisco IOS router forwards packets from ingress to egress interfaces.

Certain QoS features configurable in Security Manager, such as Class-Based Policing and Class-Based Weighted Random Early Detection, are supported only on routers that run CEF. All routers from the Cisco 800 Series to the Cisco 7200 Series require CEF for these QoS features; the Cisco 7500 Series requires distributed CEF (dCEF).

Note For a complete list, see *When is CEF Required for Quality of Service* on Cisco.com at this URL: http://www.cisco.com/en/US/tech/tk39/tk824/technologies_tech_note09186a0080094978.shtml

By default, CEF is enabled as part of the router's initial configuration. To verify whether CEF is enabled on your router, use the **show ip cef** command. You can configure CEF using the CEF interface settings policy (see CEF Interface Settings on Cisco IOS Routers). Be aware, however, that if your router does not have CEF enabled, activating CEF could have a significant impact on your router's packet streaming. Consult your router documentation before enabling CEF.

Related Topics

- Quality of Service on Cisco IOS Routers , on page 1

Understanding Matching Parameters

You define matching parameters by identifying the traffic on which QoS is performed, that is, classifying the interesting packets. Various classification tools are available, including protocol type, IP precedence (IPP) value, Differentiated Service Code Point (DSCP) value, and ACLs.

Traffic classes consist of a series of match criteria and a means of evaluating these criteria. For example, you might define a class with matching criteria based on several specified protocols and a DSCP value. You can

then specify that a packet must match only one of these defined criteria to be considered part of this class. Your other option is to specify that packets must match all defined criteria considered part of the traffic class.

Packets that are members of a defined traffic class are forwarded according to the QoS specifications that you defined in the policy map. Packets that fail to meet any of the matching criteria are classified as members of the default traffic class.

For information about defining matching parameters in a QoS policy, see Defining QoS Class Matching Parameters , on page 14.

Related Topics

- Defining QoS Policies , on page 11
- Quality of Service on Cisco IOS Routers , on page 1

Understanding Marking Parameters

Marking parameters enable you to classify packets, which entails using a traffic descriptor to categorize a packet within a specific group. This defines the packet and makes it accessible for QoS handling on the network. Both traffic policers and traffic shapers use the packet classification to ensure adherence to the contracted level of service agreed upon between the source and your network. Additionally, marking parameters enable you to take packets that might have arrived at the device with one QoS classification and reclassify them. Downstream devices use this new classification to identify the packets and apply the appropriate QoS functions to them.

Security Manager uses two types of marking for IPv4 packets—one based on IPP classes and one based on DSCP values. IPP is based on the three most significant bits in the Type of Service (ToS) byte of each packet, which means you can partition traffic into eight classes. For historical reasons, each precedence value corresponds with a name, as defined in RFC 791. Table 1: IP Precedence Classes , on page 3 lists the numbers and their corresponding names, from least to most important.

Table 1: IP Precedence Classes

Class	Name
0	routine
1	priority
2	immediate
3	flash
4	flash-override
5	critical
6	internet
7	network

Note Classes 6 and 7 are generally reserved for network control information, such as routing updates.

DSCP is based on the six most significant bits in the ToS byte (the remaining two bits are used for flow control), with values ranging from 0 to 63. The DSCP bits contains the IPP bits, which makes DSCP backward-compatible with IPP.

Marking is generally used on devices that are close to the network edge or administrative domain so that subsequent devices can provide service based on the classification mark.

For information about defining marking parameters in a QoS policy, see Defining QoS Class Marking Parameters , on page 16.

Related Topics

- Understanding Queuing Parameters , on page 4
- Understanding Policing and Shaping Parameters , on page 7
- Defining QoS Policies , on page 11
- Quality of Service on Cisco IOS Routers , on page 1

Understanding Queuing Parameters

Queuing manages congestion on traffic leaving a Cisco IOS router by determining the order in which to send packets out over an interface, based on priorities you assign to those packets. Queuing makes it possible to prioritize traffic to satisfy time-critical applications, such as desktop video conferencing, while still addressing the needs of less time-dependent applications, such as file transfer.

During periods of light traffic, that is, when no congestion exists, packets are sent out as soon as they arrive at an interface. However, during periods of transmission congestion at the outgoing interface, packets arrive faster than the interface can send them. By using congestion management features such as queuing, packets accumulating at the interface are queued until the interface is free to send them. They are then scheduled for transmission according to their assigned priority and the queuing mechanism configured for the interface. The router determines the order of packet transmission by controlling which packets are placed in which queue and how queues are serviced with respect to one another.

Security Manager uses a form of queuing called Class-Based Weighted Fair Queuing (CBWFQ). With CBWFQ, you define traffic classes based on match criteria. Packets matching the criteria constitute the traffic for this class. A queue is reserved for each class, containing the traffic belonging to that class. You assign characteristics to queues, such as the bandwidth (fixed or minimum) assigned to it and the queue limit, which is the maximum number of packets allowed to accumulate in the queue.

When you use CBWFQ, the sum of all bandwidth allocation on an interface cannot exceed 75 percent of the total available interface bandwidth. The remaining 25 percent is used for other overhead, including Layer 2 overhead, routing traffic, and best-effort traffic. Bandwidth for the CBWFQ default class, for instance, is taken from the remaining 25 percent.

For more information about queuing, see:

- Tail Drop vs. WRED , on page 5
- Low-Latency Queuing , on page 6

- Default Class Queuing , on page 6

For information about defining queuing parameters in a QoS policy, see Defining QoS Class Queuing Parameters , on page 16.

Related Topics

- Understanding Marking Parameters , on page 3
- Understanding Policing and Shaping Parameters , on page 7
- Defining QoS Policies , on page 11
- Quality of Service on Cisco IOS Routers , on page 1

Tail Drop vs. WRED

After a queue reaches its configured queue limit, the arrival of additional packets causes tail drop or packet drop to take effect, depending on how you configured the QoS policy. Tail drop, which is the default response, treats all traffic equally and does not differentiate between different classes of service. When tail drop is in effect, packets are dropped from full queues until the congestion is eliminated and the queue is no longer full. This often leads to global synchronization, in which a period of congestion is followed by a period of underutilization, as multiple TCP hosts reduce their transmission rates simultaneously.

A more sophisticated approach to managing queue congestion is offered by Cisco's implementation of Random Early Detection, called Weighted Random Early Detection, or WRED. As shown in Figure 1: Weighted Random Early Detection, on page 5, WRED reduces the chances of tail drop by selectively dropping packets when the output interface begins to show signs of congestion. By dropping some packets early instead of waiting until the queue is full, WRED avoids dropping large numbers of packets at once and allows the transmission line to be used fully at all times.

Figure 1: Weighted Random Early Detection

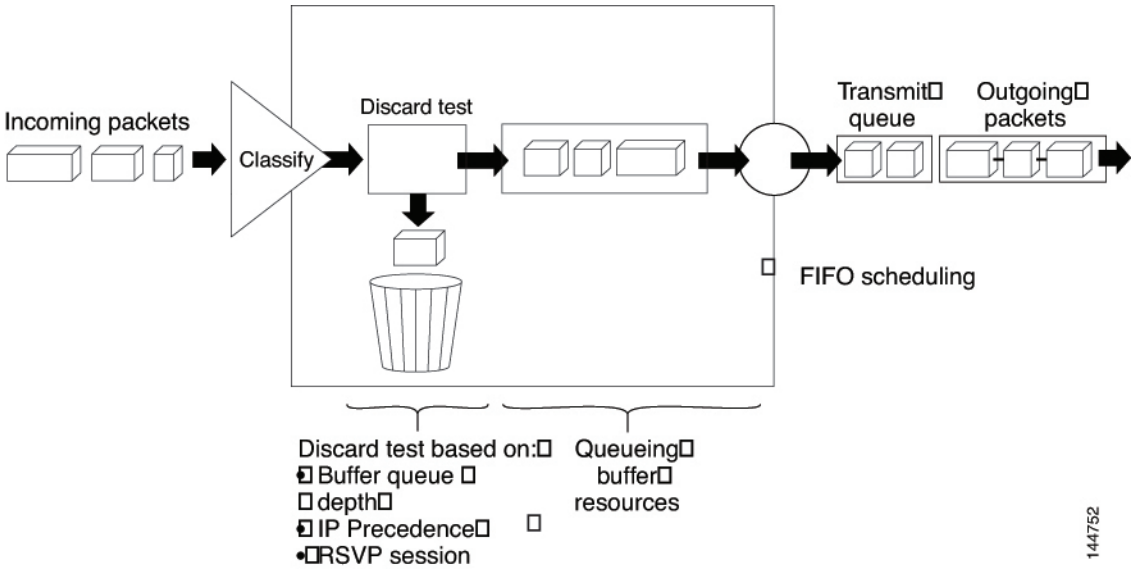

WRED is useful only when the bulk of the traffic is TCP/IP traffic, because TCP hosts reduce their transmission rate in response to congestion. With other protocols, packet sources might not respond, or might resend dropped packets at the same rate. As a result, dropping packets does not decrease congestion.

Note WRED treats non-IP traffic as precedence 0, the lowest precedence value. Therefore, non-IP traffic is more likely to be dropped than IP traffic.

Related Topics

- Low-Latency Queuing , on page 6
- Default Class Queuing , on page 6
- Understanding Queuing Parameters , on page 4

Low-Latency Queuing

The low-latency queuing (LLQ) feature brings strict priority queuing to CBWFQ. Strict priority queuing gives delay-sensitive data, such as voice traffic, preference over other traffic.

Note Although it is possible to assign various types of real-time traffic to the strict priority queue, we strongly recommend that you direct only voice traffic to it.

LLQ defines the maximum bandwidth that you can allocate to priority traffic during times of congestion. Setting a maximum ensures that nonpriority traffic does not starve (meaning that this traffic is also provided with bandwidth). When the device is not congested, the priority class traffic is allowed to exceed its allocated bandwidth. Policing drops packets from the priority queue; therefore, neither WRED nor tail drop (as configured in the Queue Limit field) is used.

When LLQ is not used, CBWFQ provides weighted fair queuing based on defined classes, with no strict priority queue available for real-time traffic.

Related Topics

- Tail Drop vs. WRED , on page 5
- Default Class Queuing , on page 6
- Understanding Queuing Parameters , on page 4

Default Class Queuing

You use the Fair Queue field to define the number of dynamic queues that should be reserved for the default class to use. This is the class to which traffic that does not satisfy the match criteria of other classes is directed. By default, the number of queues that are created is based on the interface bandwidth.

Table 2: Default Number of Queues for Default Class , on page 7 lists the default number of dynamic queues that CBWFQ uses when it is enabled on an interface:

Table 2: Default Number of Queues for Default Class

Bandwidth Range	Number of Dynamic Queues
Less than or equal to 64 kbps	16
More than 64 kbps and less than or equal to 128 kbps	32
More than 128 kbps and less than or equal to 256 kbps	64
More than 256 kbps and less than or equal to 512 kbps	128
More than 512 kbps	256

Related Topics

- Tail Drop vs. WRED , on page 5
- Default Class Queuing , on page 6
- Understanding Queuing Parameters , on page 4

Understanding Policing and Shaping Parameters

Security Manager offers two kinds of traffic regulation mechanisms:

- The rate-limiting feature of Class-Based Policing for policing traffic. Policing limits traffic flow to a configured rate. Policing can be performed on a selected interface or on the control plane. See Understanding Control Plane Policing , on page 10.
- Distributed Traffic Shaping (DTS) for shaping traffic. Traffic shaping enables you to control the traffic leaving an interface (output traffic) in order to match its flow to the speed of the remote target interface and to ensure that the traffic conforms to the policies defined for it. By shaping traffic to meet downstream requirements, you can eliminate bottlenecks in topologies with data-rate mismatches. Shaping can either be performed on selected QoS classes or at the interface level (hierarchical shaping).

Both policing and shaping mechanisms use the traffic descriptor for a packet—indicated by the classification of the packet (see Understanding Marking Parameters , on page 3)—to ensure adherence to the agreed upon level of service. Although policers and shapers usually identify traffic descriptor violations in the same way, they differ in the way they respond to violations, as shown in Figure 2: Traffic Policing vs. Traffic Shaping, on page 8:

- A policer typically drops excess traffic. In other cases, it transmits the traffic with a different (usually lower) priority.
- A shaper typically delays excess traffic using a buffer, or queuing mechanism, to hold packets and shape the flow when the data rate of the source is later than expected.

Figure 2: Traffic Policing vs. Traffic Shaping

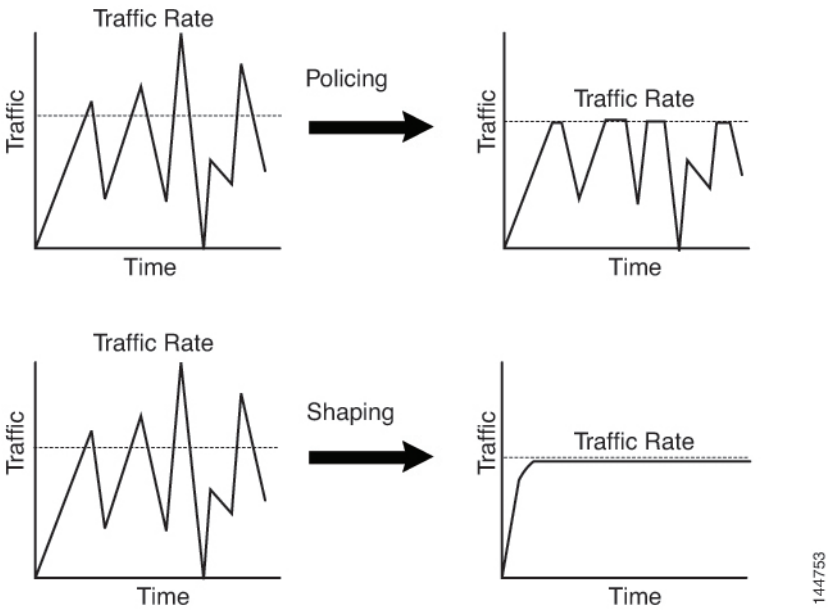

For information about defining policing and shaping parameters in a QoS policy, see Defining QoS Class Policing Parameters , on page 18 and Defining QoS Class Shaping Parameters , on page 19.

Related Topics

- Understanding the Token-Bucket Mechanism , on page 8
- Understanding Marking Parameters , on page 3
- Understanding Queuing Parameters , on page 4
- Defining QoS Policies , on page 11
- Quality of Service on Cisco IOS Routers , on page 1

Understanding the Token-Bucket Mechanism

Both policing and shaping use a token-bucket mechanism to regulate data flow. A token bucket is a formal definition of a rate of transfer. It has three components: a burst size, a mean rate, and a time interval (Tc). Any two values may be derived from the third using this formula:

mean rate = burst size / time interval

These terms are defined as follows:

- Mean rate—Also called the committed information rate (CIR), it specifies how much data can be sent or forwarded per unit time on average. The CIR is defined either as an absolute value or as a percentage of the available bandwidth on the interface. When defined as a percentage, the equivalent value in bits per second (bps) is calculated after deployment based on the interface bandwidth and the percent value defined in the policy.

Note If the interface bandwidth changes (for example, more bandwidth is added), the bps value of the CIR is recalculated based on the revised amount of bandwidth.

- Burst size—Also called the committed burst (Bc) size, it specifies for each burst how much data can be sent within a given time without creating scheduling concerns. When you use percentages to calculate the CIR, burst size is measured in milliseconds.
- Time interval—Also called the measurement interval, it specifies the amount of time in seconds per burst. Over any integral multiple of this interval, the bit rate of the interface does not exceed the mean rate. The bit rate, however, might be arbitrarily fast within the interval.

In the token-bucket metaphor, tokens are put into the bucket at a certain rate. These tokens represent permission for the source to send a certain number of bits into the network. To send a packet, the regulator (policer or shaper) must remove a number of tokens from the bucket that equals the packet size.

Security Manager uses a two-bucket algorithm, as shown in Figure 3: Two-Token Bucket Algorithm, on page 9. The first bucket is the conform bucket and the second bucket is the exceed bucket. The full size of the conform bucket is the number of bytes specified as the normal burst size. The full size of the exceed bucket is the number of bytes specified in the maximum burst size. Both buckets are initially full, and they are updated based on the token arrival rate, which is determined by the CIR. If the number of bytes in the arriving packet is less than the number of bytes in the conform bucket, the packet conforms. The required number of tokens are removed from the conform bucket and the defined conform action is taken (for example, the packet is transmitted). The exceed bucket is unaffected.

If the conform bucket does not contain sufficient tokens, the excess token bucket is checked against the number of bytes in the packet. If enough tokens are present in the two buckets combined, the exceed action is taken on the packet and the required number of bytes are removed from each bucket. If the exceed bucket contains an insufficient number of bytes, the packet is in violation of the burst limits and the violate action is taken on the packet.

Figure 3: Two-Token Bucket Algorithm

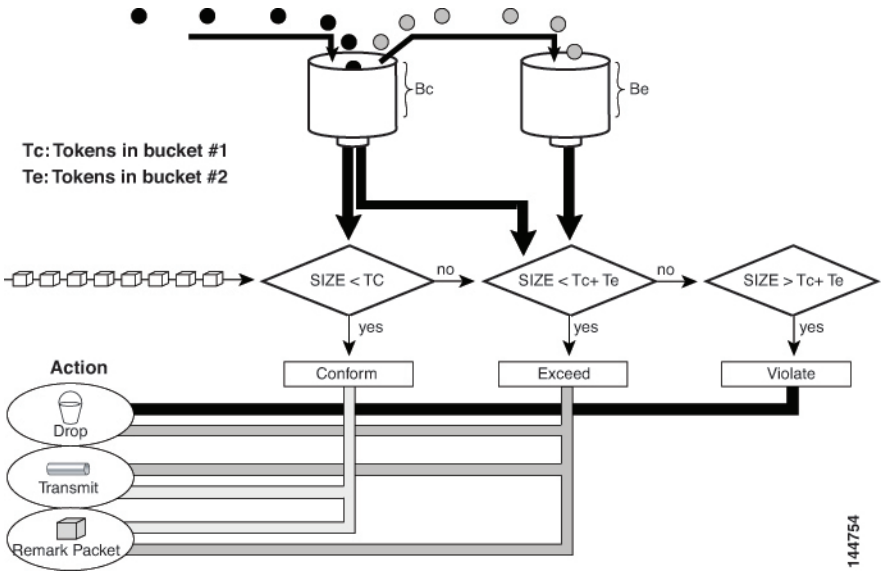

When you use traffic policing, the token-bucket algorithm provides three actions for each packet: a conform action, an exceed action, and an optional violate action. For instance, packets that conform can be configured to be transmitted, packets that exceed can be configured to be sent with a decreased priority, and packets that violate can be configured to be dropped.

Traffic policing is often configured on interfaces at the edge of a network to limit the rate of traffic entering or leaving the network. In the most common traffic policing configurations, traffic that conforms is transmitted and traffic that exceeds is sent with a decreased priority or is dropped. You can change these configuration options to suit your network needs.

When you use traffic shaping, the token-bucket mechanism includes a data buffer for holding packets that cannot be sent immediately. (Policers do not have such a buffer.) The token buckets permit packets to be sent in bursts, but places bounds on this capability so that the flow is never faster than the capacity of the buckets plus the time interval multiplied by the refill rate. The buffer also guarantees that the long-term transmission rate does not exceed the CIR.

Related Topics

- Understanding Control Plane Policing , on page 10
- Understanding Policing and Shaping Parameters , on page 7

Understanding Control Plane Policing

The Control Plane Policing feature enables you to manage input traffic entering the control plane (CP) of the router. The CP is a collection of processes that run at the process level on the route processor. These processes collectively provide high-level control for most Cisco IOS functions. Control plane policing protects the CP of Cisco IOS routers and switches against reconnaissance and denial-of-service (DoS) attacks, enabling the CP to maintain packet forwarding and protocol states despite an attack or heavy traffic load on the router or switch.

The Control Plane Policing feature treats the CP as a separate entity with its own ingress (input) and egress (output) ports, enabling you to use Security Manager to configure QoS policies on input. These policies are applied when a packet enters the CP. You can configure a QoS policy to prevent unwanted packets from progressing after a specified rate limit is reached. For example, a system administrator can limit all TCP/SYN packets that are destined for the CP to a maximum rate of 1 megabit per second. Additional packets beyond this limit are silently discarded.

The following types of Layer 3 packets are forwarded to the CP and processed by aggregate control plane policing:

- Routing protocol control packets
- Packets destined for the local IP address of the router
- Packets from management protocols, such as SNMP, Telnet, and secure shell (SSH).

Note Support for output policing is available only in Cisco IOS Release 12.3(4)T and later T-train releases.

For information about how to define Control Plane Policing, see Defining QoS on the Control Plane , on page 13. For more information about this feature, refer to the document, *Control Plane Policing* on Cisco.com at this URL:

http://www.cisco.com/en/US/docs/ios/qos/configuration/guide/ctrl_plane_policng.html

Related Topics

- Understanding the Token-Bucket Mechanism , on page 8
- Understanding Policing and Shaping Parameters , on page 7

Defining QoS Policies

When you define QoS policies, you must first decide whether to configure the policy on specific interfaces or on the control plane. This initial choice determines how you configure the rest of the policy, as described in the following topics:

- Defining QoS on Interfaces , on page 11
- Defining QoS on the Control Plane , on page 13

Note If you define a QoS policy on both the interfaces and the control plane of the same device, only the control plane configuration is deployed.

Related Topics

- Quality of Service on Cisco IOS Routers , on page 1

Defining QoS on Interfaces

You can create multiple QoS interface definitions, each of which applies to either input traffic (entering the router) or output traffic (exiting the router).

When you create a QoS interface definition on output traffic, you have the option of configuring hierarchical shaping on the interface as a whole instead of configuring shaping on individual QoS classes.

After you create your interface definitions, you must define one or more QoS classes on each interface. QoS classes contain the matching criteria that determine which packets are included in the class and the QoS functions (marking, queuing, policing, and shaping) to apply to that traffic. You can configure each interface (or interface role) with up to 16 QoS classes, each containing its own set of matching criteria and a defined set of QoS functions to apply to the traffic in that class.

For each interface, we recommend that for each interface you define at least one QoS class and a default class. If you do not configure a default class, packets that do not match the criteria of the other defined classes are treated as members of a default class that has no configured QoS functionality. Packets assigned to this class are placed in a simple first-in first-out (FIFO) queue, and are forwarded at a rate determined by the available underlying link bandwidth. This FIFO queue is managed by tail drop, which avoids congestion by dropping packets from the queue until it is no longer full.

Note QoS is applied to packets on a first-match basis. The router examines the table of QoS classes starting from the top and applies the properties of the first class whose matching criteria matches the packet. Therefore, it is important that you define and order your classes carefully. The default class should be placed last to prevent traffic that matches a specific class from being treated as unmatched traffic.

Before You Begin

Ensure that Cisco Express Forwarding (CEF) is enabled on the router. For more information, see CEF Interface Settings on Cisco IOS Routers.

Related Topics

- Defining QoS Policies , on page 11
- Defining QoS on the Control Plane , on page 13
- Quality of Service on Cisco IOS Routers , on page 1

Step 1 Do one of the following:

- (Device view) Select **Platform > Quality of Service** from the Policy selector.
- (Policy view) Select **Router Platform > Quality of Service** from the Policy Type selector. Select an existing policy or create a new one.

The Quality of Service page is displayed. See <Table 3: Quality of Service Page , on page 21 for a description of the fields on this page.

Step 2 In the Applied to field, select **Interfaces** to define QoS parameters for specific interfaces on the selected router.

Step 3 Click the **Add** button under the upper table to display the QoS Policy dialog box. See Table 4: QoS Policy Dialog Box , on page 22 for a description of the fields in this dialog box.

Step 4 In the Interface field, enter the name of an interface or interface role, or click **Select** to display a selector.

Tip If the interface role you want is not listed in the selector, click the **Create** button or the **Edit** button to open the Interface Role Dialog Box. From here you can define an interface role to use in the policy.

Step 5 Select the traffic direction on which you want to apply the QoS definition, Output (traffic exiting the interface) or Input (traffic entering the interface). Queuing and shaping can be applied only to output traffic.

Step 6 (Optional) Define interface-level (hierarchical) shaping parameters. See Table 4: QoS Policy Dialog Box , on page 22 for details.

Note When you enable hierarchical shaping on an interface, you cannot define shaping parameters for specific QoS classes. Shaping can be used only on output traffic. See Understanding Policing and Shaping Parameters , on page 7 for more information about shaping.

Step 7 Click **OK**. The QoS interface definition is displayed in the upper table of the Quality of Service page.

Note To edit a QoS interface definition, select an interface from the upper table, then click the **Edit** button. To remove an interface definition, select it from the table, then click the **Delete** button. You cannot delete an interface that has defined classes.

Step 8 With the interface selected in the upper table, click the **Add** button beneath the QoS Classes table. The QoS Class dialog box is displayed. See Table 5: QoS Class Dialog Box , on page 25 for a description of the fields in this dialog box.

The QoS Class dialog box enables you to determine which traffic over the selected interface is included in the QoS class and how to handle that traffic.

Step 9 (Optional) Select the **Default class** check box if you are defining the properties of the default QoS class for this interface. The default class is assigned to all traffic that does not match the criteria of the other defined classes.

Step 10 Define the QoS class using one or more tabs in the QoS Class dialog box, as described in:

- Defining QoS Class Matching Parameters , on page 14
- Defining QoS Class Marking Parameters , on page 16
- Defining QoS Class Queuing Parameters , on page 16
- Defining QoS Class Policing Parameters , on page 18
- Defining QoS Class Shaping Parameters , on page 19

Step 11 Repeat Step 8, on page 13 through Step 10, on page 13 to add QoS classes to the interface defined in Step 3, on page 12. If required, use the **Up Row** and **Down Row** buttons to reorder the classes.

Note To edit a QoS class, select the relevant interface from the upper table to display its defined classes in the QoS Class table. Select the class to edit, then click the **Edit** button. To remove a class, select it from the table, then click the **Delete** button.

Step 12 Repeat Step 3, on page 12 through Step 11, on page 13 to define QoS classes for a different interface on the selected router.

Defining QoS on the Control Plane

When you configure QoS on input traffic entering the control plane, you can define multiple QoS classes, including a default class for traffic that does not match the criteria you define for the other classes. After defining the matching criteria for a particular class, you can configure a policing definition for that class. (Marking, queuing, and shaping are not available.) For more information, see Understanding Control Plane Policing , on page 10.

QoS policies defined on the control plane override any QoS parameters defined on an interface of the same device.

Note QoS is applied to packets on a first-match basis. The router examines the table of QoS classes starting from the top and applies the properties of the first class whose matching criteria matches the packet. Therefore, it is important that you define and order your classes carefully. The default class should be placed last to prevent traffic that matches a specific class from being treated as unmatched traffic.

Before You Begin

Ensure that Cisco Express Forwarding (CEF) is enabled on the router. For more information, see CEF Interface Settings on Cisco IOS Routers.

Related Topics

- Defining QoS Policies , on page 11
- Defining QoS on Interfaces , on page 11
- Quality of Service on Cisco IOS Routers , on page 1

Step 1 Do one of the following:

- (Device view) Select **Platform > Quality of Service** from the Policy selector.
- (Policy view) Select **Router Platform > Quality of Service** from the Policy Type selector. Select an existing policy or create a new one.

The Quality of Service page is displayed. See Table 3: Quality of Service Page , on page 21 for a description of the fields on this page.

Step 2 In the Applied to field, select **Control Plane** to define QoS policing on input traffic entering the control plane.

Step 3 Click the **Add** button beneath the Control Plane QoS Classes table. The QoS Class dialog box is displayed. See Table 5: QoS Class Dialog Box , on page 25 for a description of the fields in this dialog box.

The QoS Class dialog box enables you to determine which traffic over the selected interface is included in the QoS class and how to handle that traffic.

Step 4 (Optional) Select the **Default class** check box if you are defining the properties of the default QoS class for the control plane. The default class is assigned to all traffic that does not match the criteria of the other defined classes.

Step 5 Define the QoS class using the tabs in the QoS Class dialog box, as described in the following sections:

- Defining QoS Class Matching Parameters , on page 14
- Defining QoS Class Policing Parameters , on page 18

Step 6 Repeat Step 3, on page 14 through Step 5, on page 14 to add QoS classes to the control plane. If required, use the **Up Row** and **Down Row** buttons to reorder the classes.

Defining QoS Class Matching Parameters

When you define matching parameters, you must define matching criteria and specify whether packets must meet one or all of the criteria to be considered part of the class. See Understanding Matching Parameters , on page 2 for more information.

Note You do not define matching parameters when configuring the default class.

Related Topics

- Defining QoS Class Marking Parameters , on page 16
- Defining QoS Class Queuing Parameters , on page 16
- Defining QoS Class Policing Parameters , on page 18
- Defining QoS Class Shaping Parameters , on page 19

- Defining QoS Policies , on page 11
- Quality of Service on Cisco IOS Routers , on page 1

Step 1 On the Quality of Service page, click the **Add** button beneath the QoS Classes table, or select a class and then click the **Edit** button. The QoS Class dialog box is displayed.

Step 2 Click the **Matching** tab. SeeTable 5: QoS Class Dialog Box , on page 25 for a description of the fields on this tab.

Step 3 Select a matching method:

- Any—Traffic matching any of the defined parameters is included in this class.
- All—Only traffic matching all of the defined parameters is included in this class.

Step 4 (Optional) Under Protocol, click **Add** to display a selector for choosing the protocols to include in this class. Select one or more items from the Available Protocols list, then click >> to add them to the Selected Protocols list.

Note When configuring QoS on the control plane, only the ARP protocol can be selected.

When you finish, click **OK** to save your definitions and return to the QoS Class dialog box. Your selections are displayed in the Protocol field.

Step 5 (Optional) Under Precedence, click **Add** to display a selector for choosing which IP precedence values (from 0 to 7) to include in this class. Select one or more items from the Available Precedences list, then click >> to add them to the Selected Precedences list. Traffic that arrives marked with one of these values matches this criterion.

Note For more information about IP precedence values, see Table 1: IP Precedence Classes , on page 3.

When you finish, click **OK** to save your definitions and return to the QoS Class dialog box. Your selections are displayed in the Precedences field.

Step 6 (Optional) Under DSCP, click **Add** to display a selector for choosing which DSCP values (from 0 to 63) to include in this class. Select one or more items (up to eight) from the Available DSCPs list, then click >> to add them to the Selected DSCPs list. Traffic that arrives marked with one of these values matches this criterion.

When you finish, click **OK** to save your definitions and return to the QoS Class dialog box. Your selections are displayed in the DSCP field.

Step 7 (Optional) Under ACL, define ACLs as part of the matching criteria for this class:

a) Click **Edit** to display the Edit ACLs dialog box. Use this dialog box to define which ACLs to include in this class.
b) Enter one or more ACLs, or click **Select** to select an ACL object from a list or to create a new one. Traffic that matches these ACL definitions matches this criterion.
c) When you finish, click **OK** twice to save your definitions and return to the QoS Class dialog box. Your selections are displayed in the ACL field.

Tip Use the up and down arrows to order the ACLs. We recommend placing more frequently used ACLs at the top of the list to optimize the matching process.

Step 8 Go to another tab or click **OK** to save your definitions locally on the client and close the dialog box. The defined class is displayed in the QoS Classes table on the Quality of Service page.

Step 9 Do one of the following:

- When defining QoS on interfaces, continue as described in Defining QoS on Interfaces , on page 11.

- When defining control plane policing, continue as described in Defining QoS on the Control Plane , on page 13.

Defining QoS Class Marking Parameters

When you define marking parameters, you can mark the packets in this QoS class with either a precedence value or a DSCP value. See Understanding Marking Parameters , on page 3 for more information.

Note Marking is not available when you configure QoS on the control plane.

Related Topics

- Defining QoS Class Matching Parameters , on page 14
- Defining QoS Class Queuing Parameters , on page 16
- Defining QoS Class Policing Parameters , on page 18
- Defining QoS Class Shaping Parameters , on page 19
- Defining QoS Policies , on page 11
- Quality of Service on Cisco IOS Routers , on page 1

Step 1 On the Quality of Service page, click the **Add** button beneath the QoS Classes table, or select a class and then click the **Edit** button. The QoS Class dialog box is displayed.

Step 2 Click the **Marking** tab. See Table 7: QoS Class Dialog Box—Marking Tab , on page 28 for a description of the fields on this tab.

Step 3 Select the **Enable Marking** check box.

Step 4 Select one of the following marking options:

- Precedence—Select an IP precedence value (0 to 7) from the displayed list. For more information about these values, see Table 1: IP Precedence Classes , on page 3.
- DSCP—Select a DSCP value (0 to 63) from the displayed list.

Step 5 Go to another tab or click **OK** to save your definitions locally on the client and close the dialog box. The defined class is displayed in the QoS Classes table on the Quality of Service page.

Step 6 Continue as described in Defining QoS Policies , on page 11.

Defining QoS Class Queuing Parameters

When you define queuing parameters, you can specify the amount of available bandwidth to provide to the traffic in this QoS class. You can also define a fixed amount of bandwidth that must be provided to high-priority traffic; you can define the priority parameter on only one class per interface. In addition, you must specify the type of queue management to perform on this class. See Understanding Queuing Parameters , on page 4 for more information.

Note Queuing is not available when you configure QoS on the control plane.

Related Topics

- Defining QoS Class Matching Parameters , on page 14
- Defining QoS Class Marking Parameters , on page 16
- Defining QoS Class Policing Parameters , on page 18
- Defining QoS Class Shaping Parameters , on page 19
- Defining QoS Policies , on page 11
- Quality of Service on Cisco IOS Routers , on page 1

Step 1 On the Quality of Service page, click the **Add** button beneath the QoS Classes table, or select a class and then click the **Edit** button. The QoS Class dialog box is displayed.

Step 2 Click the **Queuing and Congestion Avoidance** tab. See Table 8: QoS Class Dialog Box—Queuing and Congestion Avoidance Tab , on page 29 for a description of the fields on this tab.

Step 3 Click the **Enable Queuing and Congestion Avoidance** check box.

Queuing options depend on whether you are defining the default class or a different class:

- When you define any class other than the default class, select one of the following queuing options:
  - Priority—Define the amount of bandwidth to make available to high-priority traffic. Low-Latency Queuing , on page 6 (LLQ) ensures that this traffic receives this fixed amount of bandwidth at all times. This is particularly useful for voice traffic, which requires low latency. You can define this amount by percentage or by an absolute value of kilobits per second.

Note You can define this option for only one class per interface.

  - Bandwidth—Enter the amount of bandwidth to allocate to this class. You can define this amount by percentage or by an absolute value of kilobits per second.

Note The sum of all class bandwidth allocations on an interface cannot exceed 100 percent of the total available bandwidth.

- When you define the default class, select one of the following queuing options:
  - Fair queue—Enter the number of queues to reserve for the default class. Values range in powers of 2 from 16 to 4096. By default, the number of queues is based on the available bandwidth of the selected interface. For more information, see Table 2: Default Number of Queues for Default Class , on page 7.
  - Bandwidth—Enter the amount of bandwidth to allocate to this class. You can define this amount by percentage or by an absolute value of kilobits per second.

Step 4 (Optional) Define *one* of the following queue length management options:

- Queue Limit—(Default) Specify the maximum number of packets allowed. If you select this option, tail drop drops excess packets when the queue reaches its capacity.
- WRED Weight for Mean Queue Depth—WRED proactively drops packets until the transmitting protocol (usually TCP) responds by dropping its transmission rate, thereby alleviating congestion. Configure WRED by entering an exponential weight factor that is used to calculate the average queue size.

For more information, see Tail Drop vs. WRED , on page 5.

Note You should change the default only if you are certain that your applications will benefit from a different value.

Note Do not use WRED with protocols that are not sufficiently robust to reduce their transmission rates in response to packet loss, such as IPX or AppleTalk. WRED cannot be configured when you select the Priority percent option.

Step 5 Go to another tab or click **OK** to save your definitions locally on the client and close the dialog box. The defined class is displayed in the QoS Classes table on the Quality of Service page.

Step 6 Continue as described in Defining QoS Policies , on page 11.

Defining QoS Class Policing Parameters

When you define policing parameters, you must specify the average data rate, which determines the amount of traffic that can be transmitted. In addition, you must specify the action to take on traffic bursts that exceed this data rate.

You can configure policing for all QoS classes, including the default class. For more information about policing, see Understanding Policing and Shaping Parameters , on page 7.

You can also configure policing on the control plane. For more information, see Understanding Control Plane Policing , on page 10.

Related Topics

- Defining QoS Class Matching Parameters , on page 14
- Defining QoS Class Marking Parameters , on page 16
- Defining QoS Class Queuing Parameters , on page 16
- Defining QoS Class Shaping Parameters , on page 19
- Defining QoS Policies , on page 11
- Quality of Service on Cisco IOS Routers , on page 1

Step 1 On the Quality of Service page, click the **Add** button beneath the QoS Classes table, or select a class and then click the **Edit** button. The QoS Class dialog box is displayed.

Step 2 Click the **Policing** tab. See Table 5: QoS Class Dialog Box , on page 25 for a description of the fields on this tab.

Step 3 Select the **Enable Policing** check box.

Step 4 Define CIR, confirm burst, and excess burst values. You can define the CIR by percentage or by an absolute value of bits per second. The option you choose determines how you define the burst values.

Step 5 Select the action to perform on packets that conform to the rate limit:

- transmit—Transmit the packet.
- set-prec-transmit—Set the IP precedence to a defined value, then send the packet. This option is not available when configuring QoS on the control plane.
- set-dscp-transmit—Set the DSCP to a defined value, then send the packet. This option is not available when configuring QoS on the control plane.
- drop—Drop the packet.

Step 6 Select the action to perform on exceed packets. The list of available actions depends on the selected conform action.

For example, if transmit is performed on conforming packets, you can select any of the actions listed in Step 5, on page 18 for exceeding packets. However, if you selected one of the set actions for conforming packets, you can select only a set action or the drop action for exceeding packets. If you selected drop as the conform action, you must select drop as the exceed action.

Step 7 Select the action to perform on violate packets. The list of available actions depends on the selected exceed action.

For example, if transmit is performed on exceeding packets, you can select any of the actions listed in Step 5, on page 18 for violating packets. However, if you selected one of the set actions for exceeding packets, you can select only a set action or the drop action for violating packets. If you selected drop as the exceed action, you must select drop as the violate action.

Step 8 Go to another tab, or click **OK** to save your definitions locally on the client and close the dialog box. The defined class is displayed in the QoS Classes table on the Quality of Service page.

Step 9 Do one of the following:

- When defining QoS on interfaces, continue as described in Defining QoS Policies , on page 11.
- When defining control plane policing, continue as described in Defining QoS on the Control Plane , on page 13.

Defining QoS Class Shaping Parameters

When you define shaping parameters, you must specify whether to base traffic shaping on the average data rate or on the average data rate plus the excess burst rate that occurs during traffic peaks. In both cases, traffic that exceeds these definitions is buffered until the rate lowers, allowing the packets to be sent.

The following conditions pertain:

- Shaping can be used only on output traffic.
- Shaping can be configured for all QoS classes, including the default class.
- Shaping is not available when you configure the QoS class for priority traffic.
- Shaping is not available when you configure QoS on the control plane.

For more information about shaping, see Understanding Policing and Shaping Parameters , on page 7.

Tip To configure shaping on all the QoS classes defined for the interface (hierarchical shaping), see Defining QoS on Interfaces , on page 11.

Related Topics

- Defining QoS Class Matching Parameters , on page 14
- Defining QoS Class Marking Parameters , on page 16
- Defining QoS Class Queuing Parameters , on page 16
- Defining QoS Class Policing Parameters , on page 18
- Defining QoS Policies , on page 11
- Quality of Service on Cisco IOS Routers , on page 1

Step 1 On the Quality of Service page, click the **Add** button beneath the QoS Classes table, or select a class and then click the **Edit** button. The QoS Class dialog box is displayed.

Step 2 Click the **Shaping** tab. See Table 10: QoS Class Dialog Box—ShapingTab, on page 33 for a description of the fields on this tab.

Step 3 Select the **Enable Shaping** check box.

Step 4 Select the shaping type (Average or Peak).

Step 5 Define CIR, sustained burst, and excess burst values. You can define the CIR by percentage or by an absolute value of bits per second. The option you choose determines how you define the burst values.

Step 6 Proceed to another tab or click **OK** to save your definitions locally on the client and close the dialog box. The defined class is displayed in the QoS Classes table on the Quality of Service page.

Step 7 Continue as described in Defining QoS Policies , on page 11.

Quality of Service Policy Page

Use the Quality of Service page to view, create, and edit QoS classes on specific interfaces of the selected device or on the control plane. QoS policies enable you to define techniques for managing the delay, delay variation (jitter), bandwidth, and packet loss parameters on a network. In addition, you can use the Quality of Service page to configure hierarchical shaping on an interface as an alternative to configuring shaping parameters for individual QoS classes.

For more information, see Quality of Service on Cisco IOS Routers , on page 1.

Navigation Path

- (Device view) Select **Platform > Quality of Service** from the Policy selector.
- (Policy view) Select **Router Platform > Quality of Service** from the Policy Type selector. Create a new policy or select an existing policy from the Shared Policy selector.

Related Topics

- Defining QoS Policies , on page 11
- Table Columns and Column Heading Features
- Filtering Tables

Field Reference

Table 3: Quality of Service Page

Element	Description
Apply To	The router component on which to define the QoS policy: • Interfaces—Configures QoS classes on specific interfaces. • Control Plane—Configures QoS on the router control plane. See Understanding Control Plane Policing , on page 10. **Note** If you configure QoS on both the interfaces and the control plane of the same device, only the control plane configuration is deployed.
Interface Table	If you are defining classes on interfaces, the upper table lists the interfaces on which you are defining QoS classes. The direction column indicates the direction of traffic through the interface to which the classes apply (Output or Input). The classes you can define vary based on the direction. The other fields indicate whether you defined shaping on the interface, and if shaping is defined, the type of hierarchical shaping (average or peak), the committed information rate (CIR), and the sustained and excess burst size. For detailed information about the attributes, see QoS Policy Dialog Box , on page 22 . • To add an interface to the table, click the Add button. • To edit the settings for an interface, select it and click the Edit button. • To delete an interface, select it and click the Delete button.
QoS Classes Table	The classes defined for the interface selected in the upper table, or for the control plane. Each row represents a separate class. The No. column indicates the order of the classes, and is very important: QoS is applied to packets on a first-match basis, based on class order. The Default Class column indicates whether this class is the default for all packets on the interface that do not match the criteria of the other defined classes. Make this the last class in the list. The remaining columns indicate the match criteria for the class, and the packet marking, queuing and congestion avoidance, policing, and shaping defined for the class, if any. For detailed information about the attributes, see QoS Policy Dialog Box , on page 22. • To add class to the table, click the Add button. • To edit the settings for a class, select it and click the Edit button. • To delete a class, select it and click the Delete button. • To change the order of a class, select it and click the Up and Down arrow buttons to reposition it.

QoS Policy Dialog Box

Use the QoS Policy dialog box to select an interface on which you want to define QoS parameters. In addition, you can use this dialog box to configure a single set of shaping parameters for all the traffic on the selected interface (known as hierarchical shaping). Using hierarchical shaping eliminates the need to configure shaping parameters for each QoS class defined on the interface.

Note This dialog box is not applicable when defining a QoS policy on the control plane. For more information, see Defining QoS on the Control Plane , on page 13.

After you create your QoS interface definitions, you can define one or more QoS classes for each interface. For more information, see QoS Class Dialog Box , on page 24.

Navigation Path

Go to the Quality of Service Policy Page , on page 20, then click the **Add** or **Edit** button beneath the upper table to define a QoS interface definition.

Related Topics

- Defining QoS Policies , on page 11
- Quality of Service on Cisco IOS Routers , on page 1
- Basic Interface Settings on Cisco IOS Routers
- Understanding Interface Role Objects

Field Reference

Table 4: QoS Policy Dialog Box

Element	Description
Interface	The interface on which QoS is defined. Enter the name of an interface or interface role, or click **Select** to select an object from a list or to create a new object.
Direction	The direction of the traffic on which to configure QoS: • Output—Traffic that exits the interface. • Input—Traffic that enters the interface.
Hierarchical Shaping settings	
Enable Shaping	When selected, configures hierarchical traffic shaping on the selected interface. When deselected, hierarchical shaping is not used. **Note** Shaping can be performed only on output traffic.

Element	Description
Type	The type of shaping to perform: • Average—Limits the data rate for each interval to the sustained burst rate (also known as the Committed Burst rate or Bc), achieving an average rate no higher than the committed information rate (CIR). Additional packets are buffered until they can be sent. • Peak—Limits the data rate for each interval to the sustained burst rate plus the excess burst rate (Be). Additional packets are buffered until they can be sent.
CIR	The average data rate (also known as the committed information rate or CIR). You can define this amount by: • Percentage—Valid values range from 0 to 100% of the overall available bandwidth. • Bit/sec—Valid values range from 8000 to 1000000000 bits per second, and must be in multiples of 8000. Although data bursts during an interval may exceed this rate, the average data rate over any multiple integral of the interval will not exceed this rate.
Sustained Burst	The normal burst size. If you select average as the shaping type, data bursts during an interval are limited to this value. The range of valid values is determined by the CIR: • When the CIR is defined by percentage—Valid values range from 10 to 2000 milliseconds. • When the CIR is defined by an absolute value—Valid values range from 1000 to 154400000 bytes, in multiples of 128 bytes. **Note** We recommend that you leave this field blank when the CIR is defined by an absolute value. This allows the algorithms used by the device to determine the optimal sustained burst value.
Excess Burst	The excess burst size. If you select peak as the shaping type, data bursts during an interval can equal the sum of the sustained burst value plus this value. The average data rate over multiple intervals, however, will continue to conform to the CIR. The range of valid values is determined by the CIR: • When the CIR is defined by percentage—Valid values range from 10 to 2000 milliseconds. • When the CIR is defined by an absolute value—Valid values range from 1000 to 154400000 bytes, in multiples of 128 bytes. **Note** If you do not configure this field when the CIR is defined by an absolute value, the sustained burst value is used.

QoS Class Dialog Box

Use the QoS Class dialog box to create or edit a QoS class on a selected interface or control plane of a Cisco IOS router. You can define up to 16 classes on a single interface and 256 classes for the device as a whole.

Note QoS is applied to packets on a first-match basis. The router examines the table of QoS classes starting from the top and applies the properties of the first class whose matching criteria matches the packet. Therefore, it is important that you define and order your classes carefully. The default class should be placed last to prevent traffic that matches a specific class from being treated as unmatched traffic.

Navigation Path

Go to the Quality of Service Policy Page , on page 20. Complete the options at the top of the page, then do one of the following:

- To create a QoS class, select an interface from the upper table, then click the **Add** button beneath the QoS Class table. When creating a QoS class for the control plane, just click the **Add** button beneath the table.
- To edit a QoS class:
  - Select the interface whose class you want to edit from the upper table (Not required when selecting the control plane.).
  - Select the relevant class defined for that interface in the QoS Classes table. (Not required when selecting the control plane.)
  - Click the **Edit** button under the QoS Class table.

Related Topics

- QoS Policy Dialog Box , on page 22
- Defining QoS Policies , on page 11
- Defining QoS on Interfaces , on page 11
- Defining QoS on the Control Plane , on page 13

Field Reference

Table 5: QoS Class Dialog Box

Element	Description
Set as Default Class	When selected, enables you to define the default class for all traffic that does not match the other QoS classes on this interface. When deselected, enables you to define a specific QoS class on this interface. **Note** When you define the default class, you do not configure any matching parameters; by definition the class consists of all traffic that does not match any of the other classes. Therefore, the Matching tab is disabled.
Matching tab	Defines the traffic that is included in this QoS class. See QoS Class Dialog Box—Matching Tab , on page 25.
Marking tab	Marks the traffic in this class so that downstream devices can properly identify it. See QoS Class Dialog Box—Marking Tab , on page 27.
Queuing and Congestion Avoidance tab	Defines how to queue the output traffic in this class. See QoS Class Dialog Box—Queuing and Congestion Avoidance Tab , on page 28.
Policing tab	Limits the traffic flow for this class to a configured rate. See QoS Class Dialog Box—Policing Tab , on page 30.
Shaping tab	Controls the flow of output traffic for this class so that it conforms with the requirements of downstream devices. See QoS Class Dialog Box—Shaping Tab , on page 32.

Note When you configure a QoS policy on the control plane, only the Matching tab and Policing tab are available.

QoS Class Dialog Box—Matching Tab

Use the Matching tab of the QoS Class dialog box to define which traffic over the selected interface is considered to be part of this class.

Note When you define the default class, the Matching tab is disabled.

Navigation Path

Go to the QoS Class Dialog Box , on page 24, then click the **Matching** tab.

Related Topics

- Defining QoS Class Matching Parameters , on page 14
- Defining QoS on Interfaces , on page 11

- Defining QoS on the Control Plane , on page 13
- Quality of Service Policy Page , on page 20
- Creating Access Control List Objects

Field Reference

Table 6: QoS Class Dialog Box—Matching Tab

Element	Description
Match Method	The traffic matching option used for this class: • Any—Assigns traffic matching any of the defined class map criteria to this QoS class. • All—Assigns only traffic matching all of the defined class map criteria to this QoS class.
Protocol	One or more protocols included in this class map. Click **Add** to display a selector. Select one or more items from the Available Protocols list, then click >> to add them to the Selected Protocols list. The only protocol available for the control plane is ARP; ARP and CDP are not available for input classes configured on an interface. When you finish, click **OK** to return to the QoS Class dialog box. Your selections are displayed in the Protocol field. **Note** To remove a protocol from the QoS class, select it from the Protocol field, then click **Delete**.
Precedence	One or more IP Precedence (IPP) values included in this class map. Click **Add** to display a selector. Select one or more items from the Available Precedences list, then click >> to add them to the Selected Precedences list. For more information about IP precedence values, see Table 1: IP Precedence Classes , on page 3. When you finish, click **OK** to return to the QoS Class dialog box. Your selections are displayed in the Precedence field. **Note** To remove an IPP value from the QoS class, select it from the Precedence field, then click **Delete**.
DSCP	One or more Differentiated Services Code Point (DSCP) values included in this class map. Click **Add** to display a selector. Select one or more items (up to eight) from the Available DSCPs list, then click >> to add them to the Selected DSCPs list. When you finish, click **OK** to return to the QoS Class dialog box. Your selections are displayed in the DSCP field. **Note** To remove a DSCP value from the QoS class, select it from the DSCP field, then click **Delete**.

Element	Description
ACL	The ACLs that are used for defining which traffic requires QoS. Click **Edit** to add or remove ACL objects. Use the up and down arrows to order the ACLs in the list. We recommend that you place frequently used ACLs at the top of the list to optimize the matching process.

Edit ACLs Dialog Box—QoS Classes

When configuring a QoS policy on a Cisco IOS router, use the Edit ACLs dialog box to specify which ACLs should be included in the matching criteria for the selected QoS class. Traffic matching this criteria is included as part of the class.

Enter the names of the extended ACLs or click **Select** to select an ACL object from a list or to create a new one. Separate multiple ACL objects with commas and place them in priority order.

For more information, see Creating Extended Access Control List Objects.

Navigation Path

Go to the QoS Class Dialog Box—Matching Tab , on page 25, then click **Edit** in the ACL field.

Related Topics

- Defining QoS Class Matching Parameters , on page 14
- Defining QoS on Interfaces , on page 11
- Defining QoS on the Control Plane , on page 13
- Quality of Service Policy Page , on page 20
- Selecting Objects for Policies

QoS Class Dialog Box—Marking Tab

Use the Marking tab of the QoS Class dialog box to classify packets. Traffic policers and shapers use these classifications to ensure adherence to the contracted level of service. Downstream devices use this classification to identify the packets and apply the appropriate QoS functions to them.

Note The Marking tab is unavailable when you define a QoS policy on the control plane.

Navigation Path

Go to the QoS Class Dialog Box , on page 24, then click the **Marking** tab.

Related Topics

- Defining QoS Class Marking Parameters , on page 16
- Defining QoS on Interfaces , on page 11

- Defining QoS on the Control Plane , on page 13
- Quality of Service Policy Page , on page 20

Field Reference

Table 7: QoS Class Dialog Box—Marking Tab

Element	Description
Enable Marking	When selected, enables you to mark the traffic in this QoS class with a specific precedence or DSCP value (regardless of any value the traffic might have had when it first entered the device). This mark enables downstream devices to identify the traffic and apply the appropriate QoS features to it. When deselected, disables all marking options for the selected QoS class. The traffic in this QoS class maintains its original precedence or DSCP value, if any.
Precedence	The precedence value with which to mark the traffic in this class: • network (7) • internet match (6) • critical (5) • flash-override (4) • flash (3) • immediate (2) • priority (1) • routine (0)
DSCP	The DSCP value (0 to 63) with which to mark the traffic in this class.

QoS Class Dialog Box—Queuing and Congestion Avoidance Tab

Use the Queuing and Congestion Avoidance tab of the QoS Class dialog box to perform Class-Based Weighted Fair Queuing (CBWFQ) on the output traffic in the selected QoS class. Queuing prioritizes traffic and manages congestion on your network by determining the order in which packets are sent out over an interface. Queuing and congestion avoidance applies only to interface classes for output traffic.

The fields displayed in the Queuing tab depend on whether you are defining a specific QoS class or the default class (by selecting **Set as Default Class**), and also by the type of router and the Cisco IOS software version.

Navigation Path

Go to the QoS Class Dialog Box , on page 24, then click the **Queuing and Congestion Avoidance** tab.

Related Topics

- Defining QoS Class Queuing Parameters , on page 16

- Defining QoS on Interfaces , on page 11
- Defining QoS on the Control Plane , on page 13
- Quality of Service Policy Page , on page 20

Field Reference

Table 8: QoS Class Dialog Box—Queuing and Congestion Avoidance Tab

Element	Description
Enable Queuing and Congestion Avoidance	Whether to configure queuing and congestion avoidance properties in the QoS class.
Priority (Non-default classes only.)	Configure low-latency queuing (LLQ) in this class to ensure that priority traffic, such as voice traffic, receives the defined bandwidth (see Low-Latency Queuing , on page 6). Specify the amount of bandwidth allocated to high-priority traffic on this interface by: • Percentage—Valid values range from 1 to 100%. • Kbit/sec—Valid values range from 8-2000000 kilobits per second. **Note** You can define this option for one class only per interface. If you select this option, the Shaping tab is disabled.
Fair Queue Number of Dynamic Queues (Default class only.)	Configure class-based weighted fair queuing in this class. If the device is running an IOS software version lower than 12.4(20)T, you must specify the number of dynamic queues to reserve for this class. You should base your number on the available bandwidth of the interface. You can specify a number between 16 and 4096 that is a power to 2. For information on the default number of queues the device uses, see Default Class Queuing , on page 6. Available bandwidth is evenly distributed among the queues unless you configure a queue limit. **Tip** Failure to provide a sufficient number of queues for the default class (a condition known as starvation) could result in the traffic not being sent.
Bandwidth	Configure the minimum bandwidth to guarantee to this class. You can define this amount by: • Percentage—Valid values range from 1 to 100% of the total available bandwidth. • Kbit/sec—Valid values range from 8-2000000 kilobits per second.
Enable Fair Queue (Non-default class only.)	When you configure bandwidth for a non-default class, whether to also enable class-based weighted fair queuing (CBWFQ). The device calculates the number of queues to configure based on the available bandwidth, and distributes the bandwidth evenly among the queues unless you configure a queue limit. This option is available only for Aggregation Services Routers (ASR) and for routers running 12.4(20)T and later.

Element	Description
Queue Limit	The maximum number of packets that can be queued for the class. Any additional packets are dropped using tail drop until the congestion is gone. This is the default option for limiting queue size unless Weighted Random Early Detection (WRED) is configured.
WRED Weight for Mean Queue Depth	The exponential weight factor to use to calculate the average queue size. Use this option when defining WRED instead of tail drop (queue limit) for this class. When the queue size exceeds the value determined by this weight factor, WRED randomly discards packets until the transmitting protocol decreases its transmission rate to ease congestion. Exponent values range from 1 to 16. The default is 9. This option is best suited for protocols like TCP, which respond to dropped packets by decreasing the transmission rate. We recommend that you do not change the default unless you determine that your applications would benefit from the change.

QoS Class Dialog Box—Policing Tab

Use the Policing tab of the QoS Class dialog box to configure rate limits on the traffic in a selected QoS class. Excess traffic is either dropped or transmitted with a different (typically lower) priority.

Navigation Path

Go to the QoS Class Dialog Box , on page 24, then click the **Policing** tab.

Related Topics

- Defining QoS Class Policing Parameters , on page 18
- Defining QoS on Interfaces , on page 11
- Defining QoS on the Control Plane , on page 13
- Quality of Service Policy Page , on page 20

Field Reference

Table 9: QoS Class Dialog Box—Policing Tab

Element	Description
Enable Policing	When selected, enables you to configure Class-Based Policing to control the maximum rate of traffic for this class. Security Manager uses a two-token bucket algorithm, which includes a defined violate action that is performed when neither bucket can accommodate the incoming packet. When deselected, disables all policing options for the selected QoS class.

Element	Description
CIR	The average data rate (also known as the committed information rate or CIR). You can define this amount by: • Percentage—Valid values range from 0 to 100% of the overall available bandwidth. • Bit/sec—Valid values range from 8000 to 2000000000 bits per second. In the token bucket algorithm, this rate represents the token arrival rate for filling both token buckets. Traffic that falls under this rate always conforms. **Note** When you configure Understanding Control Plane Policing , on page 10, you must define the CIR in bits per second.
Conform Burst	The normal burst size, which determines how large traffic bursts can be before some traffic exceeds the rate limit. In the token bucket algorithm, it represents the full size of the first (conform) token bucket. The range of valid values is determined by the CIR: • When the CIR is defined by percentage—Valid values range from 1 to 2000 milliseconds. • When the CIR is defined by an absolute value—Valid values range from 1000-512000000 bytes.
Excess Burst	The excess burst size, which determines how large traffic bursts can be before all traffic exceeds the rate limit. In the token bucket algorithm, it represents the full size of the second (exceed) token bucket. The range of valid values is determined by the CIR: • When the CIR is defined by percentage—Valid values range from 1 to 2000 milliseconds. • When the CIR is defined by an absolute value—Valid values range from 1000-512000000 bytes.
Conform action	The action to take on packets that conform to the rate limit: • transmit—Transmits the packet. • set-prec-transmit—Sets the IP precedence to a value you specify (0 to 7) and then sends the packet. Not available on the control plane. • set-dscp-transmit—Sets the DSCP to a value you specify (0 to 63) and then sends the packet. Not available on the control plane. • drop—Drops the packet.

Element	Description
Exceed action	The action to take on packets that exceed the rate limit, but can be handled using the second (exceed) token bucket. The actions available for selection depend on the defined conform action. For example, if you select one of the set options as the conform action, you cannot select transmit as the exceed action. If you select drop as the conform action, then you must also select drop as the exceed action.
Violate action	The action to take on packets that cannot be serviced by either the conform bucket or the exceed bucket. The actions available for selection depend on the defined exceed action. For example, if you select one of the set options as the exceed action, you cannot select transmit as the violate action. If you select drop as the exceed action, then you must also select drop as the violate action.

QoS Class Dialog Box—Shaping Tab

Use the Shaping tab of the QoS Class dialog box to control the rate of output traffic for the selected QoS class. Shaping typically delays excess traffic by using a buffer, or queuing mechanism, to hold packets and shape the flow when the data rate of the source is higher than expected.

Note The Shaping tab is unavailable when you define a QoS policy on the control plane, use hierarchical shaping on the interface, define a QoS class for input traffic, or perform queuing on priority traffic.

Navigation Path

Go to the QoS Class Dialog Box , on page 24, then click the **Shaping** tab.

Related Topics

- Defining QoS Class Shaping Parameters , on page 19
- Defining QoS on Interfaces , on page 11
- Defining QoS on the Control Plane , on page 13
- Quality of Service Policy Page , on page 20

Field Reference

Table 10: QoS Class Dialog Box—ShapingTab

Element	Description
Enable Shaping	When selected, enables you to configure Distributed Traffic Shaping (DTS) to control the rate of traffic for this class. DTS uses queues to buffer traffic surges that can congest the network. When deselected, disables all shaping options for the selected QoS class. **Note** Shaping can be performed only on output traffic.
Type	The type of shaping to perform: • Average—Limits the data rate for each interval to the sustained burst rate (also known as the committed burst rate or Bc), achieving an average rate no higher than the committed information rate (CIR). Additional packets are buffered until they can be sent. • Peak—Limits the data rate for each interval to the sustained burst rate plus the excess burst rate (Be). Additional packets are buffered until they can be sent.
CIR	The average data rate (also known as the committed information rate or CIR). You can define this amount by: • Percentage—Valid values range from 0 to 100% of the overall available bandwidth. • Bit/sec—Valid values range from 8000 to 1000000000 bits per second, and must be in multiples of 8000. Although data bursts during an interval may exceed this rate, the average data rate over any multiple integral of the interval will not exceed this rate.
Sustained Burst	The normal burst size. If you select average as the shaping type, data bursts during an interval are limited to this value. The range of valid values is determined by the CIR: • When the CIR is defined by percentage—Valid values range from 10 to 2000 milliseconds. • When the CIR is defined by an absolute value—Valid values range from 1000 to 154400000 bytes, in multiples of 128 bytes. **Note** We recommend that you leave this field blank when the CIR is defined by an absolute value. This allows the algorithms used by the device to determine the optimal sustained burst value.

Element	Description
Excess Burst	The excess burst size. If you select peak as the shaping type, data bursts during an interval can equal the sum of the sustained burst value plus this value. The average data rate over multiple intervals, however, will continue to conform to the CIR. The range of valid values is determined by the CIR: • When the CIR is defined by percentage—Valid values range from 10 to 2000 milliseconds. • When the CIR is defined by an absolute value—Valid values range from 1000 to 154400000 bytes, in multiples of 128 bytes. **Note** If you do not configure this field when the CIR is defined by an absolute value, the sustained burst value is used.