

Configuring Quality of Service

This chapter describes the quality of service (QoS) features built into your Layer 3 switch router and how to map QoS scheduling at both the system and interface levels. This chapter contains the following sections:

- Understanding Layer 3 Switching Quality of Service, page 9-1
- Understanding Scheduling and Weighted Round-Robin, page 9-2
- Configuring Precedence to WRR Scheduling, page 9-3
- Configuring Per-Port Traffic Conditioning, page 9-4

Understanding Layer 3 Switching Quality of Service

Extensive core quality of service (QoS) mechanisms are built into the Layer 3 switch router architecture to ensure policy enforcement and queuing, as well as weighted round-robin (WRR) scheduling at the egress port.

Layer 3-switching QoS is based on IP precedence for partitioning traffic into multiple classes of service. IP precedence uses the three type-of-service precedence bits in the IP header to specify the class of service assignment for each packet.

IP precedence can be mapped into adjacent technologies (for example, tag switching or ATM) to deliver end-to-end network QoS policies. This feature enables service classes to be established with no changes to existing applications and with no complicated network signaling requirements.

The system gathers IP precedence information from the IP header type-of-service field. For an incoming IP packet, the first two (most significant) bits of the service type field determine the delay priority. Layer 3 switching recognizes four QoS classes, Q-0 to Q-3, as summarized in Table 9-1.

Table 9-1 QoS Delay Priorities and Queues

IP Precedence Bits	Delay Priority	Queue Selected
0 0 0	0 0	Q-0
0 0 1	0 0	Q-0
0 1 0	0 1	Q-1
0 1 1	0 1	Q-1
1 0 0	1 0	Q-2

Table 9-1 QoS Delay Priorities and Queues (continued)

IP Precedence Bits	Delay Priority	Queue Selected
1 0 1	1 0	Q-2
1 1 0	1 1	Q-3
1 1 1	1 1	Q-3

Your Layer 3 switch router can read the precedence field and switch the packet accordingly, but it cannot reclassify traffic. The edge router or switch is expected to set the precedence field according to its local policy.

The Layer 3 switch router queues packets to the appropriate queue based on the delay priority.

Understanding Scheduling and Weighted Round-Robin

Frame scheduling becomes increasingly important when an outgoing interface is congested. To handle this situation, network administrators can assign weights to each of the different queues. This feature provides bandwidth to higher priority applications (using IP precedence), and also grants access to lower priority queues. The frame schedule affords each queue the bandwidth allotted to it by the network administrator. This mapping is configurable both at the system and interface levels (as described later in this chapter).

The four queues on any destination interface are configured to be part of the same service class. Bandwidth is not explicitly reserved for these four queues. Each of them is assigned a different WRR-scheduling weight, which determines the way they share the interface bandwidth. The WRR weight is user configurable; you can assign a different WRR weight for each queue.



Tips The higher the WRR weight, the higher the effective bandwidth for that particular queue.

You can find the effective bandwidth (in Megabits per second [Mbps]) for a particular queue with the following formula:

$$(W/S) \times B = n \text{ Mbps},$$

where W is the WRR weight of the specified queue, S is the sum of the weight of all active queues on the outgoing interface, B is the available bandwidth in Mbps, and n is the effective bandwidth in Mbps.

For example, if W is 4, S is 15, and B is 100, the formula would be $(4/15) \times 100 = 26$ Mbps. The effective bandwidth for the specified queue in this example is 26 Mbps.

Configuring Precedence to WRR Scheduling

This section describes the Cisco IOS commands necessary to configure QoS mapping at the system and interface levels. The commands described in this section are unique to Layer 3 switching software.

Layer 3 switching software enables QoS-based forwarding by default. If disabled, enter the following command to enable QoS switching:

```
# qos switching
```

Configuring QoS Scheduling at the System Level

To configure QoS scheduling at the system level, perform the following task beginning in global configuration mode:

	Command	Purpose
Step 1	Router(config)# qos mapping precedence <i>value</i> wrr-weight <i>weight</i>	Set the mapping between IP precedence and the WRR weight. The command syntax is as follows: <ul style="list-style-type: none"> The precedence <i>value</i> (0 to 3) is the higher 2 bits of the IP precedence field. The WRR-scheduling <i>weight</i> (1 to 4). This parameter specifies the weight assigned to traffic with the given precedence.
Step 2	Router(config)# end Router#	Return to privileged EXEC mode.

To set the precedence back to the default setting for the Layer 3 switch router, enter the **no** version of the **qos mapping precedence** command.

Table 9-2 shows the default WRR weights for IP precedence.

Table 9-2 IP Precedence and Default WRR Weights

IP Precedence	WRR Weight
0	1
1	2
2	3
3	4

Configuring QoS Scheduling at the Interface Level

Configuring the QoS mapping at the interface level overrides the system-level mapping. Using the **qos mapping precedence wrr-weight** command, you can assign different WRR-scheduling weights for a particular precedence traffic destined to any interface.

To configure QoS scheduling at the interface level, perform the following task beginning in global configuration mode:

	Command	Purpose
Step 1	Router(config)# qos mapping [destination <i>dest-interface</i>] precedence <i>value</i> wrr-weight <i>weight</i>	Assign different WRR-scheduling weights for a particular precedence traffic destined to an interface.
Step 2	Router(config)# end Router#	Return to privileged EXEC mode.

The QoS commands are applicable to both Gigabit Ethernet and Fast Ethernet interfaces.

To set the precedence back to the system-level default setting for the Layer 3 switch router, enter the **no** version of the **qos mapping precedence wrr-weight** command.

The destination interface parameter is optional. When the destination interface parameter is not specified, the system-level QoS mapping is configured. If interface-specific QoS mapping is configured for a certain precedence, it takes precedence over the system-level QoS mapping for that precedence.

Monitoring and Verifying the QoS Configuration

To verify the QoS configuration, perform the following tasks:

Command	Purpose
Router# show qos switching	Show whether QoS-based switching is enabled.
Router# show qos mapping [destination <i>dest-interface</i>]	Show effective mapping at either the system level or interface level.

For more information on QoS, refer to the *Cisco IOS Quality of Service Solutions Configuration Guide*.

Configuring Per-Port Traffic Conditioning

This section describes the Cisco IOS commands necessary to configure per-port QoS traffic conditioning at the interface level. The per-port traffic conditioning commands are applicable to both Gigabit Ethernet and Fast Ethernet interfaces. The Catalyst 2948G-L3 and 4908G-L3 switch routers support per-port input rate limiting, output rate limiting, and traffic shaping based on the token bucket algorithm beginning with Release 12.0(10)W5(18e).

A token bucket is a formal definition of a rate of transfer. The mean rate determines the long-term average transmission rate and the burst size determines how large traffic bursts can be sent before some traffic exceeds the contract.

Tokens are inserted into the bucket at the mean rate. The depth of the bucket is the burst size. Traffic arriving at the bucket when sufficient tokens are available is said to conform and the corresponding number of tokens are removed from the bucket. If sufficient tokens are not available, then the traffic is said to exceed (nonconforming).

The rate limiter and the shaper identify traffic violations in an identical manner. They differ, however, in the way they respond to traffic violations. The rate limiter drops the traffic that is not conforming to the traffic contract, whereas the shaper delays excess traffic using a buffer and drops traffic only when the buffer overflows.

Token bucket-based per-port traffic conditioning features support a minimum rate of 32 kbps and have a granularity of 32 kbps. The additional processing associated with these features may affect the peak data forwarding performance that limits the peak rate at which traffic conditioning may be applied.

The per-port traffic conditioning features cannot be enabled on any of the interfaces when IPX routing is enabled; IPX routing cannot be enabled when any of the per-port traffic conditioning features are enabled on any of the interfaces. Transitioning between IPX routing and per-port traffic conditioning involves dynamic downloading of new binaries to the network processor. We recommend that you perform this dynamic downloading under light traffic conditions.



Note

You cannot configure per-port traffic conditioning features on the FEC, GEC, Bridge-Group Virtual Interface (BVI), or on subinterfaces.

Configuring Per-Port Input Rate Limiting

The per-port input rate limiting feature allows you to rate limit the input traffic of a port. The input traffic rate of the port is monitored checking for nonconforming traffic. The nonconforming input traffic is dropped and the conforming input traffic is allowed.

Rate limiting is applied to all input traffic and is not confined to IP Layer 3 traffic. However, high priority traffic (such as routing updates or BPDUs) destined to the CPU is not subjected to per-port input rate limiting.

To configure per-port input rate limiting at the interface level, perform the following task beginning in interface configuration mode:

	Command	Purpose
Step 1	Router(config-if)# rate-limit input {32000-100000000} {0-64000}	Assign different rate limits on a per-physical port basis. The first parameter specifies the mean rate in bits per second, and the second parameter is the burst size in bytes.
Step 2	Router(config-if)# end	Return to privileged EXEC mode.
Step 3	Router# show run interface {interface_name}	Display the interface configuration to verify changes.

To disable the per-port input rate limiting, enter the **no** version of the **rate-limit input** command.

The following example shows how to configure per-port input rate limiting:

```
Router(config)# interface fastethernet 1
Router(config-if)# rate-limit input 1000000 1000
Router(config-if)# end
Router# show run interface fastethernet 1
```

Configuring Per-Port Shaping

The per-port shaping feature allows you to shape the output of a port. The output traffic rate of the port is monitored to verify that the traffic leaves the interface at the user-configured rate. When excess traffic comes into the Layer 3 switch router, back pressure is applied from the modules to the switch fabric, and the excess traffic gets queued in the switch fabric. If the switch fabric queues overflow, the excess traffic is dropped.



Note Shaping is applied to all output traffic and is not confined to IP Layer 3 traffic. However, traffic originating from the CPU or process switched by the CPU is not subjected to per-port shaping.



Note You cannot configure per-port output rate limiting and shaping on an interface at the same time.

To configure per-port shaping at the interface level, perform the following task beginning in interface configuration mode:

	Command	Purpose
Step 1	Router(config-if)# traffic-shape rate {32000-100000000} {0-512000}	Assign different traffic shaping rate limits on a per-physical port basis. The first parameter specifies the mean rate in bits per second, and the second parameter is the burst size in bits.
Step 2	Router(config-if)# end	Return to privileged EXEC mode.
Step 3	Router# show run interface {interface_name}	Display the interface configuration to verify changes.

To disable per-port traffic shaping, enter the **no** version of the **traffic-shape rate** command.

The following example shows how to configure per-port shaping:

```
Router(config)# interface fastethernet 1
Router(config-if)# traffic-shape rate 1000000 1000
Router(config-if)# end
Router# show run interface fastethernet 1
```

Configuring Per-Port Output Rate Limiting

The per-port output rate limiting feature allows you to rate limit the output traffic of a port. The output traffic rate of the port is monitored checking for conformance and nonconforming output traffic is dropped.



Note Rate limiting is applied to all output traffic and is not confined to IP Layer 3 traffic. However, traffic originating from the CPU or process switched by the CPU is not subjected to per-port output rate limiting.

**Note**

You cannot configure per-port output rate limiting and shaping on an interface at the same time.

To configure per-port output rate limiting at the interface level, perform the following task beginning in interface configuration mode:

	Command	Purpose
Step 1	Router(config-if)# rate-limit output {32000-100000000} {0-64000}	Assign different rate limits on a per-physical port basis. The first parameter specifies the mean rate in bits per second, and the second parameter is the burst size in bytes.
Step 2	Router(config-if)# end Router#	Return to privileged EXEC mode.
Step 3	Router# show run interface {interface_name}	Display the interface configuration to verify changes.

To disable per-port output rate limiting, enter the **no** version of the **rate-limit output** command.

The following example shows how to configure per-port output rate limiting:

```
Router(config)# interface fastethernet 1
Router(config-if)# rate-limit output 1000000 1000
Router(config-if)# end
Router# show run interface fastethernet 1
```

**Caution**

You will be unable to use IPX routing when you enable any of the per-port traffic conditioning features on any port of a Catalyst 2948G-L3 or 4908G-L3 switch router.

■ Configuring Per-Port Traffic Conditioning