



Configuring QoS

This chapter describes how to configure quality of service (QoS) by using standard QoS commands on the Catalyst 2940 switch. With QoS, you can give preferential treatment to certain types of traffic at the expense of others. Without QoS, the switch offers best-effort service to each packet, regardless of the packet contents or size. It sends the packets without any assurance of reliability, delay bounds, or throughput.



Note

For complete syntax and usage information for the commands used in this chapter, refer to the command reference for this release.

QoS can be configured either by using the Cluster Management Suite (CMS) or through the command-line interface (CLI). Refer to the CMS online help for configuration procedures through CMS. For information about accessing and using CMS, see [Chapter 3, “Getting Started with CMS.”](#)

This chapter consists of these sections:

- [Understanding QoS, page 27-1](#)
- [Configuring QoS, page 27-3](#)
- [Displaying QoS Information, page 27-10](#)

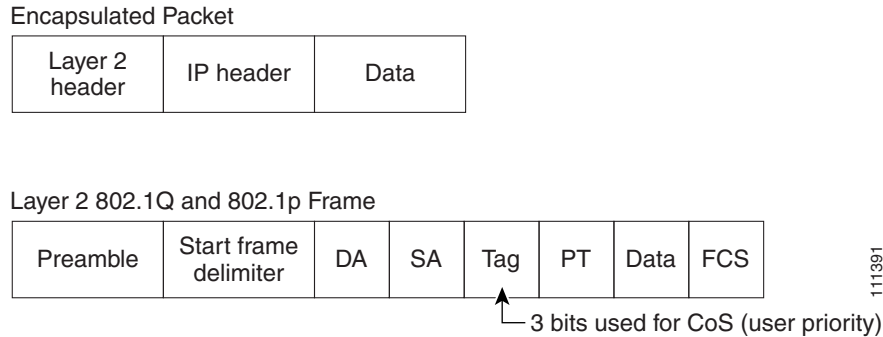
Understanding QoS

This section describes how QoS is implemented on the switch. Typically, networks operate on a best-effort delivery basis, which means that all traffic has equal priority and an equal chance of being delivered in a timely manner. When congestion occurs, all traffic has an equal chance of being dropped.

When you configure the QoS feature, you can use congestion-management and congestion-avoidance techniques to give preferential treatment. Implementing QoS in your network makes network performance more predictable and bandwidth utilization more effective.

The QoS implementation is based on the prioritization values in Layer 2 frames

Layer 2 802.1Q frame headers have a 2-byte Tag Control Information field that carries the class of service (CoS) value in the three most-significant bits, which are called the User Priority bits. On interfaces configured as Layer 2 802.1Q trunks, all traffic is in 802.1Q frames except for traffic in the native VLAN. Other frame types cannot carry Layer 2 CoS values. Layer 2 CoS values range from 0 for low priority to 7 for high priority.

Figure 27-1 QoS Classification Layers in Frames and Packets

All switches and routers that access the Internet rely on the class information to give the same forwarding treatment to packets with the same class information and different treatment to packets with different class information. The class information in the packet can be assigned by end hosts or by switches or routers along the way, based on a configured policy, detailed examination of the packet, or both. Detailed examination of the packet is expected to happen closer to the edge of the network so that the core switches and routers are not overloaded.

Switches and routers along the path can use the class information to limit the amount of resources allocated per traffic class. The behavior of an individual device when handling traffic is called per-hop behavior. If all devices along a path have a consistent per-hop behavior, you can construct an end-to-end QoS solution.

Implementing QoS in your network can be a simple or complex task and depends on the QoS features offered by your internetworking devices, the traffic types and patterns in your network, and the granularity of control that you need over incoming and outgoing traffic.

Queueing and Scheduling

The switch gives QoS-based 802.1p CoS values. QoS uses classification and scheduling to send network traffic from the switch in a predictable manner. QoS classifies frames by examining priority-indexed CoS values in them and gives preference to higher-priority traffic such as telephone calls.

How Class of Service Works

Before you set up 802.1p CoS on a Catalyst 2940 switch that operates with the Catalyst 6000 family of switches, refer to the Catalyst 6000 documentation. There are differences in the 802.1p implementation that you should understand to ensure compatibility.

Port Priority

Frames received from users in the administratively-defined VLANs are classified or *tagged* for transmission to other devices. Based on rules that you define, a unique identifier (the tag) is inserted in each frame header before it is forwarded. The tag is examined and understood by each device before any broadcasts or transmissions to other switches, routers, or end stations. When the frame reaches the last switch or router, the tag is removed before the frame is sent to the target end station. VLANs that are assigned on trunk or access ports without identification or a tag are called *native* or *untagged* frames.

For IEEE 802.1Q frames with tag information, the priority value from the header frame is used. For native frames, the default priority of the input port is used.

Each port on the switch has a single receive queue buffer (the *ingress* port) for incoming traffic. When an untagged frame arrives, it is assigned the value of the port as its port default priority. You assign this value by using the CLI or CMS. A tagged frame continues to use its assigned CoS value when it passes through the ingress port.

Egress CoS Queues

The switch supports four CoS queues for each egress port. For each queue, you can specify these types of scheduling:

- Strict priority scheduling

Strict priority scheduling is based on the priority of queues. Packets in the high-priority queue always transmit first, and packets in the low-priority queue do not transmit until all the high-priority queues become empty.

The default scheduling method is strict priority.

- Weighted round-robin (WRR) scheduling

WRR scheduling requires you to specify a number that indicates the importance (weight) of the queue relative to the other CoS queues. WRR scheduling prevents the low-priority queues from being completely neglected during periods of high-priority traffic. The WRR scheduler sends some packets from each queue in turn. The number of packets it sends corresponds to the relative importance of the queue. For example, if one queue has a weight of 3 and another has a weight of 4, three packets are sent from the first queue for every four that are sent from the second queue. By using this scheduling, low-priority queues have the opportunity to send packets even though the high-priority queues are not empty.

Configuring QoS

Before configuring QoS, you must have a thorough understanding of these items:

- The types of applications used and the traffic patterns on your network.
- Traffic characteristics and needs of your network. Is the traffic bursty? Do you need to reserve bandwidth for voice and video streams?
- Bandwidth requirements and speed of the network.
- Location of congestion points in the network.

This section describes how to configure QoS on your switch:

- [Default QoS Configuration, page 27-4](#)
- [Configuring Classification Using Port Trust States, page 27-4](#)
- [Configuring the Egress Queues, page 27-9](#)

Default QoS Configuration

This is the default QoS configuration:

- The default port CoS value is 0.
- The default port CoS value is assigned to all incoming untagged packets. The CoS value of each tagged packet remains unaltered.
- By default, the port trust state is not configured.
- All traffic is sent through one egress queue.

Configuring Classification Using Port Trust States

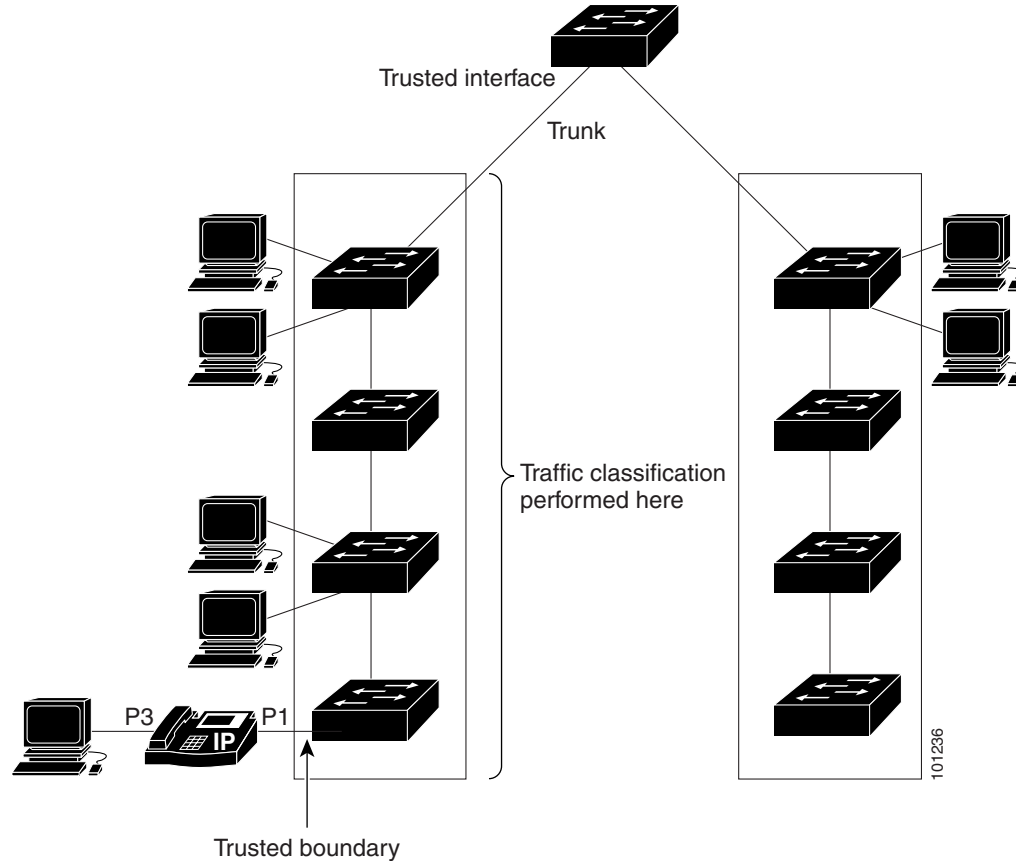
This section describes how to classify incoming traffic by using port trust states:

- [Configuring the Trust State on Ports within the QoS Domain, page 27-4](#)
- [Configuring the CoS Value for an Interface, page 27-6](#)
- [Configuring Trusted Boundary, page 27-7](#)

Configuring the Trust State on Ports within the QoS Domain

Packets entering a QoS domain are classified at the edge of the QoS domain. When the packets are classified at the edge, the switch port within the QoS domain can be configured to one of the trusted states because there is no need to classify the packets at every switch within the QoS domain. [Figure 27-2](#) shows a sample network topology.

Figure 27-2 Port Trusted States within the QoS Domain



Beginning in privileged EXEC mode, follow these steps to configure the port to trust the classification of the traffic that it receives:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Specify the interface to be trusted, and enter interface configuration mode. Valid interfaces include physical interfaces.
Step 3	mls qos trust [cos]	Configure the port trust state. By default, the port is not trusted. All traffic is sent through one egress queue. Use the cos keyword to classify ingress packets with the packet CoS values. The egress queue assigned to the packet is based on the packet CoS value. When this keyword is entered, the traffic is sent through the four QoS queues, as described in the “Configuring the Egress Queues” section on page 27-9. For more information about this command, refer to the command reference for this release.
Step 4	end	Return to privileged EXEC mode.

	Command	Purpose
Step 5	<code>show mls qos interface [interface-id]</code>	Verify your entries.
Step 6	<code>copy running-config startup-config</code>	(Optional) Save your entries in the configuration file.

After you enter the `mls qos trust cos` command, the DSCP values are changed according to the values listed in [Table 27-1](#).

Table 27-1 Default CoS-to-DSCP Map

CoS value	0	1	2	3	4	5	6	7
DSCP value	0	8	16	24	32	40	48	56



Note CoS-to-DSCP values cannot be configured.

To return a trusted port to its unconfigured state, use the `no mls qos trust` interface configuration command.

Configuring the CoS Value for an Interface

QoS assigns the CoS value specified with the `mls qos cos` interface configuration command to untagged frames received on trusted and untrusted ports.

Beginning in privileged EXEC mode, follow these steps to define the default CoS value of a port or to assign the default CoS to all incoming packets on the port:

	Command	Purpose
Step 1	<code>configure terminal</code>	Enter global configuration mode.
Step 2	<code>interface interface-id</code>	Specify the interface to be trusted, and enter interface configuration mode. Valid interfaces include physical interfaces.
Step 3	<code>mls qos cos {default-cos override}</code>	Configure the default CoS value for the port. <ul style="list-style-type: none"> For <i>default-cos</i>, specify a default CoS value to be assigned to a port. If the port is CoS trusted and packets are untagged, the default CoS value becomes the CoS value for the packet. The CoS range is 0 to 7. The default is 0. Use the override keyword to override the previously configured trust state of the incoming packets and to apply the default port CoS value to all incoming packets. By default, CoS override is disabled. <p>Use the override keyword when all incoming packets on certain ports deserve higher priority than packets entering from other ports. If an incoming packet is tagged, the CoS value of the packet is modified with the default CoS of the port at the egress port.</p>
Step 4	<code>end</code>	Return to privileged EXEC mode.
Step 5	<code>show mls qos interface</code>	Verify your entries.
Step 6	<code>copy running-config startup-config</code>	(Optional) Save your entries in the configuration file.

To return to the default setting, use the **no mls qos cos** {*default-cos* | **override**} interface configuration command.

Configuring Trusted Boundary

In a typical network, you connect a Cisco IP Phone to a switch port as shown in [Figure 27-2 on page 27-5](#). Traffic sent from the telephone to the switch is typically marked with a tag that uses the 802.1Q header. The header contains the VLAN information and the CoS 3-bit field, which determines the priority of the packet. For most Cisco IP Phone configurations, the traffic sent from the telephone to the switch is trusted to ensure that voice traffic is properly prioritized over other types of traffic in the network. By using the **mls qos trust cos** interface configuration command, you can configure the switch port to which the telephone is connected to trust the CoS labels of all traffic received on that port.

In some situations, you also might connect a PC or workstation to the IP phone. In these cases, you can use the **switchport priority extend cos** interface configuration command to configure the telephone through the switch CLI to override the priority of the traffic received from the PC. With this command, you can prevent a PC from taking advantage of a high-priority data queue.

However, if a user bypasses the telephone and connects the PC directly to the switch, the CoS labels generated by the PC are trusted by the switch (because of the trusted CoS setting) and can allow misuse of high-priority queues. The trusted boundary feature solves this problem by using the CDP to detect the presence of a Cisco IP Phone (such as the Cisco IP Phone 7910, 7935, 7940, and 7960) on a switch port. If the telephone is not detected, the trusted boundary feature disables the trusted setting on the switch port and prevents misuse of a high-priority queue.

Beginning in privileged EXEC mode, follow these steps to configure trusted boundary on a switch port:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	cdp enable	Enable CDP globally. By default, it is enabled.
Step 3	interface <i>interface-id</i>	Specify the interface to be trusted, and enter interface configuration mode. Valid interfaces include physical interfaces.
Step 4	cdp enable	Enable CDP on the interface. By default, CDP is enabled.
Step 5	mls qos trust device cisco-phone	Configure the Cisco IP Phone as a trusted device on the interface.
Step 6	mls qos trust cos	Configure the port trust state to trust the CoS value of the ingress packet. By default, the port is not trusted. For more information on this command, refer to the command reference for this release.
Step 7	end	Return to privileged EXEC mode.
Step 8	show mls qos interface [<i>interface-id</i>]	Verify your entries.
Step 9	copy running-config startup-config	(Optional) Save your entries in the configuration file.

When you enter the **no mls qos trust** interface configuration command, trusted boundary is not disabled. If this command is entered and the port is connected to a Cisco IP Phone, the port does not trust the classification of traffic that it receives. To disable trusted boundary, use the **no mls qos trust device** interface configuration command.

If you enter the **mls qos cos override** interface configuration command, the port does not trust the classification of the traffic that it receives, even when it is connected to a Cisco IP Phone.

Table 27-2 lists the port configuration when an IP phone is present or absent.

Table 27-2 Port Configurations When Trusted Boundary is Enabled

Port Configuration	When a Cisco IP Phone is Present	When a Cisco IP Phone is Absent
The port trusts the CoS value of the incoming packet.	The packet CoS value is trusted.	The packet CoS value is assigned the default CoS value.
The port assigns the default CoS value to incoming packets.	The packet CoS value is assigned the default CoS value.	The packet CoS value is assigned the default CoS value.

Enabling Pass-Through Mode

When the switch is in pass-through mode, it uses the CoS value of incoming packets without modifying the DSCP value and sends the packets from one of the four egress queues. By default, pass-through mode is disabled. The switch assigns a CoS value of 0 to all incoming packets without modifying the packets. The switch offers best-effort service to each packet regardless of the packet contents or size and sends it from a single egress queue.

Beginning in privileged EXEC mode, follow these steps to enable pass-through mode:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Specify the interface on which pass-through mode is enabled, and enter interface configuration mode. Valid interfaces include physical interfaces.
Step 3	mls qos trust cos pass-through dscp	Enable pass-through mode. The interface is configured to trust the CoS value of the incoming packets and to send them without modifying the DSCP value.
Step 4	end	Return to privileged EXEC mode.
Step 5	show mls qos interface [<i>interface-id</i>]	Verify your entries.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable pass-through mode, use the **no mls qos trust pass-through dscp** interface configuration command.

If you enter the **mls qos cos override** and the **mls qos trust [cos]** interface commands when pass-through mode is enabled, pass-through mode is disabled.

If you enter the **mls qos trust cos pass-through dscp** interface configuration command when the **mls qos cos override** and the **mls qos trust [cos]** interface commands are already configured, pass-through mode is disabled.

Configuring the Egress Queues

This section describes how to configure the egress queues:

- [Configuring CoS Priority Queues, page 27-9](#)
- [Configuring WRR Priority, page 27-9](#)

For more information about the egress queues, see the “Egress CoS Queues” section on page 27-3.

Configuring CoS Priority Queues

Beginning in privileged EXEC mode, follow these steps to configure the CoS priority queues:

	Command	Purpose										
Step 1	configure terminal	Enter global configuration mode.										
Step 2	wrr-queue cos-map <i>qid cos1...cosn</i>	Specify the queue ID of the CoS priority queue. (Ranges are 1 to 4 where 1 is the lowest CoS priority queue.) Specify the CoS values that are mapped to the queue id. Default values are as follows: <table border="1"> <thead> <tr> <th>CoS Value</th> <th>CoS Priority Queues</th> </tr> </thead> <tbody> <tr> <td>0, 1</td> <td>1</td> </tr> <tr> <td>2, 3</td> <td>2</td> </tr> <tr> <td>4, 5</td> <td>3</td> </tr> <tr> <td>6, 7</td> <td>4</td> </tr> </tbody> </table>	CoS Value	CoS Priority Queues	0, 1	1	2, 3	2	4, 5	3	6, 7	4
CoS Value	CoS Priority Queues											
0, 1	1											
2, 3	2											
4, 5	3											
6, 7	4											
Step 3	end	Return to privileged EXEC mode.										
Step 4	show wrr-queue cos-map	Display the mapping of the CoS priority queues.										

To disable the new CoS settings and return to default settings, use the **no wrr-queue cos-map** global configuration command.

Configuring WRR Priority

Beginning in privileged EXEC mode, follow these steps to configure the WRR priority:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	wrr-queue bandwidth <i>weight1...weight4</i>	Assign WRR weights to the four CoS queues. The range for the WRR values <i>weight1</i> through <i>weight4</i> is 1 to 255.
Step 3	end	Return to privileged EXEC mode.
Step 4	show wrr-queue bandwidth	Display the WRR bandwidth allocation for the CoS priority queues.

To disable the WRR scheduling and enable the strict priority scheduling, use the **no wrr-queue bandwidth** global configuration command.

To enable one of the queues as the expedite queue and to enable the WRR scheduling for the remaining queues, see the “[Displaying QoS Information](#)” section on page 27-10.

Displaying QoS Information

To display QoS information, use one or more of the privileged EXEC commands in [Table 27-3](#):

Table 27-3 *Commands for Displaying QoS Information*

Command	Purpose
show wrr-queue cos-map	Displays the mapping of the CoS priority queues.
show wrr-queue bandwidth	Displays the WRR bandwidth allocation for the CoS priority queues.