CHAPTER 30

Configuring QoS

This chapter describes how to configure quality of service (QoS) by using automatic-QoS (auto-QoS) commands or by using standard QoS commands. With QoS, you can give preferential treatment to certain types of traffic at the expense of others. Without QoS, the Catalyst 2950 or Catalyst 2955 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.

To use the features described in this chapter, you must have the enhanced software image (EI) installed on your switch.

If you have the standard software image (SI) installed on your switch, you cannot configure some of the features. Table 30-1 lists the sections that describe the features that you can configure.

Table 30-1 Sections Describing Standard Software Features

Topic	Section
Queueing and scheduling at the egress ports	"Queueing and Scheduling" section on page 30-8
Configuring QoS	"Configuring Standard QoS" section on page 30-15
	"Default Standard QoS Configuration" section on page 30-16
	"Configuring Classification Using Port Trust States" section on page 30-18
	"Configuring the Egress Queues" section on page 30-35

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 4, "Getting Started with CMS."

You can also use these wizards to configure QoS only if your switch is running the EI:

- Priority data wizard—Lets you assign priority levels to data applications based on their TCP or UDP ports. It has a standard list of applications, and you select the ones that you want to prioritize, the priority levels, and the interfaces where the prioritization occurs. Refer to the priority data wizard online help for procedures about using this wizard.
- Video wizard—Gives traffic that originates from specified video servers a higher priority than the priority of data traffic. The wizard assumes that the video servers are connected to a single device in the cluster. Refer to the video wizard online help for procedures about using this wizard.

This chapter consists of these sections:

Understanding QoS

This section describes how QoS is implemented on the switch. If you have the SI installed on your switch, some concepts and features in this section might not apply. For a list of available features, see Table 30-1 on page 30-1.

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 select specific network traffic, prioritize it according to its relative importance, and 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 DiffServ architecture, an emerging standard from the Internet Engineering Task Force (IETF). This architecture specifies that each packet is classified upon entry into the network. The classification is carried in the IP packet header, using 6 bits from the deprecated IP type-of-service (ToS) field to carry the classification (*class*) information.

Classification can also be carried in the Layer 2 frame. These special bits in the Layer 2 frame or a Layer 3 packet are described here and shown in Figure 30-1:

- 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.

- Prioritization bits in Layer 3 packets

  Layer 3 IP packets can carry a Differentiated Services Code Point (DSCP) value. The supported DSCP values are 0, 8, 10, 16, 18, 24, 26, 32, 34, 40, 46, 48, and 56.

Figure 30-1 QoS Classification Layers in Frames and Packets

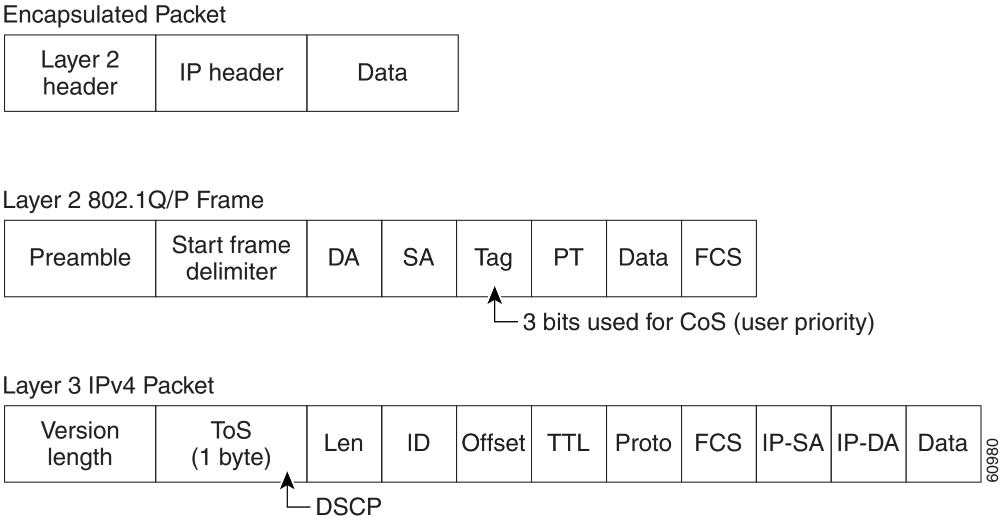

All switches and routers that access the Internet rely on the class information to give the same forwarding treatment to packets with the same class information and different treatment to packets with different class information. The class information in the packet can be assigned by end hosts or by 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 in the DiffServ architecture 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.

Basic QoS Model

Figure 30-2 shows the basic QoS model. Actions at the ingress interface include classifying traffic, policing, and marking:

Note If you have the SI installed on your switch, only the queueing and scheduling features are available.

- Classifying distinguishes one kind of traffic from another. For more information, see the "Classification" section on page 30-5.
- Policing determines whether a packet is in or out of profile according to the configured policer, and the policer limits the bandwidth consumed by a flow of traffic. The result of this determination is passed to the marker. For more information, see the "Policing and Marking" section on page 30-7.
- Marking evaluates the policer and configuration information for the action to be taken when a packet is out of profile and decides what to do with the packet (pass through a packet without modification, mark down the DSCP value in the packet, or drop the packet). For more information, see the "Policing and Marking" section on page 30-7.

Actions at the egress interface include queueing and scheduling:

- Queueing evaluates the CoS value and determines which of the four egress queues in which to place the packet.
- Scheduling services the four egress queues based on their configured weighted round robin (WRR) weights.

Figure 30-2 Basic QoS Model

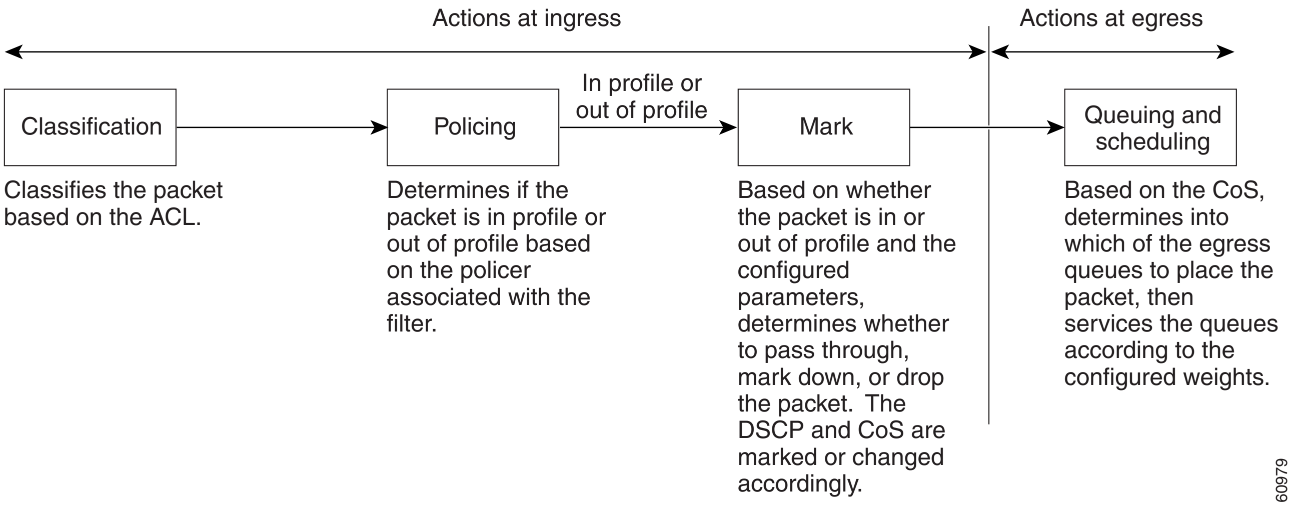
Actions at ingress
Actions at egress
Classification
Policing
In profile or out of profile
Mark
Queuing and scheduling
Classifies the packet based on the ACL.
Determines if the packet is in profile or out of profile based on the policer associated with the filter.
Based on whether the packet is in or out of profile and the configured parameters, determines whether to pass through, mark down, or drop the packet. The DSCP and CoS are marked or changed accordingly.
Based on the CoS, determines into which of the egress queues to place the packet, then services the queues according to the configured weights.
60979

Classification

Note This feature is available only if your switch is running the EI.

Classification is the process of distinguishing one kind of traffic from another by examining the fields in the packet.

Classification occurs only on a physical interface basis. No support exists for classifying packets at the VLAN level.

You specify which fields in the frame or packet that you want to use to classify incoming traffic.

For non-IP traffic, you have these classification options:

- Use the port default. If the frame does not contain a CoS value, the switch assigns the default port CoS value to the incoming frame.
- Trust the CoS value in the incoming frame (configure the port to trust CoS). Layer 2 802.1Q frame headers carry the CoS value in the three most-significant bits of the Tag Control Information field. CoS values range from 0 for low priority to 7 for high priority.

  The trust DSCP configuration is meaningless for non-IP traffic. If you configure a port with this option and non-IP traffic is received, the switch assigns the default port CoS value and classifies traffic based on the CoS value.

For IP traffic, you have these classification options:

- Trust the IP DSCP in the incoming packet (configure the port to trust DSCP). The switch assigns the same DSCP to the packet for internal use. The IETF defines the 6 most-significant bits of the 1-byte ToS field as the DSCP. The priority represented by a particular DSCP value is configurable. The supported DSCP values are 0, 8, 10, 16, 18, 24, 26, 32, 34, 40, 46, 48, and 56.
- Trust the CoS value (if present) in the incoming packet. The switch generates the DSCP by using the CoS-to-DSCP map.

Note An interface can be configured to trust either CoS or DSCP, but not both at the same time.

Classification Based on QoS ACLs

You can use IP standard, IP extended, and Layer 2 MAC access control lists (ACLs) to define a group of packets with the same characteristics (*class*). In the QoS context, the permit and deny actions in the access control entries (ACEs) have different meanings than with security ACLs:

- If a match with a permit action is encountered (first-match principle), the specified QoS-related action is taken.
- If no match with a permit action is encountered and all the ACEs have been examined, no QoS processing occurs on the packet.
- If multiple ACLs are configured on an interface, the packet matches the first ACL with a permit action, and QoS processing begins.

- Configuration of a deny action is not supported in QoS ACLs on the switch.
- System-defined masks are allowed in class maps with these restrictions:
  - A combination of system-defined and user-defined masks cannot be used in the multiple class maps that are a part of a policy map.
  - System-defined masks that are a part of a policy map must all use the same type of system mask. For example, a policy map cannot have a class map that uses the **permit tcp any any** ACE and another that uses the **permit ip any any** ACE.
  - A policy map can contain multiple class maps that all use the same user-defined mask or the same system-defined mask.

Note For more information about system-defined masks, see the "Understanding Access Control Parameters" section on page 29-4.

For more information about ACL restrictions, see the "Configuring ACLs" section on page 29-6.

After a traffic class has been defined with the ACL, you can attach a policy to it. A policy might contain multiple classes with actions specified for each one of them. A policy might include commands to classify the class as a particular aggregate (for example, assign a DSCP) or rate-limit the class. This policy is then attached to a particular port on which it becomes effective.

You implement IP ACLs to classify IP traffic by using the **access-list** global configuration command; you implement Layer 2 MAC ACLs to classify Layer 2 traffic by using the **mac access-list extended** global configuration command.

Classification Based on Class Maps and Policy Maps

A class map is a mechanism that you use to isolate and name a specific traffic flow (or class) from all other traffic. The class map defines the criteria used to match against a specific traffic flow to further classify it; the criteria can include matching the access group defined by the ACL. If you have more than one type of traffic that you want to classify, you can create another class map and use a different name. After a packet is matched against the class-map criteria, you further classify it through the use of a policy map.

A policy map specifies which traffic class to act on. Actions can include setting a specific DSCP value in the traffic class or specifying the traffic bandwidth limitations and the action to take when the traffic is out of profile. Before a policy map can be effective, you must attach it to an interface.

You create a class map by using the **class-map** global configuration command or the **class** policy-map configuration command. You should use the **class-map** global configuration command when the map is shared among many ports. When you enter the **class-map** global configuration command, the switch enters the class-map configuration mode. In this mode, you define the match criterion for the traffic by using the **match** class-map configuration command.

You create and name a policy map by using the **policy-map** global configuration command. When you enter this command, the switch enters the policy-map configuration mode. In this mode, you specify the actions to take on a specific traffic class by using the **class** policy-map configuration or **set** policy-map class configuration command. To make the policy map effective, you attach it to an interface by using the **service-policy** interface configuration command.

The policy map can also contain commands that define the policer, the bandwidth limitations of the traffic, and the action to take if the limits are exceeded. For more information, see the "Policing and Marking" section on page 30-7.

A policy map also has these characteristics:

- A policy map can contain multiple class statements.
- A separate policy-map class can exist for each type of traffic received through an interface.
- A policy-map configuration state supersedes any actions due to an interface trust state.

For configuration information, see the "Configuring a QoS Policy" section on page 30-24.

Policing and Marking

Note This feature is available only if your switch is running the EI.

Policing involves creating a policer that specifies the bandwidth limits for the traffic. Packets that exceed the limits are *out of profile* or *nonconforming*. Each policer specifies the action to take for packets that are in or out of profile. These actions, carried out by the marker, include dropping the packet or marking down the packet with a new user-defined value.

You can create an individual policer. QoS applies the bandwidth limits specified in the policer separately to each matched traffic class. You configure this type of policer within a policy map by using the **policy-map** configuration command.

When configuring policing and policers, keep these items in mind:

- By default, no policers are configured.
- Policers can only be configured on a physical port. There is no support for policing at a VLAN level.
- Only one policer can be applied to a packet in the input direction.
- Only the average rate and committed burst parameters are configurable.
- Policing occurs on the ingress interfaces:
  - 60 policers are supported on ingress Gigabit-capable Ethernet ports.
  - 6 policers are supported on ingress 10/100 Ethernet ports.
  - Granularity for the average burst rate is 1 Mbps for 10/100 ports and 8 Mbps for Gigabit Ethernet ports.
- On an interface configured for QoS, all traffic received through the interface is classified, policed, and marked according to the policy map attached to the interface. On a trunk interface configured for QoS, traffic in *all* VLANs received through the interface is classified, policed, and marked according to the policy map attached to the interface.

Note You cannot configure policers on the egress interfaces.

Mapping Tables

Note This feature is available only if your switch is running the EI.

During classification, QoS uses a configurable CoS-to-DSCP map to derive an internal DSCP value from the received CoS value. This DSCP value represents the priority of the traffic.

Before the traffic reaches the scheduling stage, QoS uses the configurable DSCP-to-CoS map to derive a CoS value from the internal DSCP value. The CoS value is used to select one of the four egress queues.

The CoS-to-DSCP and DSCP-to-CoS maps have default values that might or might not be appropriate for your network.

For configuration information, see the "Configuring CoS Maps" section on page 30-32.

Queueing and Scheduling

Note Both the SI and EI support this feature.

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 assigning priority-indexed CoS values to 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 2950 or Catalyst 2955 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.

Port Scheduling

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.

CoS configures each transmit port (the *egress* port) with a normal-priority transmit queue and a high-priority transmit queue, depending on the frame tag or the port information. Frames in the normal-priority queue are forwarded only after frames in the high-priority queue are forwarded.

The switch (802.1P user priority) has four priority queues. The frames are forwarded to appropriate queues based on the priority-to-queue mapping that you defined.

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 are always sent first, and packets in the low-priority queue are not sent 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.

- Strict priority and WRR scheduling

  Strict priority and WRR scheduling, also referred to as strict priority queueing, uses one of the egress queues as an expedite queue (queue 4). The remaining queues participate in WRR. When the expedite queue is configured, it is a priority queue and is serviced until it is empty before the other queues are serviced by WRR scheduling.

  You can enable the egress expedite queue and assign WRR weights to the other queues by using the **wrr-queue bandwidth** *weight1 weight2 weight3* **0** global configuration command.

Configuring Auto-QoS

Note This feature is available only if your switch is running the EI.

You can use the auto-QoS feature to simplify the deployment of existing QoS features. Auto-QoS makes assumptions about the network design, and as a result, the switch can prioritize different traffic flows and appropriately use the egress queues instead of using the default QoS behavior (the switch offers best-effort service to each packet regardless of the packet contents or size and sends it from a single queue).

When you enable auto-QoS, it automatically classifies traffic based on the traffic type and ingress packet label. The switch uses the resulting classification to choose the appropriate egress queue.

You use auto-QoS commands to identify ports connected to Cisco IP Phones and to identify ports that receive trusted voice over IP (VoIP) traffic through an uplink. Auto-QoS then performs these functions:

- Detects the presence or absence of IP phones
- Configures QoS classification
- Configures egress queues

These sections describe how to configure auto-QoS on your switch:

- Generated Auto-QoS Configuration, page 30-10
- Effects of Auto-QoS on the Configuration, page 30-12
- Configuration Guidelines, page 30-12
- Enabling Auto-QoS for VoIP, page 30-12

Generated Auto-QoS Configuration

When auto-QoS is enabled, it uses the ingress packet label to classify traffic and to configure the egress queues as described in Table 30-2.

Table 30-2 Traffic Types, Ingress Packet Labels, Assigned Packet Labels, and Egress Queues

	VoIP Data Traffic Only From Cisco IP Phones	**VoIP Control Traffic Only From Cisco IP Phones**	**Routing Protocol Traffic**	**STP BPDU[1] Traffic**	**All Other Traffic**
Ingress DSCP	46	26	–	–	–
Ingress CoS	5	3	6	7	–
Assigned DSCP	46	26	48	56	0
Assigned CoS	5	3	6	7	0
CoS-to-Queue Map	5	3, 6, 7			0, 1, 2, 4
Egress Queue	Expedite queue	80% WRR			20% WRR

[1] BPDU = bridge protocol data unit

Table 30-3 lists the generated auto-QoS configuration for the egress queues.

Table 30-3 Auto-QoS Configuration for the Egress Queues

Egress Queue	**Queue Number**	**CoS-to-Queue Map**	**Queue Weight**
Expedite	4	5	–
80% WRR	3	3, 6, 7	80%
20% WRR	1	0, 1, 2, 4	20%

When you enable the auto-QoS feature on the first interface, these automatic actions occur:

- When you enter the **auto qos voip trust** interface configuration command, the ingress classification on the interface is set to trust the QoS label received in the packet, and the egress queues on the interface are reconfigured (see Table 30-3).
- When you enter the **auto qos voip cisco-phone** interface configuration command, the trusted boundary feature is enabled. It uses the Cisco Discovery Protocol (CDP) to detect the presence or absence of a Cisco IP Phone. When a Cisco IP Phone is detected, the ingress classification on the interface is set to trust the QoS label received in the packet. When a Cisco IP Phone is absent, the ingress classification is set to not trust the QoS label in the packet. The egress queues on the interface are reconfigured (see Table 30-3).

  For information about the trusted boundary feature, see the "Configuring Trusted Boundary" section on page 30-21.
- The switch automatically assigns egress queue usage as shown inTable 30-3.

When you enable auto-QoS by using the **auto qos voip cisco-phone** or the **auto qos voip trust** interface configuration command, the switch automatically generates a QoS configuration based on the traffic type and ingress packet label and applies the commands listed in Table 30-4 to the interface.

Table 30-4 Generated Auto-QoS Configuration Command Equivalents

Description	Automatically Generated QoS Command Equivalent
The switch automatically enables standard QoS and configures the CoS-to-DSCP map (maps CoS values in incoming packets to a DSCP value) as shown in Table 30-2 on page 30-10.	`Switch(config)# mls qos map cos-dscp 0 8 16 26 32 46 48 56`
The switch automatically sets the ingress classification on the interface to trust the CoS value received in the packet.	`Switch(config-if)# mls qos trust cos`
If you entered the **auto qos voip cisco-phone** command, the switch automatically enables the trusted boundary feature, which uses the CDP to detect the presence or absence of a Cisco IP Phone.	`Switch(config-if)# mls qos trust device cisco-phone`

The switch automatically assigns egress queue usage (as shown in Table 30-3 on page 30-10) on this interface.
The switch enables the egress expedite queue and assigns WRR weights to queues 1 and 3. (The lowest value for a WRR queue is 1. When the WRR weight of a queue is set to 0, this queue becomes an expedite queue.)
The switch configures the CoS-to-egress-queue map:
• CoS values 0, 1, 2, and 4 select queue 1.
• CoS values 3, 6, and 7 select queue 3.
• CoS value 5 selects queue 4 (expedite queue).
Because the expedite queue (queue 4) contains the VoIP data traffic, the queue is serviced until empty.
	`Switch(config)# wrr-queue bandwidth 20 1 80 0`
	`Switch(config)# wrr-queue cos-map 1 0 1 2 4`
	`Switch(config)# wrr-queue cos-map 3 3 6 7`
	`Switch(config)# wrr-queue cos-map 4 5`

Effects of Auto-QoS on the Configuration

When auto-QoS is enabled, the **auto qos voip** interface configuration command and the generated configuration are added to the running configuration.

Configuration Guidelines

Before configuring auto-QoS, you should be aware of this information:

- In this release, auto-QoS configures the switch only for VoIP with Cisco IP Phones.
- To take advantage of the auto-QoS defaults, do not configure any standard QoS commands before entering the auto-QoS commands. If necessary, you can fine-tune the QoS configuration, but we recommend that you do so only after the auto-QoS configuration is completed.
- You can enable auto-QoS on static, dynamic-access, voice VLAN access, and trunk ports.
- By default, the CDP is enabled on all interfaces. For auto-QoS to function properly, do not disable the CDP.
- Policing is not enabled in auto-QoS. You can manually enable policing, as described in the "Configuring a QoS Policy" section on page 30-24.

Enabling Auto-QoS for VoIP

Beginning in privileged EXEC mode, follow these steps to enable auto-QoS for VoIP within a QoS domain:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**interface** *interface-id*	Enter interface configuration mode, and specify the interface that is connected to a Cisco IP Phone. You also can specify the uplink interface that is connected to another switch or router in the interior of the network.
Step 3	**auto qos voip** {**cisco-phone** | **trust**}	Enable auto-QoS. The keywords have these meanings: • **cisco-phone**—If the interface is connected to a Cisco IP Phone, the QoS labels of incoming packets are trusted only when the IP phone is detected. • **trust**—The uplink interface is connected to a trusted switch or router, and the VoIP classification in the ingress packet is trusted.
Step 4	**end**	Return to privileged EXEC mode.
Step 5	**show auto qos interface** *interface-id*	Verify your entries. This command displays the auto-QoS configuration that was initially applied; it does not display any user changes to the configuration that might be in effect.

To disable auto-QoS on the switch and return to the default port trust state set (untrusted), follow these steps:

1. Use the **no auto qos voip** interface configuration command on all interfaces on which auto-QoS is enabled. To disable auto-QoS on multiple interfaces at the same time, you can use the **interface range** global configuration command.
2. After disabling auto-QoS on all interfaces on which auto-QoS was enabled, return the egress queues and CoS-to-DSCP map to the default settings by using these global configuration commands:
   - **no wrr-queue bandwidth**
   - **no wrr-queue cos-map**
   - **no mls qos map cos-dscp**

To display the QoS commands that are automatically generated when auto-QoS is enabled or disabled, enter the **debug autoqos** privileged EXEC command before enabling auto-QoS. For more information, see the "Using the debug autoqos Command" section on page 32-21.

This example shows how to enable auto-QoS and to trust the QoS labels in incoming packets when the device connected to Fast Ethernet interface 0/1 is detected as a Cisco IP Phone:

```
Switch(config)# interface fastethernet0/1
Switch(config-if)# auto qos voip cisco-phone
```

This example shows how to enable auto-QoS and to trust the QoS labels in incoming packets when the switch or router connected to Gigabit Ethernet interface 0/1 is a trusted device:

```
Switch(config)# interface gigabitethernet0/1
Switch(config-if)# auto qos voip trust
```

Displaying Auto-QoS Information

To display the initial auto-QoS configuration, use the **show auto qos** [**interface** [*interface-id*]] privileged EXEC command. To display any user changes to that configuration, use the **show running-config** privileged EXEC command. You can compare the **show auto qos** and the **show running-config** command output to identify the user-defined QoS settings.

To display information about the QoS configuration that might be affected by auto-QoS, use one of these commands:

- **show mls qos**
- **show mls qos map cos-dscp**
- **show wrr-queue bandwidth**
- **show wrr-queue cos-map**

For more information about these commands, refer to the command reference for this release.

Auto-QoS Configuration Example

Note This example is applicable only if your switch is running the EI.

This section describes how you could implement auto-QoS in a network, as shown in Figure 30-3.

Figure 30-3 Auto-QoS Configuration Example Network

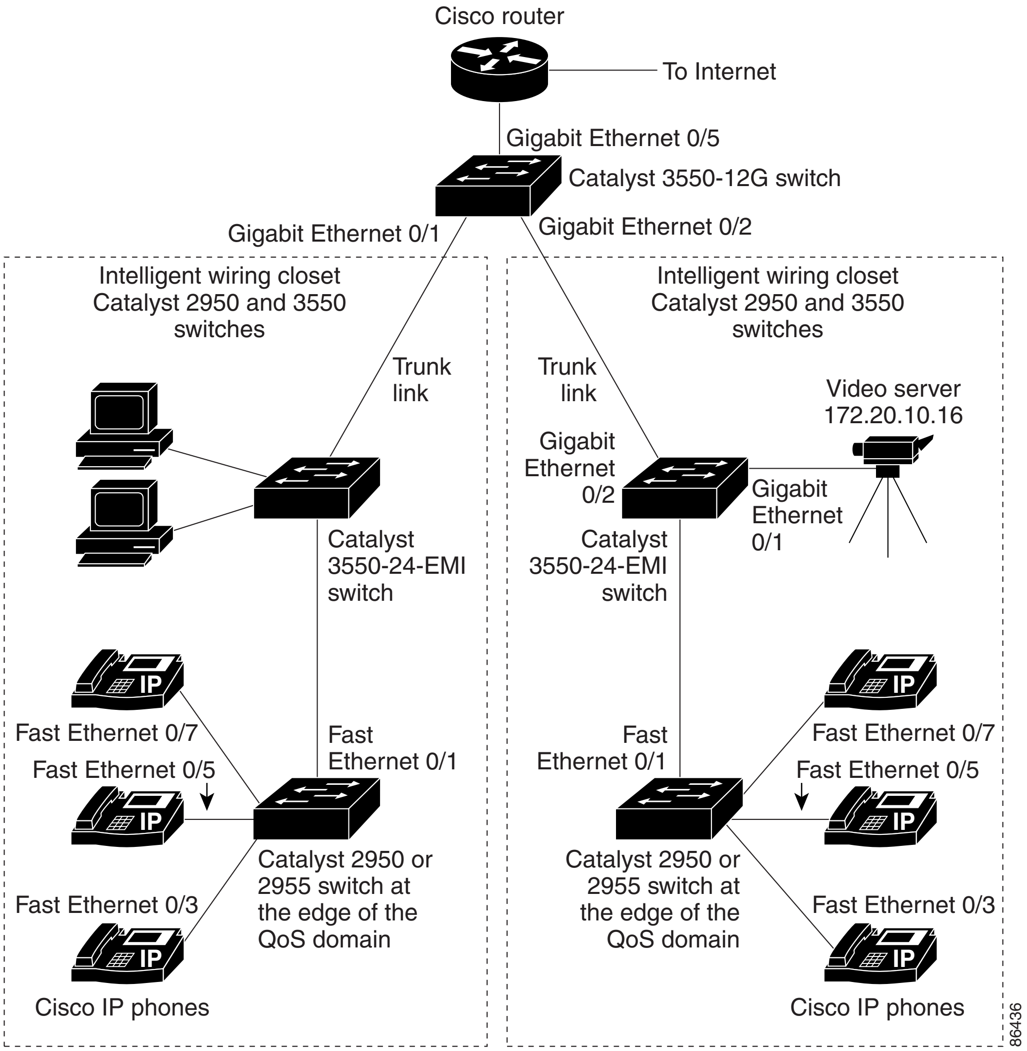

The intelligent wiring closets in Figure 30-3 are composed of Catalyst 2950 switches running the EI and Catalyst 3550 switches. The object of this example is to prioritize the VoIP traffic over all other traffic. To do so, enable auto-QoS on the switches at the edge of the QoS domains in the wiring closets.

Note You should not configure any standard-QoS commands before entering the auto-QoS commands. You can fine-tune the QoS configuration, but we recommend that you do so only after the auto-QoS configuration is completed.

Beginning in privileged EXEC mode, follow these steps to configure the switch at the edge of the QoS domain to prioritize the VoIP traffic over all other traffic:

	Command	Purpose
Step 1	**debug autoqos**	Enable debugging for auto-QoS. When debugging is enabled, the switch displays the QoS configuration that is automatically generated when auto-QoS is enabled.
Step 2	**configure terminal**	Enter global configuration mode.
Step 3	**cdp enable**	Enable CDP globally. By default, it is enabled.
Step 4	**interface fastethernet0/3**	Enter interface configuration mode.
Step 5	**auto qos voip cisco-phone**	Enable auto-QoS on the interface, and specify that the interface is connected to a Cisco IP Phone. The QoS labels of incoming packets are trusted only when the IP phone is detected.
Step 6	**interface fastethernet0/5**	Enter interface configuration mode.
Step 7	**auto qos voip cisco-phone**	Enable auto-QoS on the interface, and specify that the interface is connected to a Cisco IP Phone.
Step 8	**interface fastethernet0/7**	Enter interface configuration mode.
Step 9	**auto qos voip cisco-phone**	Enable auto-QoS on the interface, and specify that the interface is connected to a Cisco IP Phone.
Step 10	**interface ethernet0/1**	Enter interface configuration mode.
Step 11	**auto qos voip trust**	Enable auto-QoS on the interface, and specify that the interface is connected to a trusted router or switch.
Step 12	**end**	Return to privileged EXEC mode.
Step 13	**show auto qos**	Verify your entries. This command displays the auto-QoS configuration that is initially applied; it does not display any user changes to the configuration that might be in effect. For information about the QoS configuration that might be affected by auto-QoS, see the "Displaying Auto-QoS Information" section on page 26-12.
Step 14	**copy running-config startup-config**	Save the **auto qos voip** interface configuration commands and the generated auto-QoS configuration in the configuration file.

Configuring Standard QoS

Before configuring standard 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 standard QoS on your switch:

Note If your switch is running the SI, you can configure only the features described in the "Configuring Classification Using Port Trust States" and the "Configuring the Egress Queues" sections. You can also display the QoS information as described in the "Displaying Standard QoS Information" section.

- Default Standard QoS Configuration, page 30-16
- Configuration Guidelines, page 30-17
- Configuring Classification Using Port Trust States, page 30-18
- Configuring a QoS Policy, page 30-24
- Configuring CoS Maps, page 30-32
- Configuring the Egress Queues, page 30-35

Default Standard QoS Configuration

This is the default standard QoS configuration:

Note You can configure policy maps, policers, the CoS-to-DSCP map, and the DSCP-to-CoS map only if your switch is running the EI.

- The default port CoS value is 0.
- The CoS value of 0 is assigned to all incoming packets.
- The default port trust state is untrusted.
- No policy maps are configured.
- No policers are configured.
- The default CoS-to-DSCP map is shown in Table 30-7.
- The default DSCP-to-CoS map is shown in Table 30-8.
- The default scheduling method for the egress queues is strict priority.
- For default CoS and WRR values, see the "Configuring the Egress Queues" section on page 30-35.

Note In software releases earlier than Cisco IOS Release 12.1(11)EA1, the switch uses the CoS value of incoming packets without modifying the DSCP value. You can configure this by enabling pass-through mode on the port. For more information, see the "Enabling Pass-Through Mode" section on page 30-23.

Configuration Guidelines

Note These guidelines are applicable only if your switch is running the EI.

Before beginning the QoS configuration, you should be aware of this information:

- You must disable the IEEE 802.3X flowcontrol on all ports before enabling QoS on the switch. To disable it, use the **flowcontrol receive off** and **flowcontrol send off** interface configuration commands.
- If you have EtherChannel ports configured on your switch, you must configure QoS classification, policing, mapping, and queueing on the individual physical ports that comprise the EtherChannel. You must decide whether the QoS configuration should match on all ports in the EtherChannel.
- It is not possible to match IP fragments against configured IP extended ACLs to enforce QoS. IP fragments are sent as best-effort traffic. IP fragments are denoted by fields in the IP header.
- All ingress QoS processing actions apply to control traffic (such as spanning-tree bridge protocol data units [BPDUs] and routing update packets) that the switch receives.
- Only an ACL that is created for physical interfaces can be attached to a class map.
- Only one ACL per class map and only one **match** command per class map are supported. The ACL can have multiple access control entries, which are commands that match fields against the contents of the packet.
- Policy maps with ACL classification in the egress direction are not supported and cannot be attached to an interface by using the **service-policy input** *policy-map-name* interface configuration command.
- In a policy map, the class named *class-default* is not supported. The switch does not filter traffic based on the policy map defined by the **class class-default** policy-map configuration command.
- For more information about guidelines for configuring ACLs, see the "Classification Based on QoS ACLs" section on page 30-5.
- For information about applying ACLs to physical interfaces, see the "Guidelines for Applying ACLs to Physical Interfaces" section on page 29-6.
- If a policy map with a system-defined mask and a security ACL with a user-defined mask are configured on an interface, the switch might ignore the actions specified by the policy map and perform only the actions specified by the ACL. For information about masks, see the "Understanding Access Control Parameters" section on page 29-4.
- If a policy map with a user-defined mask and a security ACL with a user-defined mask are configured on an interface, the switch takes one of the actions as described in Table 30-5. For information about masks, see the "Understanding Access Control Parameters" section on page 29-4.

Table 30-5 Interaction Between Policy Maps and Security ACLs

Policy-Map Conditions	Security-ACL Conditions	Action
When the packet is in profile.	Permit specified packets.	Traffic is forwarded.
When the packet is out of profile and the out-of-profile action is to mark down the DSCP value.	Drop specified packets.	Traffic is dropped.
When the packet is out of profile and the out-of-profile action is to drop the packet.	Permit specified packets.	Traffic is dropped.
	Drop specified packets.	Traffic is dropped.

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 30-18
- Configuring the CoS Value for an Interface, page 30-21
- Configuring Trusted Boundary, page 30-21
- Enabling Pass-Through Mode, page 30-23

Note Both the SI and EI support this feature.

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 30-4 shows a sample network topology.

Figure 30-4 Port Trusted States within the QoS Domain

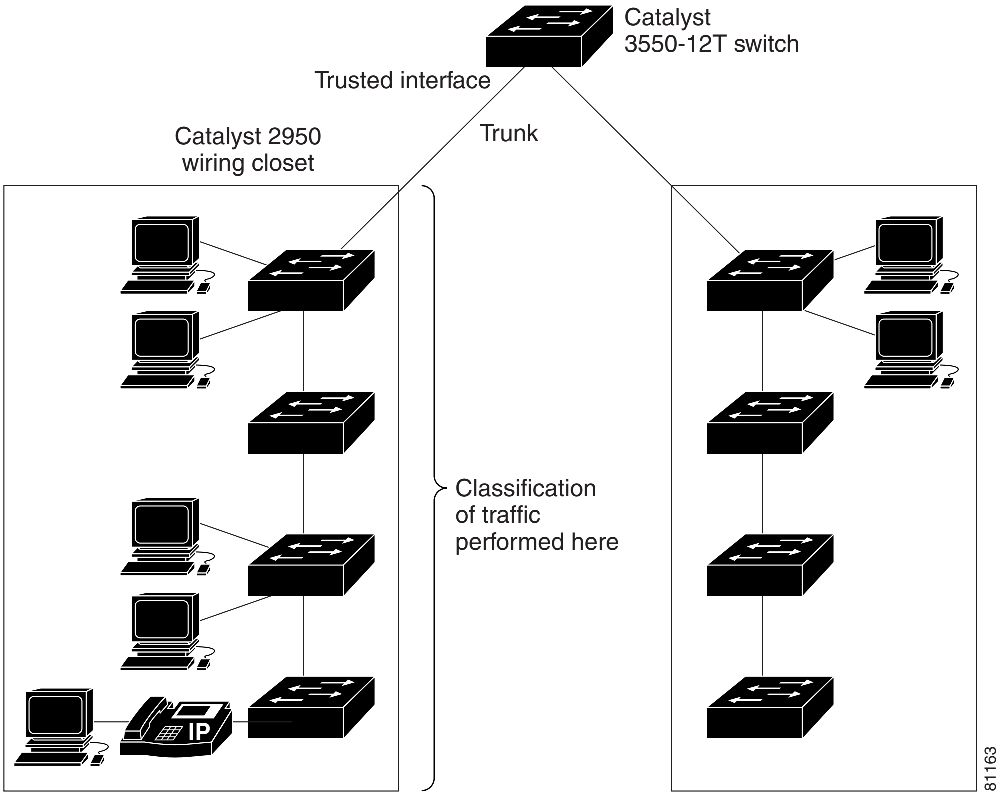

Beginning in privileged EXEC mode, follow these steps to configure the port to trust the classification of the traffic that it receives:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**interface** *interface-id*	Enter interface configuration mode, and specify the interface to be trusted. Valid interfaces include physical interfaces.
Step 3	**mls qos trust** [**cos** | **dscp**]	Configure the port trust state. By default, the port is not trusted. The keywords have these meanings: **cos**—Classifies ingress packets with the packet CoS values. For tagged IP packets, the DSCP value of the packet is modified based on the CoS-to-DSCP map. The egress queue assigned to the packet is based on the packet CoS value. **dscp**—Classifies ingress packets with packet DSCP values. For non-IP packets, the packet CoS value is set to 0 for tagged packets; the default port CoS is used for untagged packets. Internally, the switch modifies the CoS value by using the DSCP-to-CoS map. This keyword is available only if your switch is running the EI. **Note** In software releases earlier than Cisco IOS Release 12.1(11)EA1, the **mls qos trust** command is available only when the switch is running the EI. Use the **cos** keyword if your network is composed of Ethernet LANs. Use the **dscp** keyword if your network is not composed of only Ethernet LANs and if you are familiar with sophisticated QoS features and implementations. For more information about this command, refer to the command reference for this release.
Step 4	**end**	Return to privileged EXEC mode.
Step 5	**show mls qos interface** [*interface-id*] [**policers**]	Verify your entries.
Step 6	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To return a port to its untrusted state, use the **no mls qos trust** interface configuration command.

For information on how to change the default CoS value, see the "Configuring the CoS Value for an Interface" section on page 30-21. For information on how to configure the CoS-to-DSCP map, see the "Configuring the CoS-to-DSCP Map" section on page 30-33.

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	**configure terminal**	Enter global configuration mode.
Step 2	**interface** *interface-id*	Enter interface configuration mode, and specify the interface to be trusted. Valid interfaces include physical interfaces.
Step 3	**mls qos cos** {*default-cos* | **override**}	Configure the default CoS value for the port. • For *default-cos*, 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. Use the **override** keyword when all incoming packets on certain ports deserve higher priority than packets entering from other ports. Even if a port was previously set to trust DSCP, this command overrides the previously configured trust state, and all the incoming CoS values are assigned the default CoS value configured with this command. 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.
Step 4	**end**	Return to privileged EXEC mode.
Step 5	**show mls qos interface**	Verify your entries.
Step 6	**copy running-config startup-config**	(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 30-4 on page 30-19. 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** *interface-id*	Enter interface configuration mode, and specify the interface to be trusted. 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. You cannot enable both trusted boundary and auto-QoS (**auto qos voip** interface configuration command) at the same time; they are mutually exclusive.
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. **Note** In software releases earlier than Cisco IOS Release 12.1(11)EA1, the **mls qos trust cos** command is available only when the switch is running the EI. 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** [*interface-id*] [**policers**]	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.

You cannot enable trusted boundary if auto-QoS is already enabled and vice-versa. If auto-QoS is enabled and a Cisco IP Phone is absent on a port, the port does not trust the classification of traffic that it receives.

Table 30-6 lists the port configuration when an IP phone is present or absent.

Table 30-6 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 trusts the DSCP value of the incoming packet.	The packet DSCP value is trusted.	For tagged non-IP packets, the packet CoS value is set to 0. For untagged non-IP packets, 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

In software releases earlier than Cisco IOS Release 12.1(11)EA1, 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. You cannot enable or disable pass-through mode if your switch is running a software release earlier than Cisco IOS Release 12.1(11)EA1.

In Cisco IOS Release 12.1(11)EA1 or later, 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** *interface-id*	Enter interface configuration mode, and specify the interface on which pass-through mode is enabled. 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** [*interface-id*]	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** | **dscp**] 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** | **dscp**] interface commands are already configured, pass-through mode is disabled.

Configuring a QoS Policy

Note This feature is available only if your switch is running the EI.

Configuring a QoS policy typically requires classifying traffic into classes, configuring policies applied to those traffic classes, and attaching policies to interfaces.

For background information, see the "Classification" section on page 30-5 and the "Policing and Marking" section on page 30-7.

This section contains this configuration information:

- Classifying Traffic by Using ACLs, page 30-25
- Classifying Traffic by Using Class Maps, page 30-28
- Classifying, Policing, and Marking Traffic by Using Policy Maps, page 30-29

Classifying Traffic by Using ACLs

You can classify IP traffic by using IP standard or IP extended ACLs; you can classify Layer 2 traffic by using Layer 2 MAC ACLs.

Beginning in privileged EXEC mode, follow these steps to create an IP standard ACL for IP traffic:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**access-list** *access-list-number* {**permit** | **remark**} {*source source-wildcard* | **host** *source* | **any**}	Create an IP standard ACL, repeating the command as many times as necessary. For *access-list-number*, enter the ACL number. The range is 1 to 99 and 1300 to 1999. Enter **permit** to specify whether to permit access if conditions are matched. Enter **remark** to specify an ACL entry comment up to 100 characters. **Note** Deny statements are not supported for QoS ACLs. See the "Classification Based on QoS ACLs" section on page 30-5 for more details. The *source* is the source address of the network or host from which the packet is being sent, specified in one of three ways: • The 32-bit quantity in dotted decimal format. • The keyword **any** as an abbreviation for *source* and *source-wildcard* of 0.0.0.0 255.255.255.255. You do not need to enter a source wildcard. • The keyword **host** as an abbreviation for *source* and *source-wildcard* of *source* 0.0.0.0. (Optional) The *source-wildcard* variable applies wildcard bits to the source (see first bullet item).
Step 3	**end**	Return to privileged EXEC mode.
Step 4	**show access-lists**	Verify your entries.
Step 5	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

For more information about creating IP standard ACLs, see the "Guidelines for Applying ACLs to Physical Interfaces" section on page 29-6.

To delete an ACL, use the **no access-list** *access-list-number* global configuration command.

This example shows how to allow access for only those hosts on the two specified networks. The wildcard bits apply to the host portions of the network addresses. Any host with a source address that does not match the ACL statements is rejected.

```
Switch(config)# access-list 1 permit 192.5.255.0 0.0.0.255
Switch(config)# access-list 1 permit 36.0.0.0 0.0.0.255
```

Beginning in privileged EXEC mode, follow these steps to create an IP extended ACL for IP traffic:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**access-list** *access-list-number* {**permit | remark**} *protocol* {*source source-wildcard* | **host** *source* | **any}** **[operator** *port*] {*destination destination-wildcard* | **host** *destination* | **any**} [*operator port*] [**dscp** *dscp-value*] [**time-range** *time-range-name*]	Create an IP extended ACL, repeating the command as many times as necessary. For *access-list-number*, enter the ACL number. The range is 100 to 199 and 2000 to 2699. Enter **permit** to permit access if conditions are matched. Enter **remark** to specify an ACL entry comment up to 100 characters. **Note** Deny statements are not supported for QoS ACLs. See the "Classification Based on QoS ACLs" section on page 30-5 for more details. For *protocol*, enter the name or number of an IP protocol. Use the question mark (**?**) to see a list of available protocol keywords. For *source*, enter the network or host from which the packet is being sent. For *source-wildcard*, enter the wildcard bits by placing ones in the bit positions that you want to ignore. You specify the *source* and *source-wilcard* by using dotted decimal notation, by using the **any** keyword as an abbreviation for *source* 0.0.0.0 *source-wildcard* 255.255.255.255, or by using the **host** keyword for *source* 0.0.0.0. For *destination*, enter the network or host to which the packet is being sent. You have the same options for specifying the *destination* and *destination-wildcard* as those described for *source* and *source-wildcard*. Define a destination or source port. • The *operator* can be only **eq** (equal). • If *operator* is after *source source-wildcard*, conditions match when the source port matches the defined port. • If *operator* is after *destination destination-wildcard*, conditions match when the destination port matches the defined port. • The *port* is a decimal number or name of a TCP or UDP port. The number can be from 0 to 65535. • Use TCP port names only for TCP traffic. • Use UDP port names only for UDP traffic. Enter **dscp** to match packets with any of the 13 supported DSCP values (0, 8, 10, 16, 18, 24, 26, 32, 34, 40, 46, 48, and 56) or use the question mark (**?**) to see a list of available values. The **time-range** keyword is optional. For information about this keyword, see the "Applying Time Ranges to ACLs" section on page 29-15.
Step 3	**end**	Return to privileged EXEC mode.

	Command	Purpose
Step 4	**show access-lists**	Verify your entries.
Step 5	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

For more information about creating IP extended ACLs, see the "Guidelines for Applying ACLs to Physical Interfaces" section on page 29-6.

To delete an ACL, use the **no access-list** *access-list-number* global configuration command.

This example shows how to create an ACL that permits only TCP traffic from the destination IP address 128.88.1.2 with TCP port number 25:

```
Switch(config)# access-list 102 permit tcp 0.0.0.0 255.255.255.255 128.88.1.2 0.0.0.0 eq
25
```

Beginning in privileged EXEC mode, follow these steps to create a Layer 2 MAC ACL for Layer 2 traffic:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**mac access-list extended** *name*	Create a Layer 2 MAC ACL by specifying the name of the list. After entering this command, the mode changes to extended MAC ACL configuration.
Step 3	**permit** {**any** | **host** *source MAC address*} {**any** | **host** *destination MAC address*} [**aarp** | **amber** | **appletalk** | **dec-spanning** | **decnet-iv** | **diagnostic** | **dsm** | **etype-6000** | **etype-8042** | **lat** | **lavc-sca** | **mop-console** | **mop-dump** | **msdos** | **mumps** | **netbios** | **vines-echo** |**vines-ip** | **xns-idp**]	Enter **permit** to permit access if conditions are matched. **Note** Deny statements are not supported for QoS ACLs. See the "Classification Based on QoS ACLs" section on page 30-5 for more details. For *source MAC address*, enter the MAC address of the host from which the packet is being sent. You specify this by using the **any** keyword to deny any source MAC address or by using the **host** keyword and the source in the hexadecimal format (H.H.H). For *destination MAC address*, enter the MAC address of the host to which the packet is being sent. You specify this by using the **any** keyword to deny any destination MAC address or by using the **host** keyword and the destination in the hexadecimal format (H.H.H). (Optional) You can also enter these options: **aarp** | **amber** | **appletalk** | **dec-spanning** | **decnet-iv** | **diagnostic** | **dsm** | **etype-6000** | **etype-8042** | **lat** | **lavc-sca** | **mop-console** | **mop-dump** | **msdos** | **mumps** | **netbios** | **vines-echo** |**vines-ip** | **xns-idp** (a non-IP protocol).
Step 4	**end**	Return to privileged EXEC mode.
Step 5	**show access-lists** [*number* | *name*]	Verify your entries.
Step 6	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

For more information about creating MAC extended ACLs, see the "Creating Named MAC Extended ACLs" section on page 29-18.

To delete an ACL, use the **no mac access-list extended** *name* global configuration command.

This example shows how to create a Layer 2 MAC ACL with a permit statement. The statement allows traffic from the host with MAC address 0001.0000.0001 to the host with MAC address 0002.0000.0001.

```
Switch(config)# mac access-list extended maclist1
Switch(config-ext-macl)# permit host 0001.0000.0001 host 0002.0000.0001
```

Classifying Traffic by Using Class Maps

You use the **class-map** global configuration command to isolate a specific traffic flow (or class) from all other traffic and to name it. The class map defines the criteria to use to match against a specific traffic flow to further classify it. Match statements can only include ACLs. The match criterion is defined with one match statement entered within the class-map configuration mode.

Note You can also create class maps during policy map creation by using the **class** policy-map configuration command. For more information, see the "Classifying, Policing, and Marking Traffic by Using Policy Maps" section on page 30-29.

Beginning in privileged EXEC mode, follow these steps to create a class map and to define the match criterion to classify traffic:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**access-list** *access-list-number* **permit** {*source source-wildcard* | **host** *source* | **any**} or **access-list** *access-list-number* {**permit** | **remark**} *protocol* {*source source-wildcard* | **host** *source* | **any**} [*operator port*] {*destination destination-wildcard* | **host** *destination* | **any**} [*operator port*] [**dscp** *dscp-value*] [**time-range** *time-range-name*] or **mac access-list extended** *name* **permit** {**any** | **host** *source MAC address*} {**any** | **host** *destination MAC address*} [**aarp** | **amber** | **dec-spanning** | **decnet-iv** | **diagnostic** | **dsm** | **etype-6000** | **etype-8042** | **lat** | **lavc-sca** | **mop-console** | **mop-dump** | **msdos** | **mumps** | **netbios** | **vines-echo** |**vines-ip** | **xns-idp**]	Create an IP standard or extended ACL for IP traffic or a Layer 2 MAC ACL for non-IP traffic, repeating the command as many times as necessary. For more information, see the "Guidelines for Applying ACLs to Physical Interfaces" section on page 29-6 and the "Classifying Traffic by Using ACLs" section on page 30-25. For more information on the **mac access-list extended** *name* command, see the "Creating Named MAC Extended ACLs" section on page 29-18. **Note** Deny statements are not supported for QoS ACLs. See the "Classification Based on QoS ACLs" section on page 30-5 for more details.
Step 3	**class-map** *class-map-name*	Create a class map, and enter class-map configuration mode. By default, no class maps are defined. For *class-map-name,* specify the name of the class map.

	Command	Purpose
Step 4	**match** {**access-group** *acl-index* | **access-group name** *acl-name* | **ip dscp** *dscp-list*}	Define the match criterion to classify traffic. By default, no match criterion is supported. Only one match criterion per class map is supported, and only one ACL per class map is supported. For **access-group** *acl-index* or **access-group name** *acl-name*, specify the number or name of the ACL created in Step 3. For **ip dscp** *dscp-list*, enter a list of up to eight IP DSCP values for each match statement to match against incoming packets. Separate each value with a space. The supported DSCP values are 0, 8, 10, 16, 18, 24, 26, 32, 34, 40, 46, 48, and 56.
Step 5	**end**	Return to privileged EXEC mode.
Step 6	**show class-map** [*class-map-name*]	Verify your entries.
Step 7	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To delete an existing class map, use the **no class-map** *class-map-name* global configuration command. To remove a match criterion, use the **no match** {**access-group** *acl-index* | **name** *acl-name* | **ip dscp**} class-map configuration command.

This example shows how to configure the class map called *class1*. The *class1* has one match criterion, which is an ACL called *103*.

```
Switch(config)# access-list 103 permit any any tcp eq 80
Switch(config)# class-map class1
Switch(config-cmap)# match access-group 103
Switch(config-cmap)# end
Switch#
```

Classifying, Policing, and Marking Traffic by Using Policy Maps

A policy map specifies which traffic class to act on. Actions can include setting a specific DSCP value in the traffic class and specifying the traffic bandwidth limitations for each matched traffic class (policer) and the action to take when the traffic is out of profile (marking or dropping).

A policy map also has these characteristics:

- A policy map can contain multiple class statements, each with different match criteria and policers.
- A separate policy-map class can exist for each type of traffic received through an interface.

You can attach only one policy map per interface in the input direction.

Beginning in privileged EXEC mode, follow these steps to create a policy map:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**access-list** *access-list-number* **permit** {*source source-wildcard* | **host** *source* | **any**} or **access-list** *access-list-number* {**permit** | **remark**} *protocol* {*source source-wildcard* | **host** *source* | **any**} [*operator port*] {*destination destination-wildcard* | **host** *destination* | **any}** [*operator port] [dscp dscp-value*] [**time-range** *time-range-name*] or **mac access-list extended** *name* **permit** {**any** | **host** *source MAC address*} {**any** | **host** *destination MAC address*} [**aarp** | **amber** | **appletalk** |**dec-spanning** | **decnet-iv** | **diagnostic** | **dsm** | **etype-6000** | **etype-8042** | **lat** | **lavc-sca** | **mop-console** | **mop-dump** | **msdos** | **mumps** | **netbios** | **vines-echo** |**vines-ip** | **xns-idp**]	Create an IP standard or extended ACL for IP traffic or a Layer 2 MAC ACL for non-IP traffic, repeating the command as many times as necessary. For more information, see the "Classifying Traffic by Using ACLs" section on page 30-25. **Note** Deny statements are not supported for QoS ACLs. See the "Classification Based on QoS ACLs" section on page 30-5 for more details. For more information on the **mac access-list extended** *name* command, see the "Creating Named MAC Extended ACLs" section on page 29-18.
Step 3	**policy-map** *policy-map-name*	Create a policy map by entering the policy map name, and enter policy-map configuration mode. By default, no policy maps are defined. The default behavior of a policy map is to set the DSCP to 0 if the packet is an IP packet and to set the CoS to 0 if the packet is tagged. No policing is performed.
Step 4	**class** *class-map-name* [**access-group name** *acl-index-or-name*]	Define a traffic classification, and enter policy-map class configuration mode. By default, no policy map class maps are defined. If a traffic class has already been defined by using the **class-map** global configuration command, specify its name for *class-map-name* in this command. For **access-group name** *acl-index-or-name*, specify the number or name of the ACL created in Step 2. **Note** In a policy map, the class named *class-default* is not supported. The switch does not filter traffic based on the policy map defined by the **class class-default** policy-map configuration command.

	Command	Purpose
Step 5	**set** {**ip dscp** *new-dscp*}	Classify IP traffic by setting a new value in the packet. For **ip dscp** *new-dscp*, enter a new DSCP value to be assigned to the classified traffic. The supported DSCP values are 0, 8, 10, 16, 18, 24, 26, 32, 34, 40, 46, 48, and 56.
Step 6	**police** *rate-bps burst-byte* [**exceed-action** {**drop** | **dscp** *dscp-value*}]	Define a policer for the classified traffic. You can configure up to 60 policers on ingress Gigabit-capable Ethernet ports and up to 6 policers on ingress 10/100 Ethernet ports. For *rate-bps*, specify average traffic rate in bits per second (bps). The range is 1 Mbps to 100 Mbps for 10/100 Ethernet ports and 8 Mbps to 1000 Mbps for the Gigabit-capable Ethernet ports. For *burst-byte*, specify the normal burst size in bytes. The values supported on the 10/100 ports are 4096, 8192, 16384, 32768, and 65536. The values supported on the Gigabit-capable Ethernet ports are 4096, 8192, 16348, 32768, 65536, 131072, 262144, and 524288. (Optional) Specify the action to take when the rates are exceeded. Use the **exceed-action drop** keywords to drop the packet. Use the **exceed-action dscp** *dscp-value* keywords to mark down the DSCP value and send the packet.
Step 7	**exit**	Return to policy-map configuration mode.
Step 8	**exit**	Return to global configuration mode.
Step 9	**interface** *interface-id*	Enter interface configuration mode, and specify the interface to attach to the policy map. Valid interfaces include physical interfaces.
Step 10	**service-policy input** *policy-map-name*	Apply specified policy map to the input of a particular interface. Only one policy map per interface per direction is supported.
Step 11	**end**	Return to privileged EXEC mode.
Step 12	**show policy-map** [*policy-map-name* **class** *class-name*]	Verify your entries.
Step 13	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To delete an existing policy map, use the **no policy-map** *policy-map-name* global configuration command. To delete an existing class map, use the **no class** *class-map-name* policy-map configuration command. To remove an assigned DSCP value, use the **no set ip dscp** *new-dscp* policy-map configuration command. To remove an existing policer, use the **no police** *rate-bps burst-byte* [**exceed-action** {**drop** | **dscp** *dscp-value*}] policy-map configuration command. To remove the policy map and interface association, use the **no service-policy input** *policy-map-name* interface configuration command.

For details about configuring policy maps and security ACLs on the same interface, see Table 30-5 on page 30-18.

This example shows how to create a policy map and attach it to an ingress interface. In the configuration, the IP standard ACL permits traffic from network 10.1.0.0. For traffic matching this classification, the DSCP value in the incoming packet is trusted. If the matched traffic exceeds an average traffic rate of 5000000 bps and a normal burst size of 8192 bytes, its DSCP is marked down to a value of 10 and sent.

```
Switch(config)# access-list 1 permit 10.1.0.0 0.0.255.255
Switch(config)# class-map ipclass1
Switch(config-cmap)# match access-group 1
Switch(config-cmap)# exit
Switch(config)# policy-map flow1t
Switch(config-pmap)# class ipclass1
Switch(config-pmap-c)# police 5000000 8192 exceed-action dscp 10
Switch(config-pmap-c)# exit
Switch(config-pmap)# exit
Switch(config)# interface gigabitethernet0/1
Switch(config-if)# switchport mode access
Switch(config-if)# service-policy input flow1t
```

This example shows how to create a Layer 2 MAC ACL with two permit statements and attach it to an ingress interface. The first permit statement allows traffic from the host with MAC address 0001.0000.0001 destined for the host with MAC address 0002.0000.0001.

```
Switch(config)# mac access-list extended maclist1
Switch(config-ext-mac)# permit host 0001.0000.0001 host 0002.0000.0001
Switch(config-ext-mac)# exit
Switch(config)# mac access-list extended maclist2
Switch(config-ext-mac)# permit host 0001.0000.0003 host 0002.0000.0003
Switch(config-ext-mac)# exit
Switch(config)# class-map macclass1
Switch(config-cmap)# match access-group name maclist1
Switch(config-cmap)# exit
Switch(config)# policy-map macpolicy1
Switch(config-pmap)# class macclass1
Switch(config-pmap-c)# set ip dscp 56
Switch(config-pmap-c)# exit
Switch(config-pmap)# class macclass2 maclist2
Switch(config-pmap-c)# set ip dscp 48
Switch(config-pmap-c)# exit
Switch(config-pmap)# exit
Switch(config)# interface gigabitethernet0/1
Switch(config-if)# switchport mode trunk
Switch(config-if)# mls qos trust cos
Switch(config-if)# service-policy input macpolicy1
```

Configuring CoS Maps

Note This feature is available only if your switch is running the EI.

This section describes how to configure the CoS maps:

- Configuring the CoS-to-DSCP Map, page 30-33
- Configuring the DSCP-to-CoS Map, page 30-34

All the maps are globally defined.

Configuring the CoS-to-DSCP Map

You use the CoS-to-DSCP map to map CoS values in incoming packets to a DSCP value that QoS uses internally to represent the priority of the traffic.

Table 30-7 shows the default CoS-to-DSCP map.

Table 30-7 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

If these values are not appropriate for your network, you need to modify them.

Beginning in privileged EXEC mode, follow these steps to modify the CoS-to-DSCP map:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**mls qos map cos-dscp** *dscp1...dscp8*	Modify the CoS-to-DSCP map. For *dscp1...dscp8*, enter 8 DSCP values that correspond to CoS values 0 to 7. Separate each DSCP value with a space. The supported DSCP values are 0, 8, 10, 16, 18, 24, 26, 32, 34, 40, 46, 48, and 56.
Step 3	**end**	Return to privileged EXEC mode.
Step 4	**show mls qos maps cos-dscp**	Verify your entries.
Step 5	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To return to the default map, use the **no mls qos map cos-dscp** global configuration command.

This example shows how to modify and display the CoS-to-DSCP map:

```
Switch# configure terminal
Switch(config)# mls qos map cos-dscp 8 8 8 8 24 32 56 56
Switch(config)# end
Switch# show mls qos maps cos-dscp

Cos-dscp map:
        cos:  0  1  2  3  4  5  6  7
     --------------------------------
       dscp:  8  8  8  8 24 32 56 56
```

Configuring the DSCP-to-CoS Map

You use the DSCP-to-CoS map to map DSCP values in incoming packets to a CoS value, which is used to select one of the four egress queues.

The switch supports these DSCP values: 0, 8, 10, 16, 18, 24, 26, 32, 34, 40, 46, 48, and 56.

Table 30-8 shows the default DSCP-to-CoS map.

Table 30-8 Default DSCP-to-CoS Map

DSCP values	0	8, 10	16, 18	24, 26	32, 34	40, 46	48	56
CoS values	0	1	2	3	4	5	6	7

If these values are not appropriate for your network, you need to modify them.

Beginning in privileged EXEC mode, follow these steps to modify the DSCP-to-CoS map:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**mls qos map dscp-cos** *dscp-list* **to** *cos*	Modify the DSCP-to-CoS map. For *dscp-list*, enter up to 13 DSCP values separated by spaces. Then enter the **to** keyword. For *cos*, enter the CoS value to which the DSCP values correspond. The supported DSCP values are 0, 8, 10, 16, 18, 24, 26, 32, 34, 40, 46, 48, and 56. The CoS range is 0 to 7.
Step 3	**end**	Return to privileged EXEC mode.
Step 4	**show mls qos maps dscp-cos**	Verify your entries.
Step 5	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To return to the default map, use the **no mls qos map dscp-cos** global configuration command.

This example shows how the DSCP values 26 and 48 are mapped to CoS value 7. For the remaining DSCP values, the DSCP-to-CoS mapping is the default.

```
Switch(config)# mls qos map dscp-cos 26 48 to 7
Switch(config)# exit

Switch# show mls qos maps dscp-cos

Dscp-cos map:
       dscp:  0  8 10 16 18 24 26 32 34 40 46 48 56
     -----------------------------------------------
        cos:  0  1  1  2  2  3  7  4  4  5  5  7  7
```

Configuring the Egress Queues

Note This feature is supported by both the SI and EI.

This section describes how to configure the egress queues:

- Configuring CoS Priority Queues, page 30-35
- Configuring WRR Priority, page 30-36
- Enabling the Expedite Queue and Configuring WRR Priority, page 30-36

For more information about the egress queues, see the "Egress CoS Queues" section on page 30-9.

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** *qid cos1..cosn*	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: 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** *weight1...weight4*	Assign WRR weights to the four CoS queues. These are the ranges for the WRR values: • For *weight1*, *weight2*, and *weight3*, the range is 1 to 255. • For *weight4*, the range is 0 to 255. When weight4 is set to 0, queue 4 is configured as the expedite queue. **Note** In software releases earlier than Cisco IOS Release 12.1(12c)EA1, the ranges for all queues 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 "Enabling the Expedite Queue and Configuring WRR Priority" section on page 30-36.

Enabling the Expedite Queue and Configuring WRR Priority

In Cisco IOS Release 12.1(12c)EA1 or later, beginning in privileged EXEC mode, follow these steps to enable the expedite queue (queue 4) and assign WRR priority to the remaining queues:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**wrr-queue bandwidth** *weight1 weight2 weight3* **0**	Configure queue 4 as the expedite queue and assign WRR weights to the remaining egress queues. The range of WRR weights for *weight1*, *weight2*, and *weight3* 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.

Displaying Standard QoS Information

To display standard QoS information, use one or more of the privileged EXEC commands in Table 30-9:

Table 30-9 Commands for Displaying QoS Information

Command	Purpose
show class-map [*class-map-name*][1]	Display QoS class maps, which define the match criteria to classify traffic.
show policy-map [*policy-map-name* [**class** *class-name*]][1]	Display QoS policy maps, which define classification criteria for incoming traffic.
show mls qos maps [**cos-dscp** | **dscp-cos**][1]	Display QoS mapping information. Maps are used to generate an internal DSCP value, which represents the priority of the traffic.
show mls qos interface [*interface-id*] [**policers**][1]	Display QoS information at the interface level, including the configuration of the egress queues and the CoS-to-egress-queue map, which interfaces have configured policers, and ingress statistics.
show mls masks [**qos** | **security**][1]	Display details regarding the masks[2] used for QoS and security ACLs.
show wrr-queue cos-map	Display the mapping of the CoS priority queues.
show wrr-queue bandwidth	Display the WRR bandwidth allocation for the CoS priority queues.

[1] Available only on a switch running the EI.

[2] Access control parameters are called masks in the switch CLI commands and output.

Standard QoS Configuration Examples

Note These examples are applicable only if your switch is running the EI.

This section shows a QoS migration path to help you quickly implement QoS features based on your existing network and planned changes to your network, as shown in Figure 30-5. It contains this information:

- QoS Configuration for the Existing Wiring Closet, page 30-38
- QoS Configuration for the Intelligent Wiring Closet, page 30-39

Figure 30-5 QoS Configuration Example Network

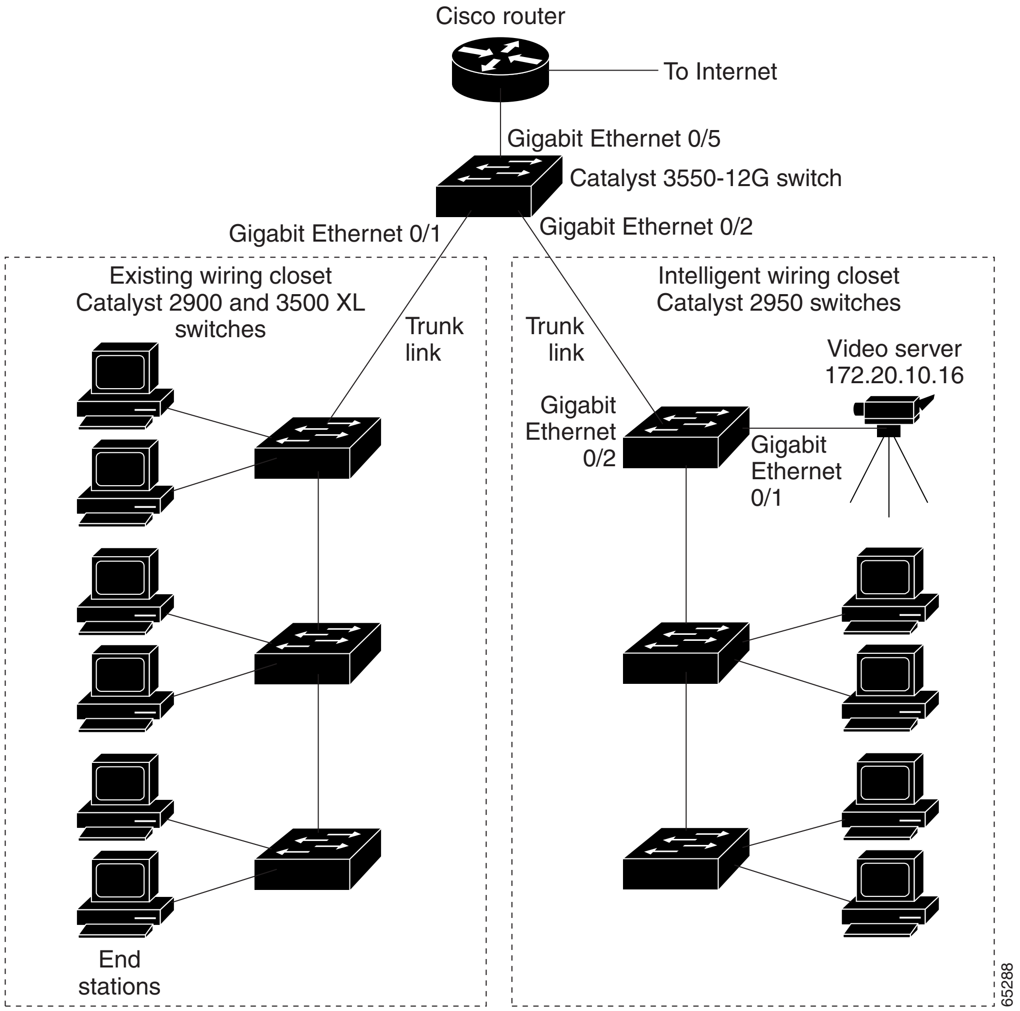

QoS Configuration for the Existing Wiring Closet

The existing wiring closet in Figure 30-5 consists of existing Catalyst 2900 XL and 3500 XL switches. These switches are running Cisco IOS Release 12.0(5)XP or later, which supports the QoS-based IEEE 802.1P CoS values. QoS classifies frames by assigning priority-indexed CoS values to them and gives preference to higher-priority traffic.

Recall that on the Catalyst 2900 and 3500 XL switches, you can classify untagged (native) Ethernet frames at the ingress ports by setting a default CoS priority (**switchport priority default** *default-priority-id* interface configuration command) for each port. For IEEE 802.1Q frames with tag information, the priority value from the header frame is used. On the Catalyst 3524-PWR XL and 3548 XL switches, you can override this priority with the default value by using the **switchport priority default override** interface configuration command. For Catalyst 2950 and Catalyst 2900 XL switches and other 3500 XL models that do not have the override feature, the Catalyst 3550-12T switch at the distribution layer can override the 802.1P CoS value by using the **mls qos cos override** interface configuration command.

For the Catalyst 2900 and 3500 XL switches, CoS configures each transmit port (the egress port) with a normal-priority transmit queue and a high-priority transmit queue, depending on the frame tag or the port information. Frames in the normal-priority queue are forwarded only after frames in the high-priority queue are forwarded. Frames that have 802.1P CoS values of 0 to 3 are placed in the normal-priority transmit queue while frames with CoS values of 4 to 7 are placed in the expedite (high-priority) queue.

QoS Configuration for the Intelligent Wiring Closet

The intelligent wiring closet in Figure 30-5 is composed of Catalyst 2950 switches. One of the switches is connected to a video server, which has an IP address of 172.20.10.16.

The object of this example is to prioritize the video traffic over all other traffic. To do so, a DSCP of 46 is assigned to the video traffic. This traffic is stored in queue 4, which is serviced more frequently than the other queues.

Beginning in privileged EXEC mode, follow these steps to configure the switch to prioritize video packets over all other traffic:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**access-list 1 permit 172.20.10.16**	Define an IP standard ACL, and permit traffic from the video server at 172.20.10.16.
Step 3	**class-map videoclass**	Create a class map called *videoclass*, and enter class-map configuration mode.
Step 4	**match access-group 1**	Define the match criterion by matching the traffic specified by ACL 1.
Step 5	**exit**	Return to global configuration mode.
Step 6	**policy-map videopolicy**	Create a policy map called *videopolicy*, and enter policy-map configuration mode.
Step 7	**class videoclass**	Specify the class on which to act, and enter policy-map class configuration mode.
Step 8	**set ip dscp 46**	For traffic matching ACL 1, set the DSCP of incoming packets to 46.
Step 9	**police 5000000 8192 exceed-action drop**	Define a policer for the classified video traffic to drop traffic that exceeds 5-Mbps average traffic rate with an 8192-byte burst size.
Step 10	**exit**	Return to policy-map configuration mode.
Step 11	**exit**	Return to global configuration mode.
Step 12	**interface gigabitethernet0/1**	Enter interface configuration mode, and specify the ingress interface.
Step 13	**service-policy input videopolicy**	Apply the policy to the ingress interface.
Step 14	**exit**	Return to global configuration mode.
Step 15	**wrr-queue bandwidth 1 2 3 4**	Assign a higher WRR weight to queue 4.
Step 16	**wrr-queue cos-map 4 6 7**	Configure the CoS-to-egress-queue map so that CoS values 6 and 7 select queue 4.
Step 17	**end**	Return to privileged EXEC mode.

	Command	Purpose
Step 18	**show class-map videoclass** **show policy-map videopolicy** **show mls qos maps** [**cos-dscp** | **dscp-cos**]	Verify your entries.
Step 19	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.