



## Configuring Quality of Service

---

This chapter describes the quality of service (QoS) features built into your ML-Series card and how to map QoS scheduling at both the system and interface levels.

This chapter contains the following major sections:

- [Understanding QoS, page 14-1](#)
- [ML-Series QoS, page 14-4](#)
- [QoS on Cisco Proprietary RPR, page 14-10](#)
- [Configuring QoS, page 14-12](#)
- [Monitoring and Verifying QoS Configuration, page 14-17](#)
- [QoS Configuration Examples, page 14-18](#)
- [Understanding Multicast QoS and Priority Multicast Queuing, page 14-24](#)
- [Configuring Multicast Priority Queuing QoS, page 14-25](#)
- [QoS not Configured on Egress, page 14-27](#)
- [ML-Series Egress Bandwidth Example, page 14-27](#)
- [Understanding CoS-Based Packet Statistics, page 14-29](#)
- [Configuring CoS-Based Packet Statistics, page 14-29](#)
- [Understanding IP SLA, page 14-31](#)

The ML-Series card employs the Cisco IOS Modular QoS command-line interface (CLI), known as the MQC. For more information on general MQC configuration, refer to the following Cisco IOS documents:

- *Cisco IOS Quality of Service Solutions Configuration Guide, Release 12.2*
- *Cisco IOS Quality of Service Solutions Command Reference, Release 12.2*

## Understanding QoS

QoS is the ability of the network to provide better or special treatment to a set of services to the detriment of less critical services. The ML-Series card uses QoS to dynamically allocate transmission bandwidth for the different services it multiplexes onto the SONET/SDH circuit. Through QoS, you can configure the ML-Series card to provide different levels of treatment to the different services. The different levels are defined through the service elements of bandwidth, including loss and delay. A service-level agreement (SLA) is a guaranteed level of these service elements.

The QoS mechanism has three basic steps. It classifies types of traffic, specifies what action to take against a type of traffic, and specifies where the action should take place. The following sections explain how the ML-Series card accomplishes these steps for unicast traffic. QoS for priority-multicast traffic and traffic with unknown destination addresses is handled with a different mechanism, detailed in the [“Understanding Multicast QoS and Priority Multicast Queuing”](#) section on page 14-24.

## Priority Mechanism in IP and Ethernet


For any QoS service to be applied to data, there must be a way to mark or identify an IP packet or an Ethernet frame. When identified, a specific priority can be assigned to each individual IP packet or Ethernet frame. The IP Precedence field or the IP Differentiated Services Code Point (DSCP) field prioritizes the IP packets, and the Ethernet class of service (IEEE 802.1p defined class of service [CoS]) is used for the Ethernet frames. IP precedence and Ethernet CoS are further described in the following sections.

## IP Precedence and Differentiated Services Code Point

IP precedence uses the three precedence bits in the IPv4 header's ToS (type of service) field to specify class of service for each IP packet (IETF RFC 1122). The most significant three bits on the IPv4 ToS field provides up to eight distinct classes, of which six are used for classifying services and the remaining two are reserved. On the edge of the network, the IP precedence is assigned by the client device or the router, so that each subsequent network element can provide services based on the determined policy or the SLA.

IP DSCP uses the six bits in the IPv4 header to specify class of service for each IP packet (IETF RFC 2474). [Figure 14-1](#) illustrates IP precedence and DSCP. The DSCP field classifies packets into any of the 64 possible classes. On the network edge, the IP DSCP is assigned by the client device or the router, so that each subsequent network element can provide services based on the determined policy or the SLA.

Figure 14-1 IP Precedence and DSCP

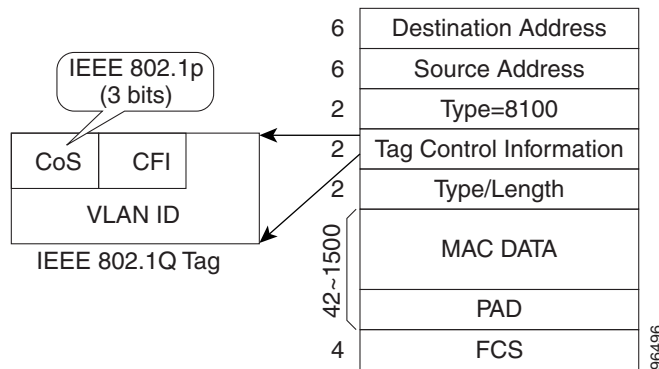


96496

## Ethernet CoS

Ethernet CoS refers to three bits within a four byte IEEE 802.1Q (VLAN) header used to indicate the priority of the Ethernet frame as it passes through a switched network. The CoS bits in the IEEE 802.1Q header are commonly referred to as the IEEE 802.1p bits. There are three CoS bits that provide eight classes, matching the number delivered by IP precedence. In many real-world networks, a packet might traverse both Layer 2 and Layer 3 domains. To maintain QoS across the network, the IP ToS can be mapped to the Ethernet CoS and vice versa, for example, in linear or one-to-one mapping, because each mechanism supports eight classes. Similarly, a set of DSCP values (64 classes) can be mapped into each of the eight individual Ethernet CoS values. Figure 14-2 shows an IEEE 802.1Q Ethernet frame, which consists of a 2-byte Ethertype and a 2-byte tag (IEEE 802.1Q tag) on the Ethernet protocol header.

Figure 14-2 Ethernet Frame and the CoS Bit (IEEE 802.1p)




96496

# ML-Series QoS

The ML-Series QoS classifies each packet in the network based on its input interface, bridge group (VLAN), Ethernet CoS, IP precedence, IP DSCP, or Cisco proprietary resilient packet ring RPR-CoS. After they are classified into class flows, further QoS functions can be applied to each packet as it traverses the card. Figure 14-3 illustrates the ML-Series QoS flow.

**Figure 14-3 ML-Series QoS Flow**



Policing provided by the ML-Series card ensures that attached equipment does not submit more than a predefined amount of bandwidth (Rate Limiting) into the network. The policing feature can be used to enforce the committed information rate (CIR) and the peak information rate (PIR) available to a customer at an interface. Policing also helps characterize the statistical nature of the information allowed into the network so that traffic engineering can more effectively ensure that the amount of committed bandwidth is available on the network, and that the peak bandwidth is over-subscribed with an appropriate ratio. The policing action is applied per classification.

Priority marking can set the Ethernet IEEE 802.1p CoS bits or RPR-CoS bits as they exit the ML-Series card. The marking feature operates on the outer IEEE 802.1p tag, and provides a mechanism for tagging packets at the ingress of a QinQ packet. The subsequent network elements can provide QoS based only on this service-provider-created QoS indicator.

Per-class flow queuing enables fair access to excess network bandwidth, allows allocation of bandwidth to support SLAs, and ensures that applications with high network resource requirements are adequately served. Buffers are allocated to queues dynamically from a shared resource pool. The allocation process incorporates the instantaneous system load as well as the allocated bandwidth to each queue to optimize buffer allocation. Congestion management on the ML-Series is performed through a tail drop mechanism along with discard eligibility on the egress scheduler.

The ML-Series uses a Weighted Deficit Round Robin (WDRR) scheduling process to provide fair access to excess bandwidth as well as guaranteed throughput to each class flow.

Admission control is a process that is invoked each time that service is configured on the ML-Series card to ensure that QoS resources are not overcommitted. In particular, admission control ensures that no configurations are accepted, where a sum of the committed bandwidths on an interface exceeds total bandwidth on the interface.

## Classification


Classification can be based on any single packet classification criteria or a combination (logical AND and OR). A total of 254 classes, not including the class default, can be defined on the card. Classification of packets is configured using the Modular CLI **class-map** command. For traffic transiting the Cisco Proprietary RPR, only the input interface and/or the RPR-CoS can be used as classification criteria.

## Policing

Dual leaky bucket policer is a process where the first bucket (CIR bucket) is filled with tokens at a known rate (CIR), which is a parameter that can be configured by the operator. Figure 14-4 illustrates the dual leaky bucket policer model. The tokens fill the bucket up to a maximum level, which is the amount of burstable committed (BC) traffic on the policer. The nonconforming packets of the first bucket are the overflow packets, which are passed to the second leaky bucket (the PIR bucket). The second leaky bucket is filled with these tokens at a known rate (PIR), which is a parameter that can be configured by the operator. The tokens fill the PIR bucket up to a maximum level (BP), which is the amount of peak burstable traffic on the policer. The nonconform packets of the second bucket are the overflow packets, which can be dropped or marked according to the policer definition.

On the dual leaky bucket policer, the packets conforming to the CIR are conform packets, the packets not conforming to CIR but conforming to PIR are exceed packets, and the packets not conforming to either the PIR or CIR are violate packets.

Figure 14-4 Dual Leaky Bucket Policer Model



## Marking and Discarding with a Policier

On the ML-Series card's policier, the conform packets can be transmitted or marked and transmitted. The exceed packets can be transmitted, marked and transmitted, or dropped. The violating packets can be transmitted, marked and transmitted, or dropped. The primary application of the dual-rate or three-color policier is to mark the conform packets with CoS bit 21, mark the exceed packet with CoS bit 1, and discard the violated packets so all the subsequent network devices can implement the proper QoS treatment per frame/packet basis based on these priority marking without knowledge of each SLA.

In some cases, it might be desirable to discard all traffic of a specific ingress class. This can be accomplished by using a police command of the following form with the class: **police 96000 conform-action drop exceed-action drop**.

If a marked packet has a provider-supplied Q-tag inserted before transmission, the marking only affects the provider Q-tag. If a Q-tag is received, it is re-marked. If a marked packet is transported over the Cisco proprietary RPR ring, the marking also affects the RPR-CoS bit.

If a Q-tag is inserted (QinQ), the marking affects the added Q-tag. If the ingress packet contains a Q-tag and is transparently switched, the existing Q-tag is marked. In the case of a packet without any Q-tag, the marking does not have any significance.

The local scheduler treats all nonconforming packets as discard eligible regardless of their CoS setting or the global CoS commit definition. For Cisco proprietary RPR implementation, the discard eligible (DE) packets are marked using the DE bit on the Cisco proprietary RPR header. The discard eligibility based on the CoS commit or the policing action is local to the ML-Series card scheduler, but it is global for the Cisco proprietary RPR ring.

## Queuing

ML-Series card queuing uses a shared buffer pool to allocate memory dynamically to different traffic queues. The ML-Series card uses a total of 12 MB of memory for the buffer pool. Ethernet ports share 6 MB of the memory, and packet-over-SONET/SDH (POS) ports share the remaining 6 MBs of memory. Memory space is allocated in 1500-byte increments.

Each queue has an upper limit on the allocated number of buffers based on the class bandwidth assignment of the queue and the number of queues configured. This upper limit is typically 30 percent to 50 percent of the shared buffer capacity. Dynamic buffer allocation to each queue can be reduced based on the number of queues that need extra buffering. The dynamic allocation mechanism provides fairness in proportion to service commitments as well as optimization of system throughput over a range of system traffic loads.

The Low Latency Queue (LLQ) is defined by setting the weight to infinity or by committing 100 percent of the bandwidth. When a LLQ is defined, a policer should also be defined on the ingress for that specific class to limit the maximum bandwidth consumed by the LLQ; otherwise there is a potential risk of LLQ occupying the whole bandwidth and starving the other unicast queues.

The ML-Series includes support for 400 user-definable queues, which are assigned according to the classification and bandwidth allocation definition. The classification used for scheduling classifies the frames/packet after the policing action, so if the policer is used to mark or change the CoS bits of the ingress frames/packet, the new values are applicable for the classification of traffic for queuing and scheduling. The ML-Series provides buffering for 4000 packets.

## Scheduling

Scheduling is provided by a series of schedulers that perform a WDRR as well as by priority scheduling mechanisms from the queued traffic associated with each egress port.


Though ordinary round robin servicing of queues can be done in constant time, unfairness occurs when different queues use different packet sizes. Deficit Round Robin (DRR) scheduling solves this problem. If a queue was not able to send a packet in its previous round because its packet size was too large, the remainder from the previous amount of credits a queue gets in each round (quantum) is added to the quantum for the next round.

WDRR extends the quantum idea from the DRR to provide weighted throughput for each queue. Different queues have different weights, and the quantum assigned to each queue in its round is proportional to the relative weight of the queue among all the queues serviced by that scheduler.

Weights are assigned to each queue as a result of the service provisioning process. When coupled with policing and policy mapping provisioning, these weights and the WDRR scheduling process ensure that QoS commitments are provided to each service flow.

Figure 14-5 illustrates the ML-Series card's queuing and scheduling.

**Figure 14-5** Queuing and Scheduling Model



The weighting structure allows traffic to be scheduled at 1/2048 of the port rate. This equates to approximately 488 kbps for traffic exiting a Gigabit Ethernet port, approximately 293 kbps for traffic exiting an OC-12c port, and approximately 49 kbps for traffic exiting a FastEthernet port.

The unicast queues are created as the output service policy implementation on the egress ports. Each unicast queue is assigned with a committed bandwidth and the weight of the queue is determined by the normalization of committed bandwidth of all defined unicast queues for that port. The traffic beyond the committed bandwidth on any queue is treated by the scheduler according to the relative weight of the queue.

The LLQ is created as the output service policy implementation on the egress ports. Each LLQ queue is assigned with a committed bandwidth of 100 percent and is served with lower latency. To limit the bandwidth usage by the LLQ, a strict policer needs to be implemented on the ingress for the LLQ traffic classes.

The DE allows some packets to be treated as committed and some as discard-eligible on the scheduler. For Ethernet frames, the CoS (IEEE 802.1p) bits are used to identify committed and discard eligible packets, where the RPR-CoS and the DE bits are used for Cisco proprietary RPR traffic. When congestion occurs and a queue begins to fill, the DE packets hit a lower tail-drop threshold than the committed packets. Committed packets are not dropped until the total committed load exceeds the interface output. The tail-drop thresholds adjust dynamically in the card to maximize use of the shared buffer pool while guaranteeing fairness under all conditions.

## Control Packets and L2 Tunneled Protocols

The control packets originated by the ML-Series card have a higher priority than data packets. The external Layer 2 and Layer 3 control packets are handled as data packets and assigned to broadcast queues. Bridge protocol data unit (BPDU) prioritization in the ML-Series card gives Layer 2-tunneled BPDU sent out the multicast/broadcast queue a higher discard value and therefore a higher priority than other packets in the multicast/broadcast queue. The Ethernet CoS (IEEE 802.1p) for Layer 2-tunneled protocols can be assigned by the ML-Series card.

## Egress Priority Marking

Egress priority marking allows the operator to assign the IEEE 802.1p CoS bits of packets that exit the card. This marking allows the operator to use the CoS bits as a mechanism for signaling to downstream nodes the QoS treatment the packet should be given. This feature operates on the outer-most IEEE 802.1p CoS field. When used with the QinQ feature, priority marking allows the user traffic (inner Q-tag) to traverse the network transparently, while providing a means for the network to internally signal QoS treatment at Layer 2.

Priority marking follows the classification process, and therefore any of the classification criteria identified earlier can be used as the basis to set the outgoing IEEE 802.1p CoS field. For example, a specific CoS value can be mapped to a specific bridge group.

Priority marking is configured using the MQC **set-cos** command. If packets would otherwise leave the card without an IEEE 802.1Q tag, then the **set-cos** command has no effect on that packet. If an IEEE 802.1Q tag is inserted in the packet (either a normal tag or a QinQ tag), the inserted tag has the set-cos priority. If an IEEE 802.1Q tag is present on packet ingress and retained on packet egress, the priority of that tag is modified. If the ingress interface is a QinQ access port and the **set-cos** policy-map classifies based on ingress tag priority, this classifies based on the user priority. This is a way to allow the user-tag priority to determine the SP tag priority. When a packet does not match any **set-cos** policy-map, the priority of any preserved tag is unchanged and the priority of any inserted IEEE 802.1Q tag is set to 0.

The **set-cos** command on the output service policy is only applied to unicast traffic. Priority marking for multicast/broadcast traffic can only be achieved by the **set-cos** action of the policing process on the input service policy.

## Ingress Priority Marking

Ingress priority marking can be done for all input packets of a port, for all input packets matching a classification, or based on a measured rate. Marking of all packets of an input class can also be done with a policing command of the form **police 96000 conform-action set-cos-transmit exceed-action set-cos-transmit**. Using this command with a policy map that contains only the "class-default" will mark all ingress packets to the value. Rate based priority marking is discussed in the ["Marking and Discarding with a Policer"](#) section on page 14-5.

## QinQ Implementation

The hierarchical VLAN or IEEE 802.1Q tunneling feature enables the service provider to transparently carry the customer VLANs coming from any specific port (UNI) and transport them over the service provider network. This feature is also known as QinQ, which is performed by adding an additional IEEE 802.1Q tag on every customer frame.




Using the QinQ feature, service providers can use a single VLAN to support customers with multiple VLANs. QinQ preserves customer VLAN IDs and segregates traffic from different customers within the service-provider infrastructure, even when traffic from different customers originally shared the same VLAN ID. The QinQ also expands VLAN space by using a VLAN-in-VLAN hierarchy and tagging the tagged packets. When the service provider (SP) tag is added, the QinQ network typically loses any visibility to the IP header or the customer Ethernet IEEE 802.1Q tag on the QinQ encapsulated frames.

On the ML-Series cards, the QinQ access ports (IEEE 802.1Q tunnel ports or QinQ UNI ports) have visibility to the customer CoS and the IP precedence or IP DSCP values; therefore, the SP tag can be assigned with the proper CoS bit, which would reflect the customer IP precedence, IP DSCP, or CoS bits. In the QinQ network, the QoS is then implemented based on the IEEE 802.1p bit of the SP tag. The ML-Series cards do not have visibility into the customer CoS, IP precedence, or DSCP values after the packet is double-tagged (because it is beyond the entry point of the QinQ service).

Figure 14-6 illustrates the QinQ implementation on the ML-Series card.

**Figure 14-6** QinQ



The ML-Series cards can be used as the IEEE 802.1Q tunneling device for the QinQ network and also provide the option to copy the customer frame's CoS bit into the CoS bit of the added QinQ tag. This way, the service provider QinQ network can be fully aware of the necessary QoS treatment for each individual customer frame.

## Flow Control Pause and QoS

When flow control and policy-map are enabled for an interface and the policy-map is configured only with 'class-default' having policer action, flow control handles the bandwidth. For all the packets, which match policers drop criteria, PAUSE frames are sent upstream so that far end device can reduce its transmit rate accordingly. If the far end device honours the received PAUSE frames, there will not be any drops on ML card due to policer configuration. However, if the policer gets noncompliant flow, the packets are dropped or demarked using the policer definition of the interface.



### Note

The above statement is valid for an interface, which has a policer with not only a class-default (i.e. with non-default class) configured. When the policy-map is configured only with class-default, the policer acts instead of allowing the flow control to drop or demark the frames.

**Note**


---

QoS and policing are not supported on the ML-Series card interface when link aggregation is used.

---

**Note**


---

Egress shaping is not supported on the ML-Series cards.

---

## QoS on Cisco Proprietary RPR

For VLAN bridging over Cisco proprietary RPR, all ML-Series cards on the ring must be configured with the base Cisco proprietary RPR and Cisco proprietary RPR QoS configuration. SLA and bridging configurations are only needed at customer Cisco proprietary RPR access points, where IEEE 802.1Q VLAN CoS is copied to the Cisco proprietary RPR CoS. This IEEE 802.1Q VLAN CoS copying can be overwritten with a **set-cos action** command. The CoS commit rule applies at Cisco proprietary RPR ring ingress.

- If the packet does not have a VLAN header, the Cisco proprietary RPR CoS for non-VLAN traffic is set using the following rules:
- The default CoS is 0.
- If the packet comes in with an assigned CoS, the assigned CoS replaces the default. If an IP packet originates locally, the IP precedence setting replaces the CoS setting.
- The input policy map has a **set-cos** action.
- The output policy map has a **set-cos** action (except for broadcast or multicast packets).

The Cisco proprietary RPR header contains a CoS value and DE indicator. The Cisco proprietary RPR DE is set for noncommitted traffic.

The ML-Series card Cisco proprietary RPR transit traffic, which is defined as traffic going from POS port to POS port around the Cisco proprietary RPR, can only be classified by Layer 2 CoS. Other match rules are ignored. This is a ML-Series card specific implementation of QoS on Cisco proprietary RPR designed for the CoS based QoS model of the Cisco Metro Ethernet Solution.

This Layer 2 CoS dependence prevents DSCP-based output policy maps from working properly with Cisco proprietary RPR on the ML-Series card. Using a DSCP based policy-map causes all transit traffic to be incorrectly treated as class-default. This results in a discard of the transit traffic without any regard for the DSCP priority when transit station congestion occurs.

The DSCP based output policy map limitation has a work around. Each Cisco proprietary RPR frame has its own three bit CoS marking, which is normally copied from the VLAN CoS. This is the field on which "match cos" classification is done for transit Cisco proprietary RPR traffic. The Cisco proprietary RPR CoS can be marked based on the DSCP match at the input station, and then classified based on the Cisco proprietary RPR CoS at transit stations. This method can support a maximum of eight classes. If you are using nine classes (including class-default), two of them would need to be combined to use this work-around.

[Example 14-1](#) shows a class and policy-map definition configuration that would overcome the DSCP limitation. The example also changes nine classes into eight by combining the Voice and Call-Sig classes.

**Caution**

"Match cos 0" should not be included in the definition of any class-map, because non-VLAN-tagged Ethernet packets are always treated as CoS 0 on input from Ethernet. Using "match cos 0" might incorrectly match all traffic coming from Ethernet.

**Example 14-1 Class and Policy-map Definition Configuration Overcoming the DSCP Limitation**

```
class-map match-any Bulk-Data
  match ip dscp af11
  match cos 3
class-map match-any Crit-Data
  match ip dscp af21 af31
  match cos 7
class-map match-any Net-Management
  match ip dscp cs2
  match cos 2
class-map match-any Video
  match ip dscp cs4 af41
  match cos 4
class-map match-any Voice
  description Includes Voice and Call Signalling
  match ip dscp ef
  match ip dscp cs3
  match cos 5
class-map match-any Routing
  match ip dscp cs6
  match cos 6
class-map match-any Scavenger
  match ip dscp cs1
  match cos 1
policy-map MAN-QoS-DSCP
  class Voice
    priority percent 4
    set cos 5
  class Bulk-Data
    bandwidth percent 20
    set cos 3
  class Crit-Data
    bandwidth percent 20
    set cos 7
  class Net-Management
    bandwidth percent 2
    set cos 2
  class Video
    bandwidth percent 5
    set cos 4
  class Routing
    bandwidth percent 2
    set cos 6
  class Scavenger
    bandwidth percent 1
    set cos 1
  class class-default
    bandwidth percent 45
    set cos 0
```

## Configuring QoS

This section describes the tasks for configuring the ML-Series card QoS functions using the MQC. The ML-Series card does not support the full set of MQC functionality.

To configure and enable class-based QoS features, perform the procedures described in the following sections:

- [Creating a Traffic Class, page 14-12](#)
- [Creating a Traffic Policy, page 14-13](#)
- [Attaching a Traffic Policy to an Interface, page 14-16](#)
- [Configuring CoS-Based QoS, page 14-17](#)

For QoS configuration examples, see the “QoS Configuration Examples” section on page 14-18.

### Creating a Traffic Class

The **class-map** global configuration command is used to create a traffic class. The syntax of the **class-map** command is as follows:

```
class-map [match-any | match-all] class-map-name
no class-map [match-any | match-all] class-map-name
```

The **match-all** and **match-any** options need to be specified only if more than one match criterion is configured in the traffic class. The **class-map match-all** command is used when all of the match criteria in the traffic class must be met for a packet to match the specified traffic class. The **class-map match-any** command is used when only one of the match criterion in the traffic class must be met for a packet to match the specified traffic class. If neither the **match-all** nor the **match-any** keyword is specified, the traffic class behaves in a manner consistent with the **class-map match-all** command.

To create a traffic class containing match criteria, use the **class-map** global configuration command to specify the traffic class name, and then use the **match** commands in [Table 14-1](#), as needed.

**Table 14-1** Traffic Class Commands

Command	Purpose
Router(config)# <b>class-map</b> <i>class-map-name</i>	Specifies the user-defined name of the traffic class. Names can be a maximum of 40 alphanumeric characters. If neither <b>match-all</b> nor <b>match-any</b> is specified, traffic must match all the match criteria to be classified as part of the traffic class.  There is no default-match criteria.  Multiple match criteria are supported. The command matches either all or any of the criteria, as controlled by the <b>match-all</b> and <b>match-any</b> subcommands of the <b>class-map</b> command.
Router(config)# <b>class-map match-all</b> <i>class-map-name</i>	Specifies that all match criteria must be met for traffic entering the traffic class to be classified as part of the traffic class.
Router(config)# <b>class-map match-any</b> <i>class-map-name</i>	Specifies that one of the match criteria must be met for traffic entering the traffic class to be classified as part of the traffic class.
Router(config-cmap)# <b>match any</b>	Specifies that all packets will be matched.
Router(config-cmap)# <b>match bridge-group</b> <i>bridge-group-number</i>	Specifies the bridge-group-number against whose contents packets are checked to determine if they belong to the class.

**Table 14-1** Traffic Class Commands (continued)

Command	Purpose
Router(config-cmap)# <b>match cos</b> <i>cos-number</i>	Specifies the CoS value against whose contents packets are checked to determine if they belong to the class.
Router(config-cmap)# <b>match input-interface</b> <i>interface-name</i>	Specifies the name of the input interface used as a match criterion against which packets are checked to determine if they belong to the class.  The shared packet ring (SPR) interface used in Cisco proprietary RPR (SPR1) is a valid interface-name for the ML-Series card. For more information on the SPR interface, see <a href="#">Chapter 17, “Configuring Cisco Proprietary Resilient Packet Ring.”</a>  The <b>input-interface</b> choice is not valid when applied to the INPUT of an interface (redundant).
Router(config-cmap)# <b>match ip dscp</b> <i>ip-dscp-value</i>	Specifies up to eight DSCP values used as match criteria. The value of each service code point is from 0 to 63.
Router (config-cmap)# <b>match ip precedence</b> <i>ip-precedence-value</i>	Specifies up to eight IP precedence values used as match criteria.

## Creating a Traffic Policy

To configure a traffic policy, use the **policy-map** global configuration command to specify the traffic policy name, and use the following configuration commands to associate a traffic class, which was configured with the **class-map** command and one or more QoS features. The traffic class is associated with the traffic policy when the **class** command is used. The **class** command must be issued after entering policy-map configuration mode. After entering the **class** command, you are automatically in policy-map class configuration mode, which is where the QoS policies for the traffic policy are defined.

When the bandwidth or priority action is used on any class in a policy map, then there must be a class defined by the **match-any** command, which has a bandwidth or priority action in that policy map. This is to ensure that all traffic can be classified into a default class that has some assigned bandwidth. A minimum bandwidth can be assigned if the class is not expected to be used or no reserved bandwidth is desired for default traffic.

The QoS policies that can be applied in the traffic policy in policy-map class configuration mode are shown in [Example 14-2](#) and [Example 14-3](#).

### Example 14-2 Policy-map syntax

```
policy-map policy-name
no policy-map policy-name
```

### Example 14-3 Class command syntax

```
class class-map-name
no class class-map-name
```

All traffic that fails to meet the matching criteria belongs to the default traffic class. The default traffic class can be configured by the user, but cannot be deleted.

To create a traffic policy, use the commands in [Table 14-2](#) as needed.

**Table 14-2** Traffic Policy Commands

Command	Purpose
Router (config)# <b>policy-map</b> <i>policy-name</i>	Specifies the name of the traffic policy to configure. Names can be a maximum of 40 alphanumeric characters.
Router (config-pmap)# <b>class</b> <i>class-map-name</i>	Specifies the name of a predefined traffic class, which was configured with the <b>class-map</b> command, used to classify traffic to the traffic policy.
Router (config-pmap)# <b>class class-default</b>	Specifies the default class to be created as part of the traffic policy.
Router (config-pmap-c)# <b>bandwidth</b> { <i>bandwidth-kbps</i>   <b>percent</b> <i>percent</i> }	<p>Specifies a minimum bandwidth guarantee to a traffic class in periods of congestion. A minimum bandwidth guarantee can be specified in kbps or by a percentage of the overall available bandwidth.</p> <p>Valid choices for the ML-Series cards are:</p> <ul style="list-style-type: none"> <li>• Rate in kilobits per second</li> <li>• Percent of total available bandwidth (1 to 100)</li> </ul> <p>If multiple classes and bandwidth actions are specified in a single policy map, they must use the same choice in specifying bandwidth (kilobits or percent).</p> <p><b>Note</b> When using the <b>bandwidth</b> command, excess traffic (beyond the configured commit) is allocated any available bandwidth in proportion to the relative bandwidth commitment of its traffic class compared to other traffic classes. Excess traffic from two classes with equal commits has equal access to available bandwidth. Excess traffic from a class with a minimum commit might receive only a minimum share of available bandwidth compared to excess bandwidth from a class with a high commit.</p> <p><b>Note</b> The true configurable bandwidth in kilobits or megabits per second is per port and depends on how the ML-Series card is configured. The <b>show interface</b> command shows the maximum bandwidth of a port (for example, BW 100000 Kbit). The sum of all bandwidth and priority actions applied to the interface, plus the cos priority-mcast bandwidth, is not allowed to exceed the maximum bandwidth of the port.</p>

Table 14-2 Traffic Policy Commands (continued)

Command	Purpose
<pre>Router (config-pmap-c)# <b>police</b> <i>cir-rate-bps</i> <i>normal-burst-byte</i> [<i>max-burst-byte</i>] [<b>pir</b> <i>pir-rate-bps</i>] [<b>conform-action</b> {<b>set-cos-transmit</b>   <b>transmit</b>   <b>drop</b>}] [<b>exceed-action</b> {<b>set-cos-transmit</b>   <b>drop</b>}] [<b>violate-action</b> {<b>set-cos-transmit</b>   <b>drop</b>}]</pre>	<p>Defines a policer for the currently selected class when the policy map is applied to input. Policing is supported only on ingress, not on egress.</p> <ul style="list-style-type: none"> <li>• For <i>cir-rate-bps</i>, specify the average committed information rate (cir) in bits per second (bps). The range is 96000 to 800000000.</li> <li>• For <i>normal-burst-byte</i>, specify the cir burst size in bytes. The range is 8000 to 64000.</li> <li>• (Optional) For <i>maximum-burst-byte</i>, specify the peak information rate (pir) burst in bytes. The range is 8000 to 64000.</li> <li>• (Optional) For <i>pir-rate-bps</i>, specify the average pir traffic rate in bps where the range is 96000 to 800000000.</li> <li>• (Optional) Conform action options are: <ul style="list-style-type: none"> <li>– Set a CoS priority value and transmit</li> <li>– Transmit packet (default)</li> <li>– Drop packet</li> </ul> </li> <li>• (Optional) Exceed action options are: <ul style="list-style-type: none"> <li>– Set a CoS value and transmit</li> <li>– Drop packet (default)</li> </ul> </li> <li>• (Optional) The violate action is only valid if pir is configured. Violate action options are: <ul style="list-style-type: none"> <li>– Set a CoS value and transmit</li> <li>– Drop packet (default)</li> </ul> </li> </ul>

Table 14-2 Traffic Policy Commands (continued)

Command	Purpose
Router (config-pmap-c)# <b>priority</b> <i>kbps</i>	<p>Specifies low latency queuing for the currently selected class. This command can only be applied to an output. When the policy-map is applied to an output, an output queue with strict priority is created for this class. The only valid rate choice is in kilobits per second.</p> <p><b>Note</b> This <b>priority</b> command does not apply to the default class.</p> <p><b>Note</b> When using the priority action, the traffic in that class is given a 100 percent CIR, regardless of the rate entered as the priority rate. To ensure that other bandwidth commitments are met for the interface, a policer must be configured on the input of all interfaces that might deliver traffic to this output class, limiting the peak rate to the priority rate entered.</p> <p><b>Note</b> The true configurable bandwidth in kilobits or megabits per second is per port and depends on how the ML-Series card is configured. The <b>show interface</b> command shows the maximum bandwidth of a port (for example, BW 100000 Kbit). The sum of all bandwidth and priority actions applied to the interface, plus the cos priority-mcast bandwidth, is not allowed to exceed the maximum bandwidth of the port.</p>
Router (config-pmap-c)# <b>set cos</b> <i>cos-value</i>	<p>Specifies a CoS value or values to associate with the packet. The number is in the range from 0 to 7.</p> <p>This command can only be used in a policy-map applied to an output. It specifies the VLAN CoS priority to set for the outbound packets in the currently selected class. If QinQ is used, the top-level VLAN tag is marked. If outbound packets have no VLAN tag, the action has no effect. This action is applied to the packet after any set-cos action done by a policer, and therefore overrides the CoS set by a policer action.</p> <p>If a packet is marked by the policer and forwarded out an interface that also has a set-cos action assigned for the traffic class, the value specified by the police action takes precedence in setting the IEEE 802.1p CoS field.</p> <p>This command also sets the CoS value in the Cisco proprietary RPR header for packets exiting the ML-Series on the Cisco proprietary RPR interface.</p>

## Attaching a Traffic Policy to an Interface

Use the **service-policy** interface configuration command to attach a traffic policy to an interface and to specify the direction in which the policy should be applied (either on packets coming into the interface or packets leaving the interface). Only one traffic policy can be applied to an interface in a given direction.



Use the **no** form of the command to detach a traffic policy from an interface. The **service-policy** command syntax is as follows:

```
service-policy {input | output} policy-map-name
no service-policy {input | output} policy-map-name
```

To attach a traffic policy to an interface, use the following commands in global configuration mode, as needed:

<b>Step 1</b>	Router(config)# <b>interface</b> <i>interface-id</i>	Enters interface configuration mode, and specifies the interface to apply the policy map.  Valid interfaces are limited to physical Ethernet and POS interfaces.  <b>Note</b> Policy maps cannot be applied to SPR interfaces, subinterfaces, port channel interfaces, or Bridge Group Virtual Interfaces (BVI).
<b>Step 2</b>	Router(config-if)# <b>service-policy</b> <b>output</b> <i>policy-map-name</i>	Specifies the name of the traffic policy to be attached to the output direction of an interface. The traffic policy evaluates all traffic leaving that interface.
<b>Step 3</b>	Router(config-if)# <b>service-policy</b> <b>input</b> <i>policy-map-name</i>	Specifies the name of the traffic policy to be attached to the input direction of an interface. The traffic policy evaluates all traffic entering that interface.

## Configuring CoS-Based QoS

The global **cos commit** *cos-value* command allows the ML-Series card to base the QoS treatment for a packet coming in on a network interface on the attached CoS value, rather than on a per-customer-queue policer.

CoS-based QoS is applied with a single global **cos commit** *cos-value* command, as shown in [Table 14-3](#).

**Table 14-3** CoS Commit Command

Command	Purpose
Router(config)# <b>cos-commit</b> <i>cos-value</i>	Labels packets that come in with a CoS equal to or higher than the <i>cos-value</i> as CIR and packets with a lower CoS as DE.

## Monitoring and Verifying QoS Configuration

After configuring QoS on the ML-Series card, the configuration of class maps and policy maps can be viewed through a variety of **show** commands. To display the information relating to a traffic class or traffic policy, use one of the commands in [Table 14-4](#) in EXEC mode, as needed. [Table 14-4](#) describes the commands that are related to QoS status.

**Table 14-4** Commands for QoS Status

Command	Purpose
Router# <b>show class-map</b> <i>name</i>	Displays the traffic class information of the user-specified traffic class.
Router# <b>show policy-map</b>	Displays all configured traffic policies.
Router# <b>show policy-map</b> <i>name</i>	Displays the user-specified policy map.
Router# <b>show policy-map interface</b> <i>interface</i>	Displays configurations of all input and output policies attached to an interface. Statistics displayed with this command are unsupported and show zero.

[Example 14-4](#) show examples of the QoS commands.

#### **Example 14-4** QoS Status Command Examples

```
Router# show class-map
Class Map match-any class-default (id 0)
  Match any
Class Map match-all policer (id 2)
  Match ip precedence 0

Router# show policy-map
Policy Map police_f0
  class policer
    police 1000000 10000 conform-action transmit exceed-action drop

Router# show policy-map interface
FastEthernet0
  service-policy input: police_f0
  class-map: policer (match-all)
    0 packets, 0 bytes
    5 minute offered rate 0 bps, drop rate 0 bps
  match: ip precedence 0
  class-map: class-default (match-any)
    0 packets, 0 bytes
    5 minute offered rate 0 bps, drop rate 0 bps
  match: any
    0 packets, 0 bytes
    5 minute rate 0 bps
```

## QoS Configuration Examples

This section provides the specific command and network configuration examples:

- [Traffic Classes Defined Example, page 14-19](#)
- [Traffic Policy Created Example, page 14-19](#)
- [class-map match-any and class-map match-all Commands Example, page 14-20](#)
- [match spr1 Interface Example, page 14-20](#)
- [ML-Series VoIP Example, page 14-21](#)
- [ML-Series Policing Example, page 14-21](#)
- [ML-Series CoS-Based QoS Example, page 14-22](#)

## Traffic Classes Defined Example

[Example 14-5](#) shows how to create a class map called class1 that matches incoming traffic entering interface fastethernet0.

### Example 14-5 Class Interface Command Examples

```
Router(config)# class-map class1
Router(config-cmap)# match input-interface fastethernet0
```

[Example 14-6](#) shows how to create a class map called class2 that matches incoming traffic with IP-precedence values of 5, 6, and 7.

### Example 14-6 Class IP-Precedence Command Examples

```
Router(config)# class-map match-any class2
Router(config-cmap)# match ip precedence 5 6 7
```



#### Note

If a class-map contains a match rule that specifies multiple values, such as 5 6 7 in this example, then the class-map must be match-any, not the default match-all. Without the match-any class-map, an error message is printed and the class is ignored. The supported commands that allow multiple values are **match cos**, **match ip precedence**, and **match ip dscp**.

[Example 14-7](#) shows how to create a class map called class3 that matches incoming traffic based on bridge group 1.

### Example 14-7 Class Map Bridge Group Command Examples

```
Router(config)# class-map class3
Router(config-cmap)# match bridge-group 1
```

## Traffic Policy Created Example

In [Example 14-8](#), a traffic policy called policy1 is defined to contain policy specifications, including a bandwidth allocation request for the default class and two additional classes—class1 and class2. The match criteria for these classes were defined in the traffic classes, see the [“Creating a Traffic Class” section on page 14-12](#).

### Example 14-8 Traffic Policy Created Example

```
Router(config)# policy-map policy1
Router(config-pmap)# class class-default
Router(config-pmap-c)# bandwidth 1000
Router(config-pmap)# exit

Router(config-pmap)# class class1
Router(config-pmap-c)# bandwidth 3000
Router(config-pmap)# exit

Router(config-pmap)# class class2
Router(config-pmap-c)# bandwidth 2000
Router(config-pmap)# exit
```

## class-map match-any and class-map match-all Commands Example

This section illustrates the difference between the **class-map match-any** command and the **class-map match-all** command. The **match-any** and **match-all** options determine how packets are evaluated when multiple match criteria exist. Packets must either meet all of the match criteria (**match-all**) or one of the match criteria (**match-any**) in order to be considered a member of the traffic class.

[Example 14-9](#) shows a traffic class configured with the **class-map match-all** command.

### Example 14-9 Class Map Match All Command Examples

```
Router(config)# class-map match-all cisco1
Router(config-cmap)# match cos 1
Router(config-cmap)# match bridge-group 10
```

If a packet arrives with a traffic class called cisco1 configured on the interface, the packet is evaluated to determine if it matches the cos 1 and bridge group 10. If both of these match criteria are met, the packet matches traffic class cisco1.

In traffic class called cisco2, the match criteria are evaluated consecutively until a successful match criterion is located. The packet is first evaluated to determine whether cos 1 can be used as a match criterion. If cos 1 can be used as a match criterion, the packet is matched to traffic class cisco2. If cos 1 is not a successful match criterion, then bridge-group 10 is evaluated as a match criterion. Each matching criterion is evaluated to see if the packet matches that criterion. When a successful match occurs, the packet is classified as a member of traffic class cisco2. If the packet matches none of the specified criteria, the packet is classified as a member of the traffic class.

Note that the **class-map match-all** command requires that all of the match criteria must be met in order for the packet to be considered a member of the specified traffic class (a logical AND operator). In the example, cos 1 AND bridge group 10 have to be successful match criteria. However, only one match criterion must be met for the packet in the **class-map match-any** command to be classified as a member of the traffic class (a logical OR operator). In the example, cos 1 OR bridge group 10 OR ip dscp 5 have to be successful match criteria.

[Example 14-10](#) shows a traffic class configured with the **class-map match-any** command.

### Example 14-10 Class Map Match Any Command Examples

```
Router(config)# class-map match-any cisco2
Router(config-cmap)# match cos 1
Router(config-cmap)# match bridge-group 10
Router(config-cmap)# match ip dscp 5
```

## match spr1 Interface Example

In [Example 14-11](#), the SPR interface is specified as a parameter to the **match input-interface** CLI when defining a class-map.

### Example 14-11 Class Map SPR Interface Command Examples


```
Router(config)# class-map spr1-cos1
Router(config-cmap)# match input-interface spr1
Router(config-cmap)# match cos 1
Router(config-cmap)# end
Router# sh class-map spr1-cos1
Class Map match-all spr1-cos1 (id 3)
```

```
Match input-interface SPR1
Match cos 1
```

## ML-Series VoIP Example

Figure 14-7 shows an example of ML-Series QoS configured for VoIP. The associated commands are provided in Example 14-12.

**Figure 14-7 ML-Series VoIP Example**




83181

**Example 14-12 ML-Series VoIP Commands**

```
Router(config)# class-map match-all voip
Router(config-cmap)# match ip precedence 5
Router(config-cmap)# exit
Router(config)# class-map match-any default
Router(config-cmap)# match any
Router(config-cmap)# exit
Router(config)# policy-map pos0
Router(config-pmap)# class default
Router(config-pmap-c)# bandwidth 1000
Router(config-pmap-c)# class voip
Router(config-pmap-c)# priority 1000
Router(config-pmap-c)# interface FastEthernet0
Router(config-if)# ip address 1.1.1.1 255.255.255.0
Router(config-if)# interface POS0
Router(config-if)# ip address 2.1.1.1 255.255.255.0
Router(config-if)# service-policy output pos0
Router(config-if)# crc 32
Router(config-if)# no cdp enable
Router(config-if)# pos flag c2 1
```

## ML-Series Policing Example

Figure 14-8 shows an example of ML-Series policing. The example shows how to configure a policer that restricts traffic with an IP precedence of 0 to 1,000,000 bps. The associated code is provided in Example 14-13.

**Figure 14-8 ML-Series Policing Example****Example 14-13 ML-Series Policing Commands**

```


Router(config)# class-map match-all policer
Router(config-cmap)# match ip precedence 0
Router(config-cmap)# exit
Router(config)# policy-map police_f0
Router(config-pmap)# class policer
Router(config-pmap-c)# police 1000000 10000 conform-action transmit exceed-action drop
Router(config-pmap-c)# interface FastEthernet0
Router(config-if)# service-policy input police_f0

```

## ML-Series CoS-Based QoS Example

Figure 14-9 shows an example of ML-Series CoS-based QoS. The associated code is provided in the examples that follow the figure. The CoS example assumes that the ML-Series cards are configured into an Cisco proprietary RPR and that the ML-Series card POS ports are linked by point-to-point SONET circuits. ML-Series Card A and ML-Series Card C are customer access points. ML-Series Card B is not a customer access point. For more information on configuring Cisco proprietary RPR, see [Chapter 17, “Configuring Cisco Proprietary Resilient Packet Ring.”](#)

Figure 14-9 ML-Series CoS Example



96501

Example 14-14 shows the code used to configure ML-Series card A in Figure 14-9.

**Example 14-14 ML-Series Card A Configuration (Customer Access Point)**

```
ML_Series_A(config)# cos commit 2
ML_Series_A(config)# policy-map Fast5_in
ML_Series_A(config-pmap)# class class-default
ML_Series_A(config-pmap-c)# police 5000 8000 8000 pir 10000 conform-action
set-cos-transmit 2 exceed-action set-cos-transmit 1 violate-action drop
```

Example 14-15 shows the code used to configure ML-Series card B in Figure 14-9.

**Example 14-15 ML-Series Card B Configuration (Not Customer Access Point)**

```
ML_Series_B(config)# cos commit 2
```

Example 14-16 shows the code used to configure ML-Series card C in Figure 14-9.

**Example 14-16 ML-Series Card C Configuration (Customer Access Point)**

```
ML_Series_B(config)# cos commit 2
ML_Series_B(config)# policy-map Fast5_in
ML_Series_B(config-pmap)# class class-default
ML_Series_B(config-pmap-c)# police 5000 8000 8000 pir 10000 conform-action
set-cos-transmit 2 exceed-action set-cos-transmit 1 violate-action drop
```

# Understanding Multicast QoS and Priority Multicast Queuing

ML-Series card QoS supports the creation of two priority classes for multicast traffic in addition to the default multiclass traffic class. Creating a multicast priority queuing class of traffic configures the ML-Series card to recognize an existing CoS value in ingress multicast traffic for priority treatment.

The multicast priority queuing CoS match is based on the “internal” CoS value of each packet. This value is normally the same as the egress CoS value (after policer marking if enabled) but differs in two cases. The internal CoS value is not the same as the egress value when dot1q-tunneling is used. Under dot1q-tunnel, the internal CoS value is always the value of the outer tag CoS, both when entering the dot1q tunnel and leaving the dot1q tunnel. The internal CoS value is also not the same as the egress value if a packet is transported over a VLAN, and the VLAN tag is removed on egress to send the packet untagged. In this case, the internal CoS is the CoS of the removed tag (including ingress policing and marking if enabled).

The **cos priority-mcast** command does not modify the CoS of the multicast packets, but only the bandwidth allocation for the multicast priority queuing class. The command guarantees a minimum amount of bandwidth and is queued separately from the default multicast/broadcast queue.

Creating a multicast priority queuing class allows for special handling of certain types of multiclass traffic. This is especially valuable for multicast video distribution and service provider multicast traffic. For example, a service provider might want to guarantee the protection of their own multicast management traffic. To do this, they could create a multicast priority queuing class on the ML-Series card for the CoS value of the multicast management traffic and guarantee its minimum bandwidth. For multicast video distribution, a multicast priority queuing class on the ML-Series card for the CoS value of the multicast video traffic enables networks to efficiently manage multicast video bandwidth demands on a network shared with VoIP and other Ethernet services.



## Note

Multicast priority queuing traffic uses port-based load-balancing over Cisco proprietary RPR and EtherChannel. Default multicast traffic is load-balanced over Cisco proprietary RPR, but not over EtherChannel. Multicast load balancing maps GigabitEthernet Port 0 to POS Port 0 and GigabitEthernet Port 1 to POS Port 1. Multicast load balancing maps Fast Ethernet Port 0 and all even-numbered Fast Ethernet ports to POS 0 and all odd-numbered Fast Ethernet ports to POS 1.



## Note

Multicast priority queuing bandwidth should not be oversubscribed for sustained periods with traffic from multiple sources. This can result in reduced multicast priority queuing throughput.

Priority multicast feature is not required and is not supported in ML card while it is in IEEE-RPR mode, as in this mode each queue created for a port can handle all of multicast, broadcast and unicast traffic.

## Default Multicast QoS

Default multicast traffic is any multicast traffic (including flooded traffic) that is not classified as multicast priority queuing. The default multicast class also includes broadcast data traffic, control traffic, Layer 2 protocol tunneling, and flooding traffic of the unknown MAC during MAC learning.

With no QoS configured (no multicast priority queuing and no output policy map) on the ML-Series card, the default multicast bandwidth is a 10 percent minimum of total bandwidth.



When bandwidth is allocated to multicast priority queuing but no output policy map is applied, the default multicast congestion bandwidth is a minimum of 10 percent of the bandwidth not allocated to multicast priority queuing.

When an output policy-map is applied to an interface, default multicast and default unicast share the minimum bandwidth assigned to the default class. This default class is also known as the match-any class. The minimum bandwidth of default multicast is 10 percent of the total default class bandwidth.

## Multicast Priority Queuing QoS Restrictions

The following restrictions apply to multicast priority queuing QoS:

- The bandwidth allocation and utilization configured for multicast priority queuing traffic is global and applies to all the ports on the ML-Series card, both POS and Fast Ethernet or Gigabit Ethernet, regardless of whether these ports carry multicast priority queuing traffic. The rate of traffic can be reduced for all ports on the ML-Series card when this feature is configured. Default multicast traffic uses bandwidth only on the ports where it egresses, not globally like multicast priority queuing.
- Multicast priority queuing QoS is supported only for Layer 2 bridging.
- The ML-Series card supports a maximum of two multicast priority queuing classes.
- Unlike the rest of the ML-Series card QoS, multicast priority queuing QoS is not part of the Cisco IOS MQC.
- Priority-mcast bandwidth allocation is per port and the maximum bandwidth configurable on an ML1000-2 with **cos priority-mcast** is 1000 Mbps. But the load-balancing of multicast priority queuing increases the effective bandwidth. For example, with an ML1000-2 with Gigabit EtherChannel (GEC) and STS-24c circuits, the user can allocate 1000 Mbps per port, but will be able to get 2000 Mbps total effective bandwidth due to the load-balancing.

## Configuring Multicast Priority Queuing QoS

To configure a priority class for multicast traffic, use the global configuration **cos priority-mcast** command, defined in [Table 14-5](#).

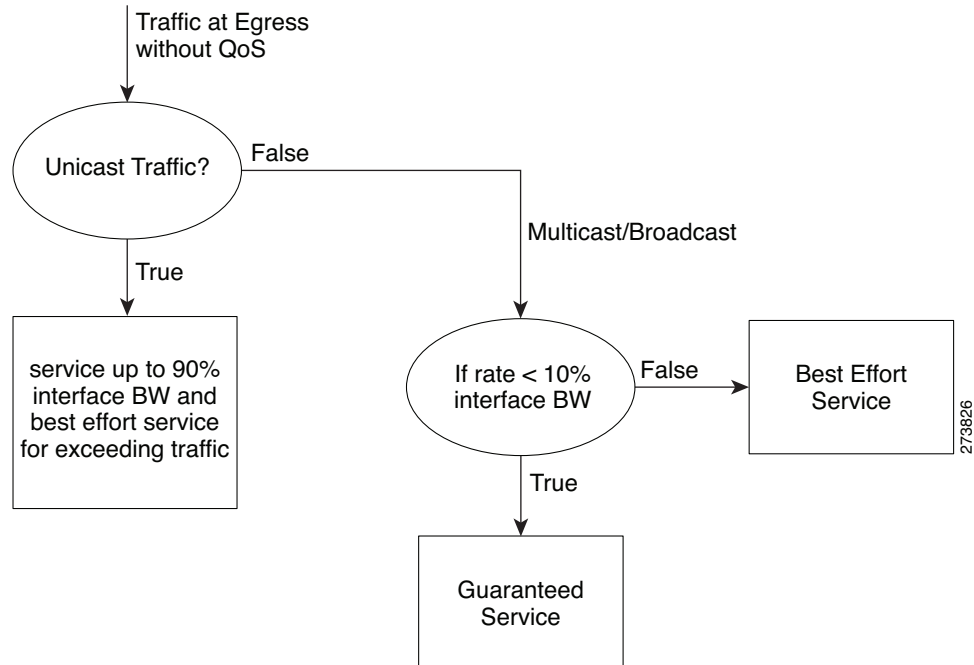
Table 14-5 CoS Multicast Priority Queuing Command

Command	Purpose
<pre>Router (config)# [no] cos priority-mcast cos-value {bandwidth-kbps   mbps bandwidth-mbps   percent percent}</pre>	<p>Creates a priority class of multicast traffic based on a multicast CoS value and specifies a minimum bandwidth guarantee to a traffic class in periods of congestion.</p> <p><i>cos-value</i> specifies the CoS value of multicast packets that will be given the bandwidth allocation. The value matches only a single CoS of traffic (not a range). The supported CoS range is 0 to 7.</p> <p>A minimum bandwidth guarantee can be specified in kbps, in Mbps, or by a percentage of the overall available bandwidth.</p> <p>Valid choices for the ML-Series card are:</p> <ul style="list-style-type: none"> <li>• Rate in kilobits per second</li> <li>• Rate in megabits per second</li> <li>• Percent of total available port bandwidth (1 to 100)</li> </ul> <p>Reentering the command with the same <i>cos-value</i> but a different bandwidth rate will modify the bandwidth of the existing class.</p> <p>Reentering the command with a different <i>cos-value</i> creates a separate multicast priority queuing class with a maximum of two multicast priority queuing classes.</p> <p>The <b>no</b> form of this command removes the multicast priority queuing class.</p> <p><b>Note</b> The true configurable bandwidth in kilobits or megabits per second is per port and depends on how the ML-Series card is configured. The <b>show interface</b> command shows the maximum bandwidth of a port (for example, BW 100000 Kbit). The sum of all bandwidth and priority actions applied to the interface, plus the <b>cos priority-mcast</b> bandwidth, is not allowed to exceed the maximum bandwidth of the port.</p> <p><b>Note</b> Attempting to configure a priority-mcast bandwidth that exceeds the true configurable bandwidth on any port will cause the priority-mcast configuration change to fail, and the multicast priority queuing bandwidth guarantee will not be changed.</p>

## QoS not Configured on Egress

The QoS bandwidth allocation of multicast and broadcast traffic is handled separately from unicast traffic. On each interface, the aggregate multicast and broadcast traffic are given a fixed bandwidth commit of 10% of the interface bandwidth. This is the optimum bandwidth that can be provided for traffic exceeding 10% of the interface bandwidth.

**Figure 14-10** QoS not Configured on Egress



## ML-Series Egress Bandwidth Example

This section explains with examples the utilization of bandwidth across different queues with or without Priority Multicast.

### Case 1: QoS with Priority and Bandwidth Configured Without Priority Multicast

Strict Priority Queue is always serviced first. The remaining interface bandwidth is utilized to service other configured traffic.

In the following example, after servicing unicast customer\_voice traffic, the remaining interface bandwidth is utilized for other WRR queues such as customer\_core\_traffic, customer\_data, and class-default in the ratio of 1:3:5.

At any given time, the sum of the bandwidth assigned cannot exceed the interface bandwidth (in kbps). The bandwidth share allocated to class-default will be utilized by default unicast traffic (in this example, unicast traffic with CoS values other than 2, 5, 7) and all multicast/broadcast traffic (all CoS values). The default unicast and all multicast/broadcast traffic will be serviced in the ratio of 9:1.

For example, if 18x bandwidth is available after servicing priority unicast traffic (CoS 5), then the remaining bandwidth will be allocated as follows:

Unicast traffic with CoS 2 : 2x

Unicast traffic with CoS 7: 6x

Unicast default (without CoS 2, CoS 5, CoS 7): 9x

All multicast/broadcast (any CoS value): 1x

**Example 14-17 QoS with Priority and Bandwidth Configured without Priority Multicast**

```

!
class-map match-all customer_voice
  match cos 5
class-map match-all customer_data
  match cos 7
class-map match-all customer_core_traffic
  match cos 2
!
!
policy-map policy_egress_bandwidth
  class customer_core_traffic
    bandwidth 1000
  class customer_voice
    priority 1000
  class customer_data
    bandwidth 3000
  class class-default
    bandwidth 5000
!
!
interface POS0
  no ip address
  crc 32
  service-policy output policy_egress_bandwidth
!

```

## Case 2: QoS with Priority and Bandwidth Configured with Priority Multicast

In this case, only multicast traffic of CoS 3 is allocated a guaranteed bandwidth. This multicast traffic will now participate in the queue along with other WRR queues. After servicing the customer\_voice traffic, the remaining interface bandwidth is utilized for WRR queues, such as customer\_core\_traffic, customer\_data, class-default, and multicast CoS 3 traffic in the ratio of 1:3:5:2.

At any given time, the sum of the bandwidth assigned cannot exceed the interface bandwidth (in kbps).

**Example 14-18 QoS with Priority and Bandwidth configured with Priority Multicast**

```

cos priority-mcast 3 2000
!
class-map match-all customer_voice
  match cos 5
class-map match-all customer_data
  match cos 7
class-map match-all customer_core_traffic
  match cos 2
!
!

```

```

policy-map policy_egress_bandwidth
  class customer_core_traffic
    bandwidth 1000
  class customer_voice
    priority 1000
  class customer_data
    bandwidth 3000
  class class-default
    bandwidth 5000
!
!
interface POS0
  no ip address
  crc 32
  service-policy output policy_egress_bandwidth
!

```

## Understanding CoS-Based Packet Statistics

Enhanced performance monitoring displays per-CoS packet statistics on the ML-Series card interfaces when CoS accounting is enabled. Per-CoS packet statistics are only supported for bridged services, not for IP routing or Multiprotocol Label Switching (MPLS). CoS-based traffic utilization is displayed at the Fast Ethernet or Gigabit Ethernet interface or subinterface (VLAN) level, or at the POS interface level. It is not displayed at the POS subinterface level. Cisco proprietary RPR statistics are not available at the SPR interface level, but statistics are available for the individual POS ports that make up the SPR interface. EtherChannel (port-channel) and BVI statistics are available only at the member port level. [Table 14-6](#) shows the types of statistics available at specific interfaces.

**Table 14-6** Packet Statistics on ML-Series Card Interfaces

Statistics Collected	Gigabit/Fast Ethernet Interface	Gigabit/Fast Ethernet Subinterface (VLAN)	POS Interface	POS Subinterface
Input—Packets and Bytes	Yes	Yes	No	No
Output—Packets and Bytes	Yes	Yes	No	No
Drop Count—Packets and Bytes <sup>1</sup>	Yes	No	Yes	No

1. Drop counts only include discards caused by output congestion and are counted at the output interface.

CoS-based packet statistics are available through the Cisco IOS CLI and Simple Network Management Protocol (SNMP), using an extension of the CISCO-PORT-QOS MIB. They are not available through CTC.

## Configuring CoS-Based Packet Statistics



### Note

CoS-based packet statistics require the enhanced microcode image to be loaded onto the ML-Series card.

**Note**

For IEEE 802.1Q (QinQ) enabled interfaces, CoS accounting is based only on the CoS value of the outer metro tag imposed by the service provider. The CoS value inside the packet sent by the customer network is not considered for CoS accounting.

For information on the enhanced microcode image, see the “[Multiple Microcode Images](#)” section on [page 3-11](#).

To enable CoS-based packet statistics on an interface, use the interface configuration level command defined in [Table 14-7](#).

**Table 14-7 CoS-Based Packet Statistics Command**

Command	Purpose
Router(config-if)# <b>cos accounting</b>	Enables CoS-based packet statistics to be recorded at the specific interface and for all the subinterfaces of that interface. This command is supported only in interface configuration mode and not in subinterface configuration mode.  The <b>no</b> form of the command disables the statistics.

After configuring CoS-based packet statistics on the ML-Series card, the statistics can be viewed through a variety of **show** commands. To display this information, use one of the commands in [Table 14-8](#) in EXEC mode.

**Table 14-8 Commands for CoS-Based Packet Statistics**

Command	Purpose
Router# <b>show interface</b> <i>type number cos</i>	Displays the CoS-based packet statistics available for an interface.
Router# <b>show interface</b> <i>type number.subinterface-number cos</i>	Displays the CoS-based packet statistics available for a FastEthernet or Gigabit Ethernet subinterface. POS subinterfaces are not eligible.

[Example 14-19](#) shows examples of these commands.

**Example 14-19 Commands for CoS-Based Packet Statistics Examples**

```
Router# show interface gigabitethernet 0.5 cos
GigabitEthernet0.5
  Stats by Internal-Cos
  Input: Packets      Bytes
    Cos 0: 31        2000
    Cos 1:
    Cos 2: 5         400
    Cos 3:
    Cos 4:
    Cos 5:
    Cos 6:
    Cos 7:
  Output: Packets     Bytes
    Cos 0: 1234567890 1234567890
    Cos 1: 31         2000
    Cos 2:
```

```

Cos 3:
Cos 4:
Cos 5:
Cos 6: 10          640
Cos 7:

```

```
Router# show interface gigabitethernet 0 cos
```

```

GigabitEthernet0
  Stats by Internal-Cos
  Input: Packets      Bytes
    Cos 0: 123        3564
    Cos 1:
    Cos 2: 3          211
    Cos 3:
    Cos 4:
    Cos 5:
    Cos 6:
    Cos 7:
  Output: Packets     Bytes
    Cos 0: 1234567890 1234567890
    Cos 1: 3           200
    Cos 2:
    Cos 3:
    Cos 4:
    Cos 5:
    Cos 6: 1           64
    Cos 7:
  Output: Drop-pkts   Drop-bytes
    Cos 0: 1234567890 1234567890
    Cos 1:
    Cos 2:
    Cos 3:
    Cos 4:
    Cos 5: 1           64
    Cos 6: 10          640
    Cos 7:

```

```
Router# show interface pos0 cos
```

```

POS0
  Stats by Internal-Cos
  Output: Drop-pkts   Drop-bytes
    Cos 0: 12          1234
    Cos 1: 31          2000
    Cos 2:
    Cos 3:
    Cos 4:
    Cos 5:
    Cos 6: 10          640
    Cos 7:

```

## Understanding IP SLA

Cisco IP SLA, formerly known as the Cisco Service Assurance Agent, is a Cisco IOS feature to assure IP service levels. Using IP SLA, service provider customers can measure and provide service level agreements, and enterprise customers can verify service levels, verify outsourced service level agreements, and understand network performance for new or existing IP services and applications. IP SLAs use unique service level assurance metrics and methodology to provide highly accurate, precise service level assurance measurements.

Depending on the specific IP SLAs operation, statistics of delay, packet loss, jitter, packet sequence, connectivity, path, server response time, and download time are monitored within the Cisco device and stored in both CLI and SNMP MIBs. The packets have configurable IP and application layer options such as source and destination IP address, User Datagram Protocol (UDP)/TCP port numbers, a ToS byte (including DSCP and IP Prefix bits), Virtual Private Network (VPN) routing/forwarding instance (VRF), and URL web address.

IP SLAs use generated traffic to measure network performance between two networking devices such as routers. IP SLAs starts when the IP SLAs device sends a generated packet to the destination device. After the destination device receives the packet, and depending on the type of IP SLAs operation, the device will respond with time-stamp information for the source to make the calculation on performance metrics. An IP SLAs operation is a network measurement to a destination in the network from the source device using a specific protocol such as UDP for the operation.

Because IP SLA is accessible using SNMP, it also can be used in performance monitoring applications for network management systems (NMSs) such as CiscoWorks2000 (CiscoWorks Blue) and the Internetwork Performance Monitor (IPM). IP SLA notifications also can be enabled through Systems Network Architecture (SNA) network management vector transport (NMVT) for applications such as NetView.

For general IP SLA information, refer to the Cisco IOS IP Service Level Agreements technology page at <http://www.cisco.com/warp/public/732/Tech/nmp/ipsla>. For information on configuring the Cisco IP SLA feature, see the “Network Monitoring Using Cisco Service Assurance Agent” chapter of the *Cisco IOS Configuration Fundamentals Configuration Guide, Release 12.2*.

## IP SLA on the ML-Series

The ML-Series card has a complete IP SLA Cisco IOS subsystem and offers all the normal features and functions available in Cisco IOS Release 12.2S. It uses the standard IP SLA Cisco IOS CLI commands. The SNMP support will be equivalent to the support provided in the IP SLA subsystem 12.2(S), which is the rttMon MIB.

## IP SLA Restrictions on the ML-Series

The ML-Series card supports only features in the Cisco IOS 12.2S branch. It does not support functions available in later Cisco IOS versions, such as the IP SLA accuracy feature or the enhanced Cisco IOS CLI support with updated IP SLA nomenclature.

Other restrictions are:

- Setting the CoS bits is supported, but set CoS bits are not honored when leaving or entering the CPU when the sender or responder is an ONS 15454, ONS 15454 SDH, or ONS 15310-CL platform. Set CoS bits are honored in intermediate ONS nodes.
- On Cisco proprietary RPR, the direction of the data flow for the IP SLA packet might differ from the direction of customer traffic.
- The system clock on the ML-Series card synchronizes with the clock on the TCC2/TCC2P card. Any NTP server synchronization is done with the TCC2/TCC2P card clock and not with the ML-Series card clock.
- The average Round Trip Time (RTT) measured on an ML-Series IP SLA feature is more than the actual data path latency. In the ML-Series cards, IP SLA is implemented in the software. The IP SLA messages are processed in the CPU of the ML-Series card. The latency time measured includes the



network latency and CPU processing time. For very accurate IP SLA measurements, it is recommended that a Cisco Router or Switch be used as an external probe or responder to measure the RTT of the ML-Series cards in a network.

