

# Configuring Circuits for QoS

---

## SYSTEM ADMINISTRATOR GUIDE

## **Copyright**

© Ericsson AB 2010-2011. All rights reserved. No part of this document may be reproduced in any form without the written permission of the copyright owner.

## **Disclaimer**

The contents of this document are subject to revision without notice due to continued progress in methodology, design and manufacturing. Ericsson shall have no liability for any error or damage of any kind resulting from the use of this document.

## **Trademark List**

**SmartEdge** is a registered trademark of Telefonaktiebolaget LM Ericsson.



# Contents

<b>1</b>	<b>Overview</b>	<b>1</b>
1.1	Circuit Configuration with QoS Policies	1
1.2	PD QoS Priority and Drop Precedence	5
1.3	Propagation of QoS Priority Settings Across Layer 3 and Layer 2 Networks	8
1.4	Circuit Groups	17
1.5	802.3ad LAG Support	22
<b>2</b>	<b>Configuration and Operations Tasks</b>	<b>25</b>
2.1	Configuration Guidelines	25
2.2	Configure an ATM PVC for QoS	25
2.3	Configure an Ethernet Circuit for QoS	26
2.4	Configure a POS Circuit for QoS	30
2.5	Configure Cross-Connected Circuits for QoS	31
2.6	Configure a Subscriber Circuit for QoS	31
2.7	Configure QoS Propagation (Optional)	32
2.8	Configure L2TP for QoS	34
2.9	Configure MPLS for QoS	36
2.10	Attach QoS Policies to a Circuit Group and Assign Members to the Group	37
2.11	Configuring Virtual Port Circuit Groups	38
2.12	Configure Nested Circuit Groups	38
2.13	Configure Ingress Port Propagation	39
2.14	Operations Tasks	40
<b>3</b>	<b>Configuration Examples</b>	<b>43</b>
3.1	Attaching Rate- and Class-Limiting Policies	43
3.2	Attaching Scheduling Policies	44
3.3	Propagating QoS	46
3.4	Attaching QoS Policies to Circuit Groups	47
3.5	QoS on 802.3ad LAGs	48





# 1 Overview

This document provides an overview of the quality of service (QoS) features supported on the SmartEdge router for circuits and describes the tasks used to configure, monitor, and administer these features on circuits. This document also provides examples of QoS configurations for circuits.

**Note:** In this document, the term, circuit, refers to a port, channel, permanent virtual circuit (PVC), or link group.

This document applies to both the Ericsson SmartEdge® and SM family routers. However, the software that applies to the SM family of systems is a subset of the SmartEdge OS; some of the functionality described in this document may not apply to SM family routers.

For information specific to the SM family chassis, including line cards, refer to the SM family chassis documentation.

For specific information about the differences between the SmartEdge and SM family routers, refer to the Technical Product Description *SM Family of Systems* (part number 5/221 02-CRA 119 1170/1) in the **Product Overview** folder of this Customer Product Information library.

For information about other QoS configuration tasks and commands, see the following documents:

- *Configuring Rate-Limiting and Class-Limiting*—Rate- and class-limiting features (metering and policing policies)
- *Configuring Queuing and Scheduling*—Queuing and scheduling features (queuing policies (also known as scheduling policies))

The Internet provides only best-effort service, offering no guarantees on when or whether a packet is delivered to the receiver. However, the SmartEdge router offers QoS differentiation based on traffic type and application. QoS policies create and enforce levels of service and bandwidth rates, and prioritize how packets are scheduled through egress queues.

## 1.1 Circuit Configuration with QoS Policies

You can attach both a metering and a policing policy to any port or channel to ATM and 802.1Q PVCs, to subscriber sessions, and to link groups. QoS metering and policing policies are described in *Configuring Rate-Limiting and Class-Limiting*.



**Note:** Metering, policing, and forwarding policies do not currently support policy ACLs for classification of IPv6 traffic. When IPv6 traffic is subject to a metering, policing, or forwarding policy that was configured using an IPv4 policy ACL, IPv6 packets do not match any of the classes, but are subject to the configured policy-level enforcement.

Child circuits can inherit the QoS metering and policing policies attached to a parent circuit below which the child circuits are configured. To enable a child circuit to inherit the metering or policing policy of its parent, configure the keyword `inherit` or `hierarchical` on the policy reference of the parent circuit. If you attach a different metering or policing policy to a child circuit, that policy overrides the metering or policing policy attached to the parent circuit unless the policy reference applied to the parent is configured with the `hierarchical` keyword. This keyword allows a child circuit that has its own metering or policing policy to inherit the policy of its parent and still be subject to its own policy. For more information about policy inheritance, see *Policy Inheritance* in *Configuring Rate-Limiting and Class-Limiting*.

The following types of inheritance are supported:

- 802.1Q PVC or tunnel from a parent Ethernet port or access link group
- 802.1Q PVC from a parent 802.1Q tunnel
- PPPoE and CLIPs sessions from a parent 802.1Q PVC, port, or access link group
- PPP and PPPoE sessions from a parent ATM PVC.
- Circuit-group members and their children from a circuit-group

You can attach a scheduling policy to individual circuits (that are not cross-connected); however, the type of scheduling policy depends on the type of traffic card. QoS scheduling policies are described in *Configuring Queuing and Scheduling*. QoS queuing policies are automatically inherited by the children of the circuit to which the policy is applied unless the child circuit has its own queuing policy binding.

**Note:** Inheritance can span multiple levels. For example, a policy configured on a port inherited by a PVC in turn is inherited by PPPoX sessions configured under the PVC unless they have a specific policy applied.

You can also attach metering, policing, and scheduling policies to subscriber circuits; the type of scheduling policy depends on the type of traffic card on which the subscriber session is initiated. Layer 2 Tunneling Protocol (L2TP) network server (LNS) subscriber sessions are limited to priority weighted fair queuing (PWFQ) policies. To attach a QoS policy of any type to a subscriber circuit, you attach it to the subscriber record or profile. The system applies the policy to the subscriber circuit (port, channel, or PVC) on which the session is initiated.



Table 1 lists the traffic cards and their circuits to which QoS scheduling policies can be attached.

**Note:** Certain restrictions apply to the attachment of a QoS scheduling policy to a port, channel, or PVC; for detailed usage guidelines for each type of circuit and policy, see the description for the *qos policy queuing* command (in the appropriate circuit configuration mode).

Restrictions also apply to the configuration of the circuit; for information about configuring traffic card ports, channels, and circuits, see *Configuring ATM, Ethernet, and POS Ports, Configuring Circuits, and Configuring Cross-Connections*.

*Table 1 QoS Scheduling Policy Support for SmartEdge Traffic Cards*

Type	Traffic Card or MIC	Circuit	Policy
Second-generation ATM OC	ATM OC-3c/STM-1c IR (2-port) MIC, SmartEdge 100	ATM PVC	ATMWfq
	ATM OC-3c/STM-1c (8-port)	ATM PVC	ATMWfq
	ATM OC-12c/STM-4c (2-port)	ATM PVC	ATMWfq
Channelized OC	Channelized OC-3/STM-1 (8/4-port) or OC-12/STM-4 (2/1-port)	Port, 802.1Q tunnel	PWFq
POS	POS OC-3c/STM-1c (8-port)	Port, Frame Relay PVC	PWFq or MDRR
	POS OC-12c/STM-4c (4-port)	Port, Frame Relay PVC	PWFq or MDRR
	POS OC-48c/STM-16c (4-port)	Port, Frame Relay PVC	PWFq or MDRR
	OC-192c/STM-64c (1-port)	Port, Frame Relay PVC	PWFq or MDRR
	10/100 TX (12-port) MIC, SmartEdge 100	Port, 802.1Q tunnel, 802.1Q PVC	PWFq or MDRR
	10/100 FX SFP-based (12-port) MIC, SmartEdge 100	Port, 802.1Q tunnel, 802.1Q PVC	PWFq or MDRR
Fast Ethernet-Gigabit Ethernet	Fast Ethernet–Gigabit Ethernet (60-port FE, 2-port GE)	Port, 802.1Q tunnel, 802.1Q PVC	PWFq or MDRR



*Table 1 QoS Scheduling Policy Support for SmartEdge Traffic Cards*

<b>Type</b>	<b>Traffic Card or MIC</b>	<b>Circuit</b>	<b>Policy</b>
Gigabit Ethernet with traffic management	Gigabit Ethernet 1020 (10-port) (ge-10-port)	Port, 802.1Q tunnel, 802.1Q PVC, hierarchical node	PWFQ or MDRR
	Gigabit Ethernet 1020 (20-port) (ge-20-port)	Port, 802.1Q tunnel, 802.1Q PVC, hierarchical node	PWFQ or MDRR
	Gigabit Ethernet (5-port) (ge-5-port)	Port, 802.1Q tunnel, 802.1Q PVC, hierarchical node	PWFQ or MDRR
	Gigabit Ethernet DDR (10-port) (ge2-10-port)	Port, 802.1Q tunnel, 802.1Q PVC, hierarchical node	PWFQ or MDRR
	Gigabit Ethernet 4 DDR (20-port) (ge4-20-port)	Port, 802.1Q tunnel, 802.1Q PVC, hierarchical node	PWFQ or MDRR
	Gigabit Ethernet (2-port, copper) native port, SmartEdge 100	Port, 802.1Q tunnel, 802.1Q PVC, hierarchical node	PWFQ or MDRR
	Gigabit Ethernet (2-port, optical) native port, SmartEdge 100	Port, 802.1Q tunnel, 802.1Q PVC, hierarchical node	PWFQ or MDRR
10-Gigabit Ethernet	10 Gigabit Ethernet Gigabit Ethernet (10GE) (1-port)	Port, 802.1Q tunnel, 802.1Q PVC	MDRR
	10 Gigabit Ethernet DDR (10GE) (4-port)	Port, 802.1Q tunnel, 802.1Q PVC	PWFQ or MDRR
10 Gigabit Ethernet / POS DDR	10 Gigabit Ethernet/OC-192c DDR (1-port)	Port, Frame Relay PVC, 802.1Q tunnel, 802.1Q PVC	MDRR



## 1.2 PD QoS Priority and Drop Precedence

The SmartEdge router associates a six-bit internal Packet Descriptor (PD) for the QoS value that is initialized when a packet is received on ingress. The PD remains associated with the packet as it transits the router. This six-bit internal PD value is referred to as the *PD QoS Codepoint*. The upper three bits hold the PD QoS priority value (ppp), and the lower three bits hold the PD QoS drop precedence value (ddd):

- The PD QoS priority value can be set from 0 to 7, where 0 receives the highest priority treatment by default and 7 the lowest.
- The PD QoS drop precedence value can be set from 0 to 7, where 0 receives the highest chance of being dropped by default and 7 the lowest chance.

For a diagram of the translation from QoS priority markings into PD QoS Code Point markings on ingress and back into external markings on egress, see Figure 1.

Default initial values depend on the circuit's binding type:

- Layer 3 (IP-routed) circuits—The upper six bits of the IP Differentiated Services Code Point (DSCP) or type of service (TOS) setting is copied to PD QoS Code Point; see Table 2.
- MPLS circuits—The EXP value from the relevant received MPLS label is translated to the upper three bits of PD QoS priority and the lower three bits of the PD QoS drop precedence bits are set to zero; see Table 3.
- Other non-IP-routed circuits—The upper three priority bits from the incoming packet are translated to the PD QoS priority (lowest priority value) and the lower three bits are translated to the PD QoS drop precedence (to zero). For example, for the default translation for 802.1Q VLAN circuits, see Table 4.

*Table 2 Default Translation of IP DSCP (or TOS) Priority to PD QoS Priority Values*

IP DSCP or TOS Value	PD QoS Priority Value	PD QoS Drop Precedence Value
111xxxb	0	0
110xxxb	1	0
101xxxb	2	0
100xxxb	3	0
011xxxb	4	0
010xxxb	5	0
001xxxb	6	0
000xxxb	7	0

*Table 3 Default Translation of MPLS EXP Priority Bits to PD QoS Values*

<b>MPLS EXP Value</b>	<b>PD QoS Priority Value</b>	<b>PD QoS Drop Precedence Value</b>
7	0	0
6	1	0
5	2	0
4	3	0
3	4	0
2	5	0
1	6	0
0	7	0

*Table 4 Default Translation of 802.1P Priority to PD QoS Settings*

<b>802.1P Value</b>	<b>PD QoS Priority Value</b>	<b>PD QoS Drop Precedence Value</b>
7	0	0
6	1	0
5	2	0
4	3	0
3	4	0
2	5	0
1	6	0
0	7	0

Table 5 shows how the DSCP values are translated to PD QoS priority and drop precedence values using the default settings. For example, a DSCP value of CS7 is translated to the PD-QoS priority bit as 0.

*Table 5 Default Translation of DSCP Values to PD-QoS Priority and Drop Precedence Values*

<b>DSCP Value</b>	<b>PD-QoS Priority Value</b>	<b>PD-QoS Drop Precedence Value</b>
CS7	0	0
CS6	1	0
CS5	2	0
CS4	3	0
CS3	4	0
CS2	5	0



*Table 5 Default Translation of DSCP Values to PD-QoS Priority and Drop Precedence Values*

CS1	6	0
CS0/DF	7	0
EF	2	6
AF41	3	2
AF42	3	4
AF43	3	6
AF31	4	2
AF32	4	4
AF33	4	6
AF21	5	2
AF22	5	4
AF23	5	6
AF11	6	2
AF12	6	4
AF13	6	6

You can modify the PD QoS settings with the following configurations:

- `qos priority` command
- `qos propagate from ip, ethernet, mpls, atm, and l2tp` commands
- Ingress class maps
- `mark priority` command and its conform, exceed, and violate variants
- `mark dscp` command and its conform, exceed, and violate variants
- `mark precedence` command and its conform, exceed, and violate variants

The following configurations reference PD QoS priority and drop precedence values:

- `qos propagate to ip, ethernet, mpls, atm, and l2tp` commands
- Egress class maps
- Class definitions
- Congestion avoidance maps
- Queue-maps (priority only)



- `exceed drop` command
- `qos-priority` and its metering or policing action (priority only)

### 1.3 Propagation of QoS Priority Settings Across Layer 3 and Layer 2 Networks

You can configure how the priority and drop precedence settings in incoming Layer 2 and Layer 3 packets are translated to PD QoS settings in packets as they transit the router, and back into those settings in egress packets.

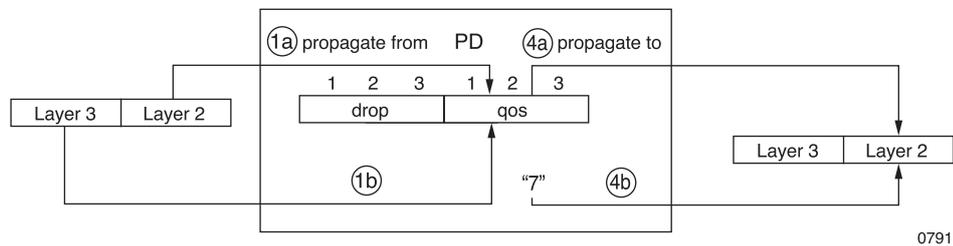


Figure 1 Propagation of QoS Across Layer 3 and Layer 2 Networks (791)

The SmartEdge router uses these PD fields to perform the following functions for an incoming Layer 2 packet:

1. Depending on configuration for an inbound circuit protocol, the SmartEdge OS populates the PD for this packet, using one of the following functions:
  - For incoming Layer 2 packets (arriving over MPLS, 802.1Q, or L2TP protocols)—If a QoS `propagate from` command is configured for the Layer 2 protocol, the SmartEdge OS translates the priority bits from the Layer 2 header to the PD QoS priority.

If no classification map is configured for the `propagate from` command (specifically setting the mapping), the default translation is used, depending on the Layer 2 protocol, to populate the PD QoS priority.

- For incoming Layer 3 packets—If a QoS `propagate from` command is configured for a Layer 3 protocol, the SmartEdge router translates the three most-significant DSCP bits from the Layer 3 header in the incoming packet to the PD QoS priority and the drop precedence settings in that header to the PD QoS drop precedence.

With a class map configured, the SmartEdge OS specifies a custom value mapping for propagating the Differentiated Services Code Point (DSCP) bits in the IP packet header to the packet descriptor (PD) priority bits for incoming IP packets. With no class map configured the default translation is used.

2. If a QoS policing policy, which can include a policy access control list (ACL), that includes a `mark` command (of any type) is attached to the



inbound circuit, the SmartEdge router modifies the bits in the priority and drop fields in the PD based on the policy.

A decision is also made whether to forward the incoming Layer 3 packet to the outbound circuit for further QoS processing.

3. If a QoS metering policy (which can include a policy ACL) that includes a mark command (of any type) is attached to the outbound circuit, the SmartEdge router modifies the bits in the qos and drop fields in the PD based on the policy.
4. The SmartEdge router encapsulates the Layer 3 packet in a Layer 2 packet, using one of the following functions:
  - If a QoS propagate to command is configured for the Layer 2 protocol, the SmartEdge router copies the priority field in the PD to the priority bits in the Layer 2 header.
  - If no QoS propagate to command is configured, the SmartEdge router sets the priority bits in the Layer 2 header to the default (lowest) priority.
5. The SmartEdge router uses the priority field in the PD to determine the egress queue for the outgoing packet, and the drop precedence field to determine the relative priority within that queue.

The following sections further describe QoS propagation:

### 1.3.1 Propagation of QoS from IP to ATM

The CLP bit in the ATM header of a cell provides a method of controlling the discarding of cells in a congested ATM environment. A CLP bit can be set to either 0 (default) or 1, and ATM cells with setting of 1 are discarded before cells with a setting of 0. When you use the `clpbit propagate qos to atm` command to propagate the QoS classification values from the internal PD to the CLP bit in cells transmitted over ATM PVCs for outgoing packets, the PD value determines when the CLP bit should be set and which ATM cells to discard in an ATM congested network. DSCP bits are mapped to the ATM CLP bit as described in Table 6.

Table 6 QoS PD Value to ATM CLP Bit Mapping

PD Priority Value	PD Drop-Precedence Value	AF Label	ATM CLP Bit
7	N/A	Network Control	0
6	N/A	Reserved	0
5	N/A	EF	0
4	0-2	AF41	0
4	3-7	AF42, AF43	1



Table 6 QoS PD Value to ATM CLP Bit Mapping

PD Priority Value	PD Drop-Precedence Value	AF Label	ATM CLP Bit
3	0-2	AF31	0
3	3-7	AF32, AF33	1
2	0-2	AF21	0
2	3-7	AF22, AF23	1
1	0-2	AF11	0
1	3-7	AF12, AF13	1
0	N/A	DF	1

**Note:** This default mapping can be modified using the `clpbit propagate qos to atm` command (in ATM profile configuration mode) in conjunction with an ATM egress class map.

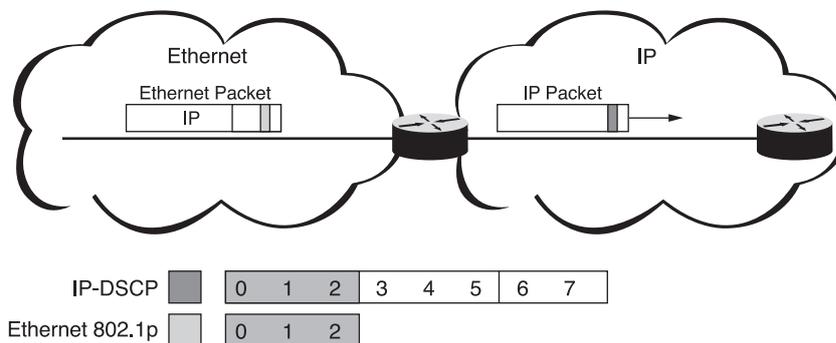
**Note:** You can also use the `mark dscp`, `mark priority`, and `mark precedence` commands (in metering policy or policing policy configuration mode) to indirectly set the ATM CLP bit when using the `clpbit propagate qos to atm` command to propagate the DSCP bits from IP packets to the CLP bit.

### 1.3.2 Propagation of QoS Between IP and Ethernet

802.1p priority is carried in virtual LAN (VLAN) tags defined in IEEE 802.1p. The VLAN tag carries one of eight priority values recognizable by Layer 2 devices in a three-bit field. These three bits are called 802.1p prioritization bits, or "p bits". P-bit marking determines the service level the packet receives when crossing an 802.1p-enabled network segment. DSCP priority bits are mapped to p bits, in either or both directions, depending on whether you configure the `qos propagate from ethernet` and `qos propagate to ethernet` commands (in dot1q profile configuration mode). As shown in Figure 2, the following steps occur for an incoming 802.1Q packet:

1. If you use the `propagate qos from ethernet` command (in dot1q profile configuration mode), and you do not specify a classification map, as an 802.1Q packet enters the SmartEdge router, its 802.1p bits are copied to the PD QoS priority field and DSCP bits.

If you do specify a classification map, the 802.1p bits are mapped to the PD QoS value and the DSCP bits are not affected.



0789

**Figure 2 Propagation of QoS Between IP and Ethernet (789)**

If you use the `propagate qos to ethernet` command (in dot1q profile configuration mode), when the SmartEdge router prepares to forward a packet on an 802.1Q virtual LAN (VLAN), the PD priority value is copied to the 802.1p field of the outgoing packet.

If you create a classification map using the `qos class-map` command and reference it in `propagate qos to ethernet` syntax, the PD priority value is mapped to the 802.1p field, rather than copied.

For outgoing Virtual Router Redundancy Protocol (VRRP) control packets, the SmartEdge router automatically writes a value of 0 into the control packet. However, the SmartEdge router supports 802.1p marking for locally-generated VRRP advertisements, making VRRP packets subject to egress QoS maps. Any configured egress Ethernet class map can be used to map the internal priority of the packet to an 802.1p bit value.

### 1.3.3

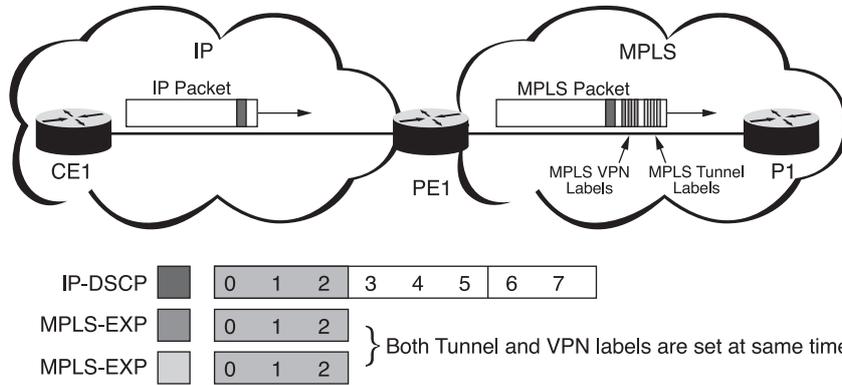
### Propagation of QoS Between MPLS and IP

MPLS EXP bits use one of eight priority values (3 bits in length), recognizable by Layer 2 devices. This marking determines the service level the packet receives when crossing an MPLS-enabled network segment. On ingress, MPLS EXP values are mapped to PD QoS priority values, by default. You can modify the default mapping by doing one of the following:

- Specifying to copy the upper three EXP bits to the DSCP bits in the IP header (and clear the lower three bits) using the `propagate qos from mpls` command without the `class-map map-name` construct (in MPLS router configuration mode).
- Specifying a custom mapping using the `propagate qos from mpls` command with the `class-map map-name` construct.
- Specifying to copy the IP header DSCP value to the PD QoS priority value using the `egress prefer dscp-qos` command (in MPLS router configuration mode).



On egress, if you use the `propagate qos to mpls` command (in MPLS router configuration mode), PD QoS priority bits are mapped to MPLS EXP bits if you specify a classification map; see Figure 3.



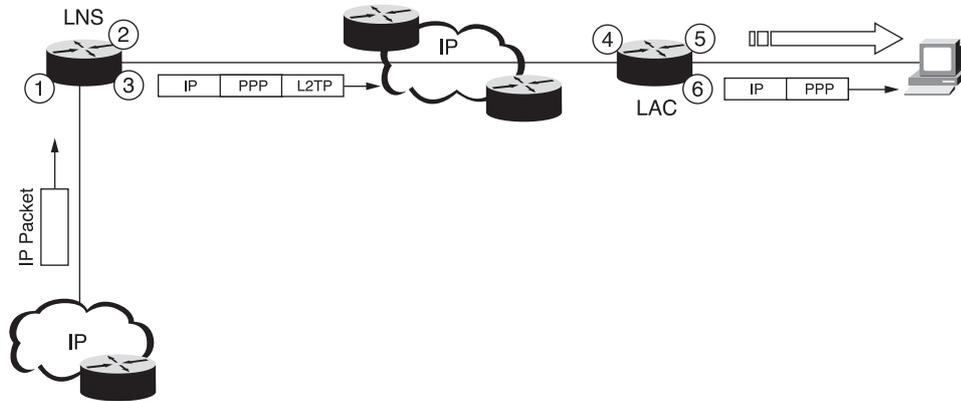
0790

Figure 3 Propagation of QoS Between IP and MPLS (790)

### 1.3.4

#### Propagation of QoS Between IP and L2TP

With L2TP packets by default, the DSCP and the precedence bits of the original IP packet are copied. The downstream process from the network to the SmartEdge router configured as an LNS to the SmartEdge router configured as an L2TP access concentrator (LAC) to the subscriber is illustrated in Figure 4



0792

Figure 4 Propagation of QoS Downstream from the Network (792)

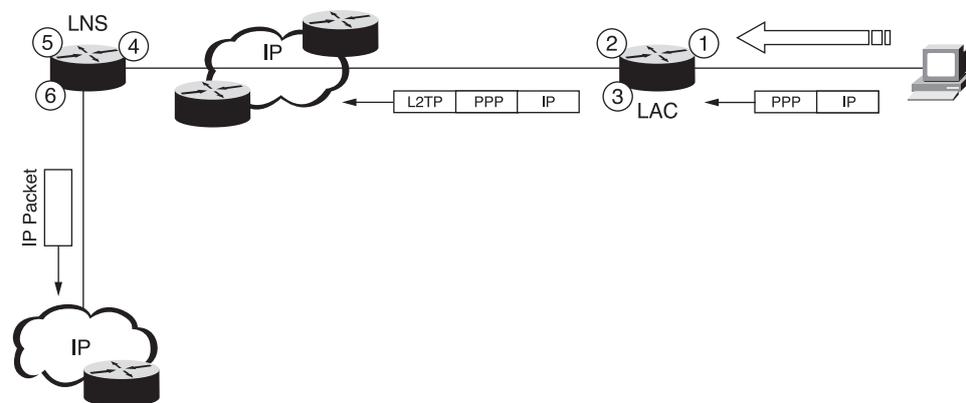
The downstream propagation process follows:

- 1 At the LNS, the SmartEdge router copies the DSCP bits from the inner subscriber IP packet header in the incoming IP packet to the PD QoS priority field.
- 2 The SmartEdge router then copies the PD QoS priority field to the DSCP bits in the outer L2TP IP packet header, using the `propagate qos to l2tp` command (in L2TP peer configuration mode), if configured. If the



- command is not configured, it sets the DSCP bits to the default (lowest) priority.
- 3 The SmartEdge router selects an egress queue for the L2TP packet, based on the PD QoS priority field.
  - 4 At the LAC, the SmartEdge router copies the DSCP bits in the outer L2TP IP packet header to the PD QoS priority field.
  - 5 The SmartEdge router then copies the DSCP bits from the inner subscriber IP packet header to the PD QoS priority field, using the `propagate qos from subscriber` command (in L2TP peer configuration mode) with the `downstream` keyword, if configured. This operation overwrites the PD QoS priority field set by step 4.
  - 6 The SmartEdge router selects an egress queue, based on the priority field in the PD.

The upstream process from the subscriber to the SmartEdge router configured as an LAC to the SmartEdge router configured as an LNS to the network is illustrated in Figure 5.



0793

Figure 5 Propagation of QoS Upstream from the Subscriber (793)

The upstream propagation process follows:

- 1 At the LAC, if the `propagate qos from subscriber` command (in L2TP peer configuration mode) with the `upstream` keyword is configured, the SmartEdge router copies the DSCP bits from the inner subscriber IP packet header in the incoming IP packet to the priority field in the PD. If the `propagate qos from subscriber` command is not configured, it sets the priority field to the default (lowest) priority.
- 2 The SmartEdge router then copies the priority field to the DSCP bits in the outer L2TP IP packet header, using the `propagate qos to l2tp` command (in L2TP peer configuration mode), if configured. If the command is not configured, it sets the DSCP bits to the default priority.



- 3 The SmartEdge router selects an egress queue for the L2TP packet based on the priority field.
- 4 At the LNS, the SmartEdge router copies the DSCP bits from the outer L2TP IP packet header in the incoming IP packet to the priority field in the PD.
- 5 The SmartEdge router then copies the priority field to the DSCP bits in the inner subscriber IP packet header, using the `propagate qos from l2tp` command (in L2TP peer configuration mode), if configured with no classification map specified. If this command is not used, the inner subscriber IP packet header is not altered.
- 6 The SmartEdge router selects an egress queue for the IP packet based on the priority field.

### 1.3.5 Priority Propagation for Oversubscribed Traffic Cards

Ingress propagation of QoS priority is available at the port level for certain oversubscribable traffic cards. Currently, only the 4-port 10 Gigabit Ethernet (10GE) card is supported. A traffic card is considered oversubscribable if it is designed such that the sum of the nominal rate of all its ports exceeds the actual packet-forwarding capacity of the card. The nominal packet-processing capacity of the 4-port 10GE traffic card is approximately 20 Gbps. When the total aggregate line bandwidth of in-service ports is greater than the nominal capacity of a traffic card, the card is oversubscribed. Any traffic received that exceeds the actual forwarding capacity of the card (which can vary based on factors such as average packet size and enabled ingress packet-processing features) is dropped.

The following port-level configuration commands are supported:

- `port-propagate qos from ethernet [class-map map-name]`
- `port-propagate qos from ip [class-map map-name]`
- `port-propagate qos from mpls [class-map map-name]`

The `port-propagate` commands enable the mapping of external priority markings of MPLS EXP, Ethernet 802.1p, and IP TOS/DSCP to SmartEdge router internal QoS priority values stored in the Packet Descriptor (PD QoS). The PD QoS value so assigned in turn determines each packet's ingress oversubscription treatment.

Each port has four ingress queues (numbered 0 to 3) for the incoming traffic, and the PD QoS priority value that is assigned to a packet determines to which of the four queues the packet is admitted. Each queue has a fixed mapping to a priority number or numbers as shown in Table 7.



*Table 7 Ingress Queue Assignment Based on PD Priority*

<b>Incoming Packets with PD QoS Priority Value</b>	<b>Are Placed in Ingress Queue</b>
0 or 1	0
2	1
3, 4, or 5	2
6 or 7	3

The PD QoS drop precedence value that is assigned to a packet determines which of the four ingress queue admittance priorities (numbered 0 to 3) is assigned to it as shown in Table 8.

*Table 8 Ingress Queue Admittance Priority Based on PD QoS Drop precedence Values*

<b>Incoming Packets with PD QoS Drop Precedence Value</b>	<b>Are Assigned an Ingress Queue Admittance Priority Value</b>
0 or 1	0
2 or 3	1
4 or 5	2
6 or 7	3

Ingress queue admittance priority determines how much of the overall capacity of a queue that an incoming packet is allowed to take advantage of. Packets of ingress queue admittance priority 0 (zero) are admitted to the queue up to the point at which it is completely full, but packets of other priorities are discarded rather than queued at successively lower thresholds of the total occupancy capacity of the queue.

The 4-port 10 GE traffic card enables different queuing thresholds for each PD priority value. Each of the four packet priorities can take advantage of a different portion of the total queue depth according to the following ratios. The following values are percentages of the maximum queue occupancy level:

- Ingress queue admittance priority 0 (highest)—100%
- Ingress queue admittance priority 1—90%
- Ingress queue admittance priority 2—80%
- Ingress queue admittance priority 3 (lowest)—70%

For example, if the target queue is 82% full, an incoming packet with a priority value of 2 or 3 is dropped, while a packet with a priority value of 0 or 1 is admitted to the queue.

The queues are serviced in strict priority order on a card-wide basis; that is, packets waiting in the highest-priority queues (queue 0) of any of the ports on



the card are serviced until those queues are empty, then packets waiting in the second-highest priority queue are serviced (queue 1), and so on. Under congestion, the highest-priority are the most likely to be forwarded, and the lowest-priority packets are the most likely to be discarded due to queue overflow. If the `port-propagate` commands are not configured, then all packets received on the port are treated as the lowest-priority traffic for ingress oversubscription purposes (they are assigned to ingress queue 3 with queue admittance priority 3).

Packets discarded because of ingress queue overflow are reported by the `show port counters detail` command in the "Receive Counters" section as "overflows" (total number of packets discarded) and "overflow bytes" (total number of octets discarded).

The PD QoS value assigned during ingress oversubscription processing always determines its ingress oversubscription treatment, and it can remain associated with the packet for the subsequent SmartEdge forwarding processing stages that it will be subject to (for example, to determine which transmit queue the packet is assigned to on the egress port). However, this PD QoS value can be overridden subsequent to ingress admission by any applicable `qos priority` or `propagate qos from` command that also applies to the circuit and the default DSCP-to-PD mapping applied to all Layer 3 (`bind interface`, `bind authentication`, `bind subscriber`) IP circuits. When more than one `port-propagate` command is enabled for the same port, the outermost priority field present in a packet for which a corresponding `port-propagate` command has been configured on the port takes precedence. For example, if both the `port-propagate from ip` and `port-propagate from ethernet` commands are configured for the same port, the outer 802.1p value determines the ingress priority for all the 802.1q-encapsulated traffic received by the port, and the IP DSCP value determines the priority for non-802.1q-encapsulated IP over Ethernet traffic.

The port-propagation ingress classifier identifies and decodes frames of different types based on the following Ethertype field values.

*Table 9 Supported Ethertype Field Values*

Ethertype Field Value	Indicates Which Protocol Is Encapsulated in the Frame
0x0800	IPv4
0x8100 <sup>(1)</sup>	802.1q
0x8847, 0x8848	MPLS

(1) Plus the Ethertype value specified by the `dot1q tunnel ethertype` command under the port, if configured.

**Note:** Frames received with EtherType field values other than ones listed in Table 10 have their PD QoS priority and drop precedence values set to 0 (zero) and receive the lowest-priority treatment for ingress oversubscription purposes (they are assigned to ingress queue 3 with queue admittance priority 3).



Table 10 lists the supported packets from which the supported PD priority values can be extracted.

*Table 10 Supported Packet Types for Ingress Port QoS Priority Propagation*

<b>IPv4 Header TOS Field DSCP Values Can Be Examined for the Following Packets</b>	<b>Ethernet 802.1p Priority Values Can Be Examined for the Following Packets</b>	<b>MPLS EXP Priority Values Can Be Examined for the Following Packets</b>
IPv4 over Ethernet	802.1q	MPLS over Ethernet
IPv4 over 802.1q	802.1q in 802.1q <sup>(1)</sup>	MPLS over 802.1q
IPv4 over 802.1q in 802.1q	—	MPLS over 802.1q in 802.1q

*(1) Only the outer 802.1p value is used in this case.*

## 1.4 Circuit Groups

Circuit groups allow you to group arbitrary PVCs or other circuits, such as subscriber sessions, for collective metering, policing, and scheduling. You can group PVCs—for example, to represent a business entity—and apply class-aware or circuit-level rate limits to the group. In this case, the traffic on all of the member circuits is collectively limited to any metering, policing, and scheduling rates configured on the circuit group.

Circuit group membership is available for 802.1Q PVCs, 802.1Q PVCs within an 802.1Q tunnel, or a mix of these circuit types. Circuit group membership is also available for subscriber circuits. When hierarchical rate limiting is applied to a circuit group, traffic is first limited according to any metering or policing policy applied to the member circuit. Subsequently, traffic is limited again according to any metering or policing policy applied to the circuit group.

Circuit group members can have child circuits configured under them. The following apply to circuit group membership:

- Circuit-group membership may be configured for 802.1Q tunnels with PVC children. Similarly, if an 802.1Q tunnel encapsulated PVC has been configured as belonging to a circuit group, child PVCs may be configured with that outer VLAN or (PVC) ID.
- Circuit-group membership may be configured for 802.1Q PVCs with authentication bindings and maximum sessions configured for a value greater than 1.
- Circuit-group membership may be configured for 802.1Q PVCs with the `circuit protocol` command configured.

When a circuit with children is assigned circuit-group membership, those child circuits remain associated with their parent and are subject to QoS inheritance of any QoS policies applied to the circuit group or the member circuit itself. However, if the child circuit is itself assigned circuit-group membership, the QoS



inheritance relationship with its parent circuit is broken and the child inherits from its own circuit group.

Counters values reported for the circuit group reflect the aggregate totals of all the members as well as their children.

Subscriber session circuits can be configured as direct members of a circuit group. You can configure the `circuit-group-member` command in the subscriber record, subscriber profile, or subscriber default configuration modes. In addition, a corresponding VSA (210, `Circuit_Group_Member`) is provided to allow configuration of circuit-group membership through RADIUS. The following apply to a subscriber session member of a circuit group:

- It is no longer subject to QoS policy inheritance from its parent port or PVC and is an equal peer of any PVC or subscriber member.
- If a subscriber session with a `Circuit-Group-Membership` attribute shares a circuit with its parent PVC due to `bind authentication maximum 1` configuration, the PVC parent of the subscriber session also becomes a member of the circuit group. When the subscriber session goes down, the PVC parent is no longer a circuit group member.

For information about the `circuit-group`, `circuit-group-member`, and `subscriber` commands, see *Command List*. For information about VSA 210, see *RADIUS Attributes*.

### 1.4.1 Virtual Port Circuit Groups

The SmartEdge router supports virtual port circuit group (VPCGs, which allow you to partition circuits or link groups of ports that support PWFQ and have a line rate greater than 1 Gbps for collective scheduling. You can use a VPCG to divide a high-speed port's traffic into multiple, lower-bandwidth scheduling domains. You can assign a circuit to a virtual port by making the circuit a member of the VPCG.

The SmartEdge router supports the following methods of VPCG assignment for virtual ports:

- Explicit assignment using standard manual provisioning with the `circuit-group-member` command or with VSA 210 (`Circuit_Group_Member`). For more information about the `circuit-group-member` command, see *circuit-group-member*. For information about VSA 210, see *RADIUS Attributes*.

**Note:** With explicit assignment of VPCG for a subscriber session, the subscriber session inherits the queuing policy and any metering and policing policies configured under the VPCG.

- Inheritance of VPCG membership. A circuit configured as an intermediate QoS scheduling node (L3 node) or with a PWFQ policy binding (L2 node) that does not have an explicit VPCG membership will inherit the



virtual-port assignment and be added to the virtual port's scheduling tree under the nearest circuit above it in the QoS inheritance hierarchy (see Circuit Configuration with QoS Policies) with an existing VPCG association (explicit, inherited, or auto-assigned). If there is no such circuit with VPCG membership in the hierarchy above, auto-assignment will be performed for the circuit (see next method).

- Auto-assignment based on TM configuration. If a circuit on a port or link group that requires virtual-port scheduling is provisioned for PWFQ (using one of the supported PWFQ configuration commands), and it does not have an explicit or inherited virtual-port assignment, the circuit is automatically assigned to one of the existing VPCGs configured for that port or link group.

Affected circuits are assigned to one of the available VPCGs based on a round robin selection of the currently available groups configured for the applicable port or access link group.

The following applies to auto-assignment of circuit group membership:

- Static circuits—Statically configured circuits, such as 802.1q PVCs, have their auto-assigned virtual-port membership reflected by the automatic addition of the `circuit-group-member` command to the router configuration.

Static circuit types:

- Implicit ranges—Each PVC in implicit PVC ranges (for example, 802.1q PVC 1 to 10) may independently be auto-assigned to a different VPCG.
- Explicit ranges—All PVCs in an explicit PVC range (for example, 802.1q PVC explicit 1 to 10) are collectively auto-assigned to a single VPCG.
- Transport ranges—All VLAN IDs in a transport PVC range (for example, 802.1q PVC transport 1 to 10) are collectively auto-assigned to a single VPCG.
- Dynamic circuits—Dynamically configured circuits, such as subscribers and on-demand PVCs, do not have their auto-assigned VPCG membership reflected in the router configuration.

Dynamic circuit types:

- On-demand ranges (also called circuit creation on demand (CCOD))—Each PVC in an on-demand PVC range (for example, 802.1q PVC on-demand 1 to 10) can be auto-assigned independently to a different VPCG when the PVC is activated.
- Subscriber sessions—Each applicable subscriber session can be auto-assigned independently to a different VPCG when the subscriber session is established or when a PWFQ policy binding is added to the subscriber (for example, through reauthentication, change of authorization (CoA), or the use of the `policy-refresh` command).



**Note:** Implicit, explicit, and on-demand range circuits create multiple circuits; transport range circuits do not.

The following applies to VPCGs:

- VPCGs are always active hierarchical nodes and are always an aggregation point for PWFQ traffic, even if no maximum rate is configured. Children always inherit the VPCG assignment of the parent.
- You can create a maximum of 10 VPCGs per port or link group. You must configure at least one VPCG per port or link group to enable PWFQ configuration on cards that require virtual ports.
- Assign circuits to VPCGs such that the maximum aggregate bandwidth of all the circuits of any VPCG is approximately 1 Gbps or less.
- A VPCG must be a top-level circuit group homed to a port or link group on an applicable card. A VPCG can only be configured on cards that support virtual ports (currently only the 10 Gigabit Ethernet 4-port card).

The following applies to a child of a VPCG:

- A circuit group created as a descendant of a VPCG implicitly inherits the home port or link group and virtual-port membership of its parent.
- A circuit group not specified as a VPCG or a descendant of VPCG, and homed to an affected port or link group, cannot accept PWFQ configuration and can only be provisioned for policing and metering.

On affected ports and link-groups, a circuit must be associated with a VPCG either directly or through inheritance before the circuit can accept any of the following PWFQ configuration commands:

- `qos policy queuing`
- `qos hierarchical mode strict`
- `qos rate`

A VPCG supports all of the functionality of a homed circuit group, including:

- Policing policy (inherited or hierarchical) binding
- Metering policy (inherited or hierarchical) binding
- PWFQ policy binding
- `qos rate maximum` command

The following configuration restrictions apply to VPCGs:

- The `qos rate minimum` and `qos weight` commands are not supported.



- Node groups are not supported. Node groups have been replaced by circuit-group membership for subscribers.
- Economical access link aggregation group (LAG) is not supported.
- L2TP Network Server (LNS) and Multilink PPP (MLPPP) subscribers support PWFQ on virtual-port cards. However, these subscribers do not support explicit or auto-assignment of VPCGs. Instead, the SmartEdge OS selects an appropriate VPCG based on the existing configuration.
- On-demand 802.1q PVCs (for example, CCoD circuits) support PWFQ queuing, but rely on automatic VPCG assignment. Explicit VPCG assignment (for example, the `circuit-group-member` command) is currently not supported for CCoD circuits.

### 1.4.2 Nested Circuit Groups

A circuit group can belong to another circuit group. Use these nested circuit groups to configure hierarchical QoS. Members of child circuit groups are subject to QoS inheritance from both the child circuit group and its parent circuit group. If a parent circuit group is homed to a port or link group, its children circuit groups are implicitly homed to the same port or link group.

No limit is enforced on the number of nested levels. However, two to four is the practical range (including VPCGs).

To configure a nested circuit group, use the `circuit-group` command with the `circuit-group-parent` keyword. For information about the `circuit-group` and `circuit-group-member` commands, see *Command List*.

### 1.4.3 Inheritance for Circuit Group Members

Natural children and grandchildren of circuit-group members are subject to inheritance of inheritable QoS policy bindings or other attributes from their natural circuit parent in the SmartEdge OS hierarchy including the following:

- Circuit-group binding > member PVC > subscriber session
- Circuit-group binding > member 802.1Q tunnel > PVC
- Circuit-group binding > member 802.1Q tunnel > PVC > subscriber session
- Member PVC binding > subscriber session
- Member 802.1Q tunnel binding > PVC
- Member 802.1Q tunnel > PVC > subscriber session

**Note:** A circuit-group member that has no explicit queuing policy applied inherits the queues and scheduling policy of its parent port or link group.



The following QoS policy bindings are inheritable:

- Queuing policy bindings.
- Overhead profile bindings.
- Metering and policing bindings configured with the `inherit` or `hierarchical` keyword.

L3 hierarchical node bindings are implicitly inherited by any L2 nodes below them in the hierarchy. Inherited bindings are overridden by any binding of the same type configured on the circuit itself or on an intermediate ancestor.

## 1.5 802.3ad LAG Support

The SmartEdge OS supports QoS for 802.3ad LAGs as follows:

- Ethernet LAGs support any QoS policy (policing, metering, queuing, forward) that is supported by the constituent ports. To provision and operate a QoS policy for an Ethernet LAG, configure the QoS service on each independent constituent port.
- Dot1q LAGs support any QoS policy (policing, metering, queuing, forward) that is supported by the constituent ports. To provision and operate a QoS policy for a dot1q LAG, configure the QoS service on each independent constituent port. All dot1q PVCs configured under the LAG automatically inherit the queuing policy. Use the `qos policy policing` command with the `inherit` keyword to apply metering and policing policies to the PVCs.

**Note:** You cannot specify individual or customized QoS for individual dot1q PVCs. QoS policies and rates that are applied to the ports also apply to all traffic under the aggregate port. You cannot apply a forward policy to the PVCs because forward policies do not support inheritance. You cannot configure 802.1p-based QoS propagation for the PVCs because the `propagate qos from ethernet` and `propagate qos to ethernet` commands are specified in dot1q profiles that cannot be applied to the PVCs configured under the dot1q LAG.

- Inheritable QoS attributes applied to constituent ports of a link group are inherited by and applied to the aggregate PVCs defined under the link group.
- Access LAGs support any QoS policing, metering, PWFQ, MDRR, or forward policy supported by the constituent ports. You can provision and operate a QoS policy for the LAG or its subordinate PVCs and subscribers. You can apply QoS policies to the LAG configuration object itself, the individual PVCs, or subscribers configured under the LAG.



**Note:** Regardless of the type of LAG and whether QoS services are configured on the individual ports or LAG, QoS is enforced on a per-link basis. For example, if a policy specifying a maximum rate of 100 kbps is configured on the port or LAG, each active port in the LAG is individually rate-limited to 100 kbps, for a possible total throughput of 100 kbps multiplied by the number of active ports.

See *802.3ad LAG Examples* for examples of how to configure QoS on 802.3ad LAGs.





## 2 Configuration and Operations Tasks

**Note:** In this section, the command syntax in the task tables displays only the root command; for the complete command syntax, see *Command List*.

To configure circuits for QoS features, perform the tasks described in the following sections:

### 2.1 Configuration Guidelines

This section includes configuration guidelines that affect more than one command or a combination of commands:

- If you attach a modified deficit round-robin (MDRR) policy to a PVC, you must also attach it to the port on which you have configured the PVC.
- The following guidelines apply to cross-connected circuits:
  - When you attach a QoS metering or policing policy to a cross-connected circuit, you can attach a policy to each individual circuit before or after you make the cross-connection.
  - You can attach a different metering or policing policy to each circuit.
  - You can attach both a metering and a policing policy to each circuit.
- The following guidelines apply to Ethernet and 802.1Q link groups:
  - You attach a policy to an Ethernet port rather than the link group of which it is a member; you attach the policy using one of the QoS policy commands (`qos policy metering`, `qos policy policing`, `qos policy queuing`) in port configuration mode.
  - You can attach any type of QoS policy that is supported by that type of Ethernet port. These include metering, policing, MDRR, and PWFQ policies. However, to preserve the operational characteristics of a link group, attach the same set of policies (metering, policing, and scheduling) to every constituent port in the link group.

### 2.2 Configure an ATM PVC for QoS

To configure an ATM PVC for QoS, perform the tasks described in the following sections:



## 2.2.1 Configure a PVC on a Second-Generation ATM OC Traffic Card

To configure an ATM PVC on a second-generation ATM OC traffic card, perform the tasks described in Table 11; enter all commands in ATM PVC configuration mode, unless otherwise noted.

Table 11 Configure a PVC on a Second-Generation ATM OC Traffic Card

Task	Root Command	Notes
For packets coming into the SmartEdge router, propagate the CLP bit to PD QoS priority values in ATM cells.	<i>clpbit propagate qos from atm</i>	Enter this command in ATM profile configuration mode.
For packets going out of the SmartEdge router, propagate PD QoS bits to the CLP bit in ATM cells.	<i>clpbit propagate qos to atm</i>	Enter this command in ATM profile configuration mode.
Attach a policing policy.	<i>qos policy policing</i>	
Attach a metering policy.	<i>qos policy metering</i>	
Optional. Specify the circuit rate, if different from the rate specified by the attached metering and policing policies.	<i>rate circuit</i>	You can specify rates for both inbound and outbound traffic.
Attach a scheduling policy to a PVC. <sup>(1)</sup>	<i>qos policy queuing</i>	Only ATMWFQ policies are supported; you can attach them only to PVCs.

(1) An ATMWFQ policy cannot be attached to a PVC that is shaped as UBRc.

## 2.3 Configure an Ethernet Circuit for QoS

To configure a circuit on any Ethernet traffic card for QoS, including any version of a Gigabit Ethernet traffic card, perform the tasks described in the following sections:

### 2.3.1 Configure Any Ethernet, Fast Ethernet-Gigabit Ethernet, or Gigabit Ethernet Circuit for QoS

To configure an Ethernet, Fast Ethernet-Gigabit Ethernet, or Gigabit Ethernet (any version) port, 802.1Q tunnel, or 802.1Q PVC, perform the tasks described in Table 12; enter all commands in port or dot1Q PVC configuration mode, unless otherwise noted.



**Table 12** *Configure Any Ethernet, Fast Ethernet-Gigabit Ethernet, or Gigabit Ethernet Circuit for QoS*

<b>Task</b>	<b>Root Command</b>	<b>Notes</b>
For packets coming into the SmartEdge router, translate Ethernet 802.1p priority bits to PD QoS and IP DSCP priority bits.	<i>propagate qos from ethernet</i>	Enter this command in dot1q profile configuration mode.  If you also configure a classification map under this command, the IP DSCP bits in the internal packet are not affected.
For packets going out of the SmartEdge router, propagate PD QoS priority bits to 802.1p priority bits.	<i>propagate qos to ethernet</i>	Enter this command in dot1q profile configuration mode.
Assign a PD QoS priority group to the port, tunnel, or PVC.	<i>qos priority</i>	The QoS bit setting for packets traveling across the ingress circuit is not changed by the PD QoS priority group assignment.  Not supported for CCOD ranges of 802.1q PVCs.
Attach an overhead profile to a port or an 802.1Q PVC.	<i>qos profile overhead</i>	
Attach a policing policy to the port, tunnel, or PVC.  Optional. Enable policy ACL rule counters.	<i>qos policy policing</i>	You can use the <b>acl-counters</b> keyword to enable IPv4 or IPv6 rule counters for the policy (may impact router performance; use with caution).
Set the rate for outgoing traffic for a Gigabit Ethernet port.	<i>qos rate</i>	
Attach a metering policy to a port, tunnel, or PVC.  Optional. Enable policy ACL rule counters.	<i>qos policy metering</i>	You can use the <b>acl-counters</b> keyword to enable IPv4 rule counters for the policy (may impact router performance; use with caution).
Attach a scheduling policy to a port, tunnel, or PVC.	<i>qos policy queuing</i>	Possible policy types are MDRR and PWFQ. <sup>(1)</sup>
Optional. Specify the circuit rate, if different from the rate specified by the attached metering and policing policies.	<i>rate circuit</i>	You can specify rates for both inbound and outbound traffic.  Not supported for CCOD ranges of 802.1q PVCs.

(1) MDRR policies are supported only on 10GE circuits.



### 2.3.2 Configure a Traffic-Managed Port for Hierarchical Scheduling

**Note:** PWFQ policies are supported only on traffic-managed ports (and the 802.1Q tunnels, 802.1Q PVCs, and hierarchical nodes configured on them). You can attach the same or different PWFQ policies to a port, its 802.1Q tunnels, its PVCs, and its hierarchical nodes.

To configure a traffic-managed port and any 802.1Q tunnels and PVCs configured on it for hierarchical scheduling with a PWFQ policy, perform the tasks described in Table 13; enter all commands in port configuration mode, unless otherwise noted. For information about the `dot1q pvc` command (in port configuration mode), see *Configuring Circuits*.

Table 13 Configure a Traffic-Managed Port for Hierarchical Scheduling

Task	Root Command	Notes
Set the maximum and minimum rates for the port.	<i>qos rate</i>	You must specify the maximum rate; the minimum rate is optional.
Specify the scheduling algorithm for the port.	<i>qos hierarchical mode strict</i>	
Attach a PWFQ policy to the port.	<i>qos policy queuing</i>	You can attach a policy to any or all 802.1Q tunnels and PVCs as well as the port.
Create one or more 802.1Q tunnels or PVCs and access dot1q PVC configuration mode.	<i>dot1q pvc</i>	
Set the maximum and minimum rates for the tunnel or PVC.	<i>qos rate</i>	Enter this command in dot1q PVC configuration mode. You must specify the maximum rate; the minimum rate is optional. You cannot set a minimum rate if you also assign a relative weight to this PVC.
Assign a relative weight to this PVC.	<i>qos weight</i>	Enter this command in dot1q PVC configuration mode. You cannot assign a relative weight if you also set a minimum rate for this PVC.
Specify the scheduling algorithm for the tunnel or PVC.	<i>qos hierarchical mode strict</i>	Enter this command in dot1q PVC configuration mode.
Attach a PWFQ policy to the tunnel or PVC.	<i>qos policy queuing</i>	Enter this command in dot1q PVC configuration mode. You can attach a policy to any or all tunnels and PVCs, as well as the port.



### 2.3.3 Configure a Traffic-Managed Port for Hierarchical Nodes

To configure a traffic-managed port for hierarchical nodes, node groups, and attach PWFQ policies to them, perform the tasks described in Table 14; enter all commands in port configuration mode, unless otherwise noted.

Table 14 Configure a Traffic-Managed Port for Hierarchical Nodes

Task	Root Command	Notes
Set the maximum and minimum rates for the port.	<i>qos rate</i>	You must specify the maximum rate; the minimum rate is optional.
Specify the scheduling algorithm for the port.	<i>qos hierarchical mode strict</i>	
Create one or more hierarchical node groups and access hierarchical node group configuration mode.	<i>qos node-group</i>	
Set the maximum and minimum rates for the node groups.	<i>qos rate</i>	Enter this command in hierarchical node group configuration mode. You must specify the maximum rate; the minimum rate is optional. You cannot set a minimum rate if you also assign a relative weight to this node group.
Assign a relative weight to this node group.	<i>qos weight</i>	Enter this command in hierarchical node group configuration mode. You cannot assign a relative weight if you also set a minimum rate for this node group.
Specify the scheduling algorithm for the node groups.	<i>qos hierarchical mode strict</i>	Enter this command in hierarchical node group configuration mode. The mode need not be the same as the one you specify for the port.
Create one or more hierarchical nodes and access hierarchical node configuration mode.	<i>qos node</i>	Enter this command in hierarchical node group configuration mode.
Set the maximum and minimum rates for these nodes.	<i>qos rate</i>	Enter this command in hierarchical node configuration mode. You must specify the maximum rate; the minimum rate is optional. You cannot set a minimum rate if you also assign a relative weight to this node.



Table 14 Configure a Traffic-Managed Port for Hierarchical Nodes

Task	Root Command	Notes
Assign a relative weight for these nodes.	<i>qos weight</i>	Enter this command in hierarchical node configuration mode. You cannot assign a relative weight if you also set a minimum rate for this node.
Specify the scheduling algorithm for these nodes.	<i>qos hierarchical mode strict</i>	Enter this command in hierarchical node configuration mode. The mode need not be the same as the one you specify for the port or node group.
Attach a PWFQ policy to these nodes.	<i>qos policy queuing</i>	Enter this command in hierarchical node configuration mode. The policy need not be the same as the one you attach to the port, tunnel, or PVC.

## 2.4 Configure a POS Circuit for QoS

To configure a circuit on a Packet over SONET/SDH (POS) traffic card for QoS, perform the tasks described in Table 15; enter all commands in port configuration mode.

Table 15 Configure a POS Circuit for QoS

Task	Root Command	Notes
Assign a PD QoS priority group.	<i>qos priority</i>	The QoS bit setting for packets traveling across the ingress circuit is not changed by the PD QoS priority group assignment.
Attach a policing policy.	<i>qos policy policing</i>	
Attach a metering policy.	<i>qos policy metering</i>	
Optional. Specify the circuit rate, if different from the rate specified by the attached metering and policing policies.	<i>rate circuit</i>	You can specify rates for both inbound and outbound traffic.
Attach a scheduling policy.	<i>qos policy queuing</i>	



## 2.5 Configure Cross-Connected Circuits for QoS

To configure a cross-connected circuit for QoS, perform the tasks described in Table 16 in. You cannot attach a scheduling policy to a cross-connected circuit; only metering and policing policies are supported on either or both circuits.

**Note:** You can perform the tasks in Table 16 in any order.

Table 16 Configure a Cross-Connected Circuit for QoS

Step	Task	Root Command	Notes
1.	Configure the inbound circuit for QoS with one of the following tasks:		
	An inbound ATM PVC.		Perform the tasks in Table 11, but do not attach a scheduling policy.
	An inbound 802.1Q PVC.		Perform the tasks in Table 13, but do not attach a scheduling policy.
2.	Configure the outbound circuit for QoS with one of the following tasks:		
	An outbound ATM PVC.		Perform the tasks in Table 11, but do not attach a scheduling policy.
	An outbound 802.1Q PVC.	<i>qos priority</i>	Perform the tasks in Table 13, but do not attach a scheduling policy.
3.	Create the cross-connection between the inbound and outbound circuits.	<i>xc</i>	Enter this command in global configuration mode. For information about this command, see the <i>Configuring Cross-Connections</i> .

## 2.6 Configure a Subscriber Circuit for QoS

You configure a subscriber circuit (or an LNS subscriber session) for QoS by configuring the subscriber record or profile; to configure a subscriber record or profile and thus any circuit on which the subscriber session is created, perform one or more of the tasks described in Table 17; enter all commands in subscriber configuration mode unless otherwise noted.



Table 17 Configure a Subscriber Circuit for QoS

Task	Root Command	Notes
Create a reference to a hierarchical node.	<i>qos node-reference</i>	
Attach a policing policy.	<i>qos policy policing</i>	
Attach a metering policy.	<i>qos policy metering</i>	
Optional. Specify the circuit rate, if different from the rate specified by the attached metering and policing policies.	<i>rate circuit</i>	You can specify rates for both inbound and outbound traffic.
Attach a scheduling policy.	<i>qos policy queuing</i>	Policy types include ATMWFQ, MDRR, and PWFQ. Only PWFQ policies are supported for LNS subscriber sessions.
Attach an overhead profile to a subscriber record.	<i>qos profile overhead</i>	Enter this command in port configuration mode.

## 2.7 Configure QoS Propagation (Optional)

To create and apply customized classification mappings for QoS bits, perform the tasks described in Table 18; enter all commands in class map configuration mode, unless otherwise noted.

Table 18 Configure QoS Propagation

Task	Root Command	Notes
Create a classification map.	<i>qos hierarchical mode strict</i>	Enter this command in global configuration mode.
Specify a set of default values for a classification map.	<i>mapping-schema</i>	
Specify an initial PD value to assign to ATM packets with the specified CLP value.	<i>atm to qos</i>	
Translate outgoing PD QoS priority values to ATM CLP values.	<i>qos to atm</i>	
For incoming packets with the specified CLP value, determine the initial PD QoS priority value from the user priority bits in the 802.1p VLAN TCI field of the packet header value.	<i>atm use-ethernet</i>	



Table 18 Configure QoS Propagation

Task	Root Command	Notes
For incoming packets with the specified CLP value, determine the initial PD QoS priority value from the IP ToS field of the packet header.	<i>atm use-ip</i>	
Translate incoming Ethernet 802.1p values to PD QoS priority values.	<i>ethernet to qos</i>	
Translate outgoing PD QoS priority values to Ethernet 802.1p values.	<i>qos to ethernet</i>	
For 802.1Q transport PVCs, use the 802.1p value from either the outer PVC header or the inner PVC header to propagate PD QoS priority values to 802.1p values in outgoing packets.	<i>propagate qos transport use-vlan-header</i>	Enter this command in dot1q profile configuration mode.
For incoming packets with the specified 802.1p value, determine the initial PD QoS value from the DSCP value in the IP header rather than the 802.1p value.	<i>ethernet use-ip</i>	
Translate incoming IP header DSCP values to PD QoS values.	<i>ip to qos</i>	
Translate outgoing PD QoS values to IP header DSCP values.	<i>qos to ip</i>	
For outgoing packets with the specified PD QoS value, determine the final header EXP or 802.1p value based on the IP header DSCP value rather than the PD QoS value.	<i>qos use-ip</i>	
Translate incoming MPLS EXP values to PD QoS values.	<i>mpls to qos</i>	
Translate outgoing PD QoS values to MPLS header EXP values.	<i>qos to mpls</i>	
For incoming packets with the specified MPLS EXP priority label, use the encapsulated Ethernet packet's 802.1p priority label to determine the PD QoS priority value for the packet.	<i>mpls use-ethernet</i>	



Table 18 Configure QoS Propagation

Task	Root Command	Notes
For incoming packets with the specified EXP value, determine the initial PD QoS priority value from the IP header DSCP value rather than the EXP value.	<i>mpls use-ip</i>	
For incoming MPLS EXP packets with an Ethernet VLAN header, specify the 802.1q-over-MPLS packets to examine for any header enclosed by the outer VLAN header.	<i>propagate qos use-vlan-ethertype</i>	Enter this command in MPLS router configuration mode.
For incoming MPLS packets, use the 802.1p value from the header of either the outer PVC header or the inner PVC for propagation from Ethernet to internal PD QoS classification values.	<i>propagate qos use-vlan-header</i>	Enter this command in MPLS router configuration mode.
Reference the classification map when configuring propagation.	<i>propagate qos from ethernet</i> <i>propagate qos from ip</i> <i>propagate qos from l2tp</i> <i>propagate qos from mpls</i> <i>propagate qos to ethernet</i> <i>propagate qos to ip</i> <i>propagate qos to l2tp</i> <i>propagate qos to mpls</i>	Specify the <b>class-map class-map</b> construct for this function.

## 2.8 Configure L2TP for QoS

To configure L2TP for QoS to propagate subscriber DSCP bits in the downstream direction, perform the tasks described in Table 19; enter all commands in L2TP peer configuration mode for the default peer.



Table 19 Configure L2TP for QoS in the Downstream Direction

Task	Root Command	Notes
For network packets the SmartEdge router sends to the LAC when the router is configured as an LNS, propagate the PD QoS priority bits to the outer DSCP value.	<i>propagate qos to l2tp</i>	
For L2TP IP packets coming into the SmartEdge router when it is configured as a LAC, propagate the subscriber DSCP bits from the inner IP packet header to the PD QoS priority bits from the LNS for the subscriber IP packet.	<i>propagate qos from subscriber</i>	Specify the <b>downstream</b> keyword for this function.

To configure L2TP for QoS to propagate DSCP bits in the upstream direction, perform the tasks described in Table 20; enter all commands in L2TP peer configuration mode for the default peer.

Table 20 Configure L2TP for QoS in the Upstream Direction

Task	Root Command	Notes
For subscriber IP packets coming into the SmartEdge router when it is configured as a LAC, propagate the subscriber DSCP bits in the IP packet header to the PD QoS priority bits for the subscriber IP packet.	<i>propagate qos from subscriber</i>	Specify the <b>upstream</b> keyword for this function.
For network packets the SmartEdge router sends to the LNS when the router is configured as a LAC, propagate the PD QoS priority bits to the outer DSCP value.	<i>propagate qos to l2tp</i>	



Table 20 Configure L2TP for QoS in the Upstream Direction

Task	Root Command	Notes
For L2TP packets coming into the SmartEdge router when it is configured as an LNS, after the outer DSCP bits have been propagated to the PD QoS priority bits, propagate the PD bits to the subscriber's inner DSCP bits.	<i>propagate qos from l2tp</i>	
For L2TP packets coming into the SmartEdge router when it is configured as an LNS, customize the default mapping from the outer IP header DSCP value to the PD QoS priority value, leaving the inner IP header DSCP value unmodified.	<i>propagate qos from l2tp</i>	Specify the optional class-map parameter.

## 2.9 Configure MPLS for QoS

To configure MPLS for QoS, perform the tasks described in one of the following sections:

- Propagate QoS Using DSCP Bits and MPLS EXP Bits
- Propagate QoS Using DSCP Bits Only

### 2.9.1 Propagate QoS Using DSCP Bits and MPLS EXP Bits

To propagate QoS using DSCP bits to MPLS EXP bits (instead of DSCP bits) and vice versa, perform the tasks described in Table 21; enter either or both commands in MPLS router configuration mode.

Table 21 Propagate QoS Using DSCP Bits and MPLS EXP Bits

Task	Root Command	Notes
For packets coming into the SmartEdge router, propagate MPLS EXP bits to PD QoS priority and IP DSCP bits.	<i>propagate qos from mpls</i>	If you also configure a classification map, a custom value mapping is used to map MPLS EXP bits to PD QoS priority values. The DSCP bits are unaffected.
For packets going out of the SmartEdge router, propagate PD QoS priority values to MPLS EXP bits.	<i>propagate qos to mpls</i>	



## 2.9.2 Propagate QoS Using DSCP Bits Only

To propagate QoS by enabling the use of DSCP bits (instead of MPLS EXP bits) only, perform the task described in Table 22.

Table 22 Propagate QoS Using DSCP Bits Only

Task	Root Command	Notes
Enable the use of IP header DSCP bits (not MPLS EXP values) when determining the initial PD priority value of incoming MPLS packets.	<i>egress prefer dscp-qos</i>	Enter this command in MPLS router configuration mode.

## 2.10 Attach QoS Policies to a Circuit Group and Assign Members to the Group

To create a circuit group, attach a QoS metering, policing or scheduling policy to it, and then assign members to the group, perform the tasks described in Table 23.

Table 23 Attach QoS Policies to a Circuit Group and Assign Members to the Group

Step	Task	Root Command	Notes
1.	Create a circuit group and assign a specified name to it.	<i>circuit group</i>	Enter this command in global configuration mode. For information about this command, see <i>Configuring Circuits</i> .
2.	Attach a (QoS) metering, policing, or scheduling policy to the circuit group.		
	Attach a metering policy.	<i>qos policy metering</i>	Enter this command in dot1q PVC configuration mode.
	Attach a policing policy.	<i>qos policy policing</i>	Enter this command in dot1q PVC configuration mode.
	Attach a queuing policy.	<i>qos policy queuing</i>	Enter this command in dot1q PVC configuration mode.
3.	Select an Ethernet port in which the members of the circuit group are to reside and access port configuration mode.	<i>port ethernet</i>	Enter this command in global configuration mode.



Table 23 Attach QoS Policies to a Circuit Group and Assign Members to the Group

Step	Task	Root Command	Notes
4.	Specify the use of 802.1Q encapsulation for the Ethernet port.	<i>encapsulation dot1q</i>	Enter this command in port configuration mode.
5.	Specify the 802.1Q tunnel or one or more static 802.1Q PVCs to be assigned to the specified circuit group and access dot1q PVC configuration mode.	<i>dot1q pvc</i>	Enter this command in port configuration mode.
6.	Specify that the 802.1Q tunnel or PVCs being configured are members of the specified circuit group.	<i>circuit-group-member</i>	Enter this command in dot1q PVC configuration mode. For information about this command, see <i>Configuring Circuits</i> .

## 2.11 Configuring Virtual Port Circuit Groups

For information on how to configure VPCGs, see *Configuring VPCGs*.

## 2.12 Configure Nested Circuit Groups

To configure nested circuit groups, perform the tasks described in Table 24; enter the commands in the specified configuration modes. It is assumed you have configured the QoS policies before performing the tasks in Table 24.

**Note:** No limit on the number of nested levels is enforced. However, two to four is the practical range.

Table 24 Configure Nested Circuit Groups

Step	Task	Root Command	Notes
1.	Define a circuit group and specify a port on which all circuits in this circuit group reside. This is the parent circuit group.	<i>circuit group</i>	Enter this command in global configuration mode.



Table 24 Configure Nested Circuit Groups

Step	Task	Root Command	Notes
2.	Optional. Attach QoS policies, QoS overhead profiles, rate, or weight to the circuit group and its members.	<ul style="list-style-type: none"> <li>• <i>qos hierarchical mode</i></li> <li>• <i>qos policy metering</i></li> <li>• <i>qos policy policing</i></li> <li>• <i>qos policy queuing</i> (for priority weighted fair queuing [PWFQ] policies only)</li> <li>• <i>qos profile overhead</i></li> <li>• <i>qos rate</i></li> <li>• <i>qos weight</i></li> </ul>	Enter commands in circuit-group configuration mode.
3.	Define a circuit group, specify a parent circuit group, and specify a port on which all circuits in this circuit group reside.	<i>circuit group</i>	Enter this command in global configuration mode. Use the <b>parent-circuit-group</b> keyword to specify the parent circuit group from which all circuits in this group are subject to inheritance from.

## 2.13 Configure Ingress Port Propagation

To configure ingress port propagation for cards that support differentiated packet treatment for ingress oversubscription scenarios, perform the tasks described in Table 25.

Table 25 Configure Ingress Port Propagation

Step	Task	Root Command	Notes
1.	Configure a card and access card configuration mode.	<i>card</i>	Enter this command in global configuration mode.
2.	Select an Ethernet port to apply ingress port propagation and access port configuration mode.	<i>port ethernet</i>	Enter this command in global configuration mode.
3.	For incoming packets, configure ingress port propagation. Perform the following steps:		Enter the <b>port-propagate</b> commands in port configuration mode.



Table 25 Configure Ingress Port Propagation

Step	Task	Root Command	Notes
	Propagate Ethernet 802.1p user priority bits to PD QoS priority bits and use the 802.1p bits to determine ingress oversubscription treatment for incoming packets.	<i>port-propagate qos from ethernet</i>	
	Propagate IP DSCP bits to the PD QoS priority bits and use DSCP bits to determine ingress oversubscription treatment for incoming packets.	<i>port-propagate qos from ip</i>	
	Propagate MPLS EXP bits to PD QoS bits and use EXP bits to determine ingress oversubscription treatment for incoming packets.	<i>port-propagate qos from mpls</i>	

## 2.14 Operations Tasks

To monitor and administer QoS features, perform the appropriate tasks described in Table 26. Enter the `debug` command in exec mode; enter the `show` commands in any mode.

You can enable policy ACL rule counters using the `acl-counters` keyword with the `qos policy policing`, `qos policy metering`, and `qos policy queueing` commands.

---



---

### Warning!

Enabling IP filtering and policy ACL counters can impact router performance. Use them carefully.

---



---

Table 26 Monitor and Administer QoS Features

Task	Command	
Enable the generation of QoS debug messages.	<i>debug qos</i>	



Table 26 Monitor and Administer QoS Features

Task	Command	
Display information about ACLs attached to QoS policies.	<i>show access-group</i>	If you have enabled counters for a subscriber, use the <b>show access group ipv6 qos cct-handle in count</b> construct to view traffic flow. Enter the command twice, two to three minutes apart to view incremented counters.  The cct-handle (in the format <i>slot/port:ch:sub:[subsub]</i> ) is the subscriber circuit where the counters where enabled.
Display cell and SAR packet-level counters for configured ATM PVCs.	<i>show atm counters</i>	
Display general counters and counters specific to the circuit type for one or more circuits in the system.	<i>show circuit counters</i>	
Display the current QoS configuration.	<i>show configuration qos</i>	
Display the information for a particular circuit that matches the specified circuit agent ID.	<i>show qos agent-circuit-id</i>	
Display the information for all circuits or a particular circuit in the system.	<i>show qos circuit</i>	
Display QoS PPA client information.	<i>show qos client</i>	
Display QoS hierarchical node information.	<i>show qos h-node</i>	
Display QoS daemon memory usage information.	<i>show qos memory</i>	
Display information for all QoS policies.	<i>show qos policy</i>	
Display the circuit information for a specific subscriber on the system.	<i>show qos username</i>	
Display information about one or more QoS metering policies.	<i>show qos policy metering</i>	



Table 26 Monitor and Administer QoS Features

<b>Task</b>	<b>Command</b>	
Display information about one or more QoS policing policies.	<i>show qos policy policing</i>	
Display active QoS port bindings for ports on one or more traffic cards.	<i>show qos port-map bind</i>	



## 3 Configuration Examples

This section provides examples of QoS configurations for attaching rate- and class-limiting policies, attaching scheduling policies, and propagating QoS.

### 3.1 Attaching Rate- and Class-Limiting Policies

Examples of configuring PVCs and subscriber records for QoS policies are provided in the following sections:

#### 3.1.1 PVC Configuration

The following example attaches a metering policy, **meter**, to an 802.1Q PVC on an Ethernet port:

```
[local]Redback(config)#port ethernet 4/2
[local]Redback(config-port)#encapsulation dot1q
[local]Redback(config-port)#dot1q pvc 200
[local]Redback(config-dot1q-pvc)#bind interface if-200 local
[local]Redback(config-dot1q-pvc)#qos policy metering meter
```

#### 3.1.2 Cross-Connected Circuit Configuration

The following example attaches a metering policy, **output**, to the inbound circuits of cross-connected 802.1Q PVCs on Ethernet ports:



```
[local] Redback(config)#port ethernet 4/1
[local] Redback(config-port)#encapsulation dot1q
[local] Redback(config-port)#dot1q pvc 2001
[local] Redback(config-dot1q-pvc)#qos policy metering output
[local] Redback(config-dot1q-pvc)#exit
[local] Redback(config-port)#dot1q pvc 2051
[local] Redback(config-dot1q-pvc)#qos policy metering output
[local] Redback(config-dot1q-pvc)#exit
[local] Redback(config-port)#dot1q pvc 2101
[local] Redback(config-dot1q-pvc)#qos policy metering output
[local] Redback(config-dot1q-pvc)#exit
!
[local] Redback(config)#port ethernet 4/1
[local] Redback(config-port)#encapsulation dot1q
[local] Redback(config-port)#dot1q pvc 2001
[local] Redback(config-dot1q-pvc)#exit
[local] Redback(config-port)#dot1q pvc 2051
[local] Redback(config-dot1q-pvc)#exit
[local] Redback(config-port)#dot1q pvc 2101
!
[local] Redback(config)#xc 4/1 vlan-id 2001 to 4/3 vlan-id 2001
[local] Redback(config)#xc 4/1 vlan-id 2051 to 4/3 vlan-id 2051
[local] Redback(config)#xc 4/1 vlan-id 2101 to 4/3 vlan-id 2101
```

### 3.1.3 Subscriber Configuration

The following example attaches a metering policy, **meter**, to a subscriber record:

```
[local] Redback(config)#subscriber name redback
[local] Redback(config-sub)#password redback
[local] Redback(config-sub)#qos policy metering meter
```

## 3.2 Attaching Scheduling Policies

Examples of configuring ports and PVCs for QoS features using scheduling policies are provided in the following sections:

### 3.2.1 Port Configuration

The following example attaches a queuing policy to a POS port:

```
[local] Redback(config)#port pos 2/1
[local] Redback(config-port)#qos policy queuing pos-qos
```



### 3.2.2 PVC Configuration

The following example attaches a scheduling policy to each of three 802.1Q PVCs:

```
[local]Redback(config)#port ethernet 4/1
[local]Redback(config-port)#encapsulation dot1q
[local]Redback(config-port)#dot1q pvc 100
[local]Redback(config-dot1q-pvc)#bind interface if-100 local
[local]Redback(config-dot1q-pvc)#qos policy queuing PerVcQueuing
[local]Redback(config-dot1q-pvc)#dot1q pvc 101
[local]Redback(config-dot1q-pvc)#bind interface if-101 local
[local]Redback(config-dot1q-pvc)#qos policy queuing PerVcQueuing
[local]Redback(config-dot1q-pvc)#dot1q pvc 102
[local]Redback(config-dot1q-pvc)#bind interface if-102 local
[local]Redback(config-dot1q-pvc)#qos policy queuing PerVcQueuing
```

### 3.2.3 Overhead Profile Configuration

The following example allows the child circuits of 802.1Q PVC **100** to inherit the **example1** overhead profile:

```
[local]Redback(config)#port ethernet 2/1
[local]Redback(config-port)#dot1q pvc 100 encapsulation lqtunnel
[local]Redback(config-port)#qos profile overhead example1 inherit
```

### 3.2.4 PWFQ Policy and Hierarchical Shaping

The following example configures a GE3 port with the **home** node group with **5 dslam** nodes and attaches a PWFQ policy to each node:

```
[local]Redback(config)#port ethernet 5/2
[local]Redback(config-port)#qos rate maximum 100000000
[local]Redback(config-port)#qos rate minimum 100000
[local]Redback(config-port)#qos hierarchical mode strict
[local]Redback(config-port)#qos node-group home 1
[local]Redback(config-h-node)#qos node dslam 1 through 5
[local]Redback(config-h-node)#qos policy queuing pwfq4
```

### 3.2.5 PWFQ Policy and Hierarchical Scheduling

The following example configures a GE3 port and its 802.1Q PVC for hierarchical scheduling and attaches a PWFQ policy to both the port (**pwfq-port**) and its PVC (**pwfq-pvc**):

```
[local]Redback(config)#port ethernet 5/1
[local]Redback(config-port)#encapsulation dot1q
[local]Redback(config-port)#qos rate maximum 100000000
[local]Redback(config-port)#qos rate minimum 100000
[local]Redback(config-port)#qos hierarchical mode strict
[local]Redback(config-port)#qos policy queuing pwfq-port
[local]Redback(config-port)#dot1q pvc 200
[local]Redback(config-dot1q-pvc)#qos rate maximum 100000000
[local]Redback(config-dot1q-pvc)#qos rate minimum 100000
[local]Redback(config-dot1q-pvc)#qos policy queuing pwfq-pvc
```



### 3.3 Propagating QoS

The following example configures 802.1q profile, **8021q-on**, to propagate QoS information from incoming packets to outgoing packets in any 802.1Q tunnel or PVC that has that profile assigned to it:

```
[local]Redback(config)#dot1q profile 8021p-on
[local]Redback(config-dot1q-profile)#propagate qos from ethernet
[local]Redback(config-dot1q-profile)#propagate qos to ethernet
[local]Redback(config-dot1q-profile)#exit
```

The following example propagates QoS priority markings on an 802.1Q PVC by configuring it with the **8021p-on** profile:

```
[local]Redback(config)#port ethernet 3/1
[local]Redback(config-port)#encapsulation dot1q
[local]Redback(config-port)#dot1q pvc 20 profile 8021p-on
[local]Redback(config-dot1q-pvc)#exit
```

The following example maps control protocols, including VRRP advertisements, to an 802.1p priority value of **3**. Control packets are always assigned an internal ToS priority of 6 (that is, CS6), so this example maps priority **CS6** to a p bit of **3**:

```
[local]Redback(config)#qos class-map pd-to-8021p ethernet out
[local]Redback(config-class-map)#qos cs6 to ethernet 3
[local]Redback(config-class-map)#exit
[local]Redback(config)#dot1q profile 8021p-on
[local]Redback(config-dot1q-profile)#propagate qos to ethernet class-map pd-to-8021p
[local]Redback(config-dot1q-profile)#exit
[local]Redback(config)#port ethernet 1/1
[local]Redback(config-port)#no shutdown
[local]Redback(config-port)#encapsulation dot1q
[local]Redback(config-port)#dot1q pvc 100 profile 8021p-on
[local]Redback(config-dot1q-pvc)#bind interface test local
[local]Redback(config-dot1q-pvc)#exit
```

The following example enables IP QoS information to be propagated to ATM on any ATM PVC or virtual path (VP) that has the profile, **clp-on**, assigned to it:

```
[local]Redback(config)#atm profile clp-on
[local]Redback(config-atm-profile)#clpbit propagate qos to atm
[local]Redback(config-atm-profile)#exit
```

The following example configures MPLS to propagate QoS priority markings in both directions:

```
[local]Redback(config)#context local
[local]Redback(config-ctx)#router mpls
[local]Redback(config-mpls)#propagate qos from mpls
[local]Redback(config-mpls)#propagate qos to mpls
[local]Redback(config-mpls)#exit
```

The following example creates a classification map **exp-to-pd** that maps MPLS experimental EXP values to QoS PD values on ingress, then applies the classification map to the **propagate qos from mpls** command:



```
[local]Redback(config)#qos class-map exp-to-pd mpls in
[local]Redback(config-class-map)#mapping-schema 7P1D
[local]Redback(config-class-map)#mpls 1 to qos 0x24
[local]Redback(config-class-map)#exit
[local]Redback(config)#context mycontext
[local]Redback(config-ctx)#router mpls
[local]Redback(config-mpls)#propagate qos from mpls class-map exp-to-dscp
[local]Redback(config-mpls)#exit
```

The following example shows how to propagate Ethernet 802.1p user priority bits, IP DSCP bits, and MPLS EXP bits to PD QoS priority bits for incoming packets arriving in slot 10, port 1 of an Ethernet port and to use these bits to determine ingress oversubscription treatment. This configuration is applied to a 4-port 10GE card that is in slot 10 of the chassis:

```
[local]Redback(config)#card 10ge-4-port 10
[local]Redback(config)#port ethernet 10/1
[local]Redback(config-port)#port-propagate qos from ip
[local]Redback(config-port)#port-propagate qos from ethernet
[local]Redback(config-port)#port-propagate qos from mpls
```

### 3.4 Attaching QoS Policies to Circuit Groups

The following example shows how to create a circuit group `salesgroup` and attach previously configured policing and metering policies (`group_policing_policy` and `group_metering_policy`) to this circuit group. This example also shows how to assign the 802.1Q PVC tunnels (50 through 60, 30, and 40) as members of the circuit group. The individual PVCs (or VLANs), 1 to 100, configured under the 802.1Q tunnel 30 `cvlan_individual_policy`. With the `hierarchical` keyword configured for the policing policy (each have their own individual policing policy specified as `group_metering_policy`), the traffic on the individual PVCs are subject to both the child circuit policy (`cvlan_individual_policy`) and the parent circuit policy (`group_metering_policy`):

```
[local]Redback(config)#circuit-group salesgroup
[local]Redback(config-circuit-group)#qos policy policing group_policing_policy hierarchical
[local]Redback(config-circuit-group)#qos policy metering group_metering_policy inherit
[local]Redback(config)#port ethernet 12/1
[local]Redback(config-port)#encapsulation dot1q
[local]Redback(config-port)#dot1q pvc 50 through 60
[local]Redback(config-dot1q-pvc)#circuit-group-member salesgroup
[local]Redback(config)#port ethernet 12/2
[local]Redback(config-port)#encapsulation dot1q
[local]Redback(config-port)#dot1q pvc 30 encapsulation lqtunnel
[local]Redback(config-dot1q-pvc)#dot1q pvc 30:1 through 100
[local]Redback(config-dot1q-pvc)#circuit-group-member salesgroup
! Specify that each 802.1Q PVC configured under the 802.1Q tunnel also
! has its own individual policing policy (specified as
! cvlan_individual_policy).
[local]Redback(config-dot1q-pvc)#qos policy policing cvlan_individual_policy
[local]Redback(config-dot1q-pvc)#dot1q pvc 40
[local]Redback(config-dot1q-pvc)#circuit-group-member salesgroup
```



## 3.5 QoS on 802.3ad LAGs

This section provides examples for configuring QoS on Ethernet, dot1q, and access LAGs.

### 3.5.1 Example: Provision QoS on an Ethernet LAG

The following example shows how to provision QoS policies on an Ethernet LAG. Note that for an Ethernet LAG, QoS service is configured on each independent constituent port within the LAG (in this example, there are two ports in the LAG called `ETHER_EXAMPLE`).

```
! Create the LAG and bind it to an interface:
[local]rock1200#configure
[local]rock1200(config)#link-group ETHER_EXAMPLE ether
[local]rock1200(config-link-group)#bind interface LINKG_GROUP local
[local]rock1200(config-link-group)#end

! Add the first port to the LAG and configure QoS policies on that port:
[local]rock1200#configure
[local]rock1200(config)#port ethernet 1/1
[local]rock1200(config-port)#link-group ETHER_EXAMPLE
[local]rock1200(config-port)#qos policy metering meter_policy
[local]rock1200(config-port)#qos policy policing policing_policy
[local]rock1200(config-port)#qos policy queuing queuing_policy
[local]rock1200(config-port)#end

! Add the second port to the LAG and configure QoS policies on that port:
[local]rock1200#configure
[local]rock1200(config)#port ethernet 2/1
[local]rock1200(config-port)#link-group ETHER_EXAMPLE
[local]rock1200(config-port)#qos policy metering meter_policy
[local]rock1200(config-port)#qos policy policing policing_policy
[local]rock1200(config-port)#qos policy queuing queuing_policy
[local]rock1200(config-port)#end
```

### 3.5.2 Example: Provision QoS on a Dot1q LAG

The following example shows how to provision QoS policies on a dot1q LAG. For dot1q LAGs, QoS service is configured on each independent constituent port within the dot1q LAG. All dot1q PVCs within the LAG automatically inherit the queuing policy, while metering and policing policies are applied to the PVCs only if the `inherit` keyword is included in the `qos policy` command syntax.



```

!Create a LAG that has two PVCs, and bind each PVC to an interface:
[local]rock1200#configure
[local]rock1200(config)#link-group DOT1Q_EXAMPLE dot1q
[local]rock1200(config-link-group)#dot1q pvc 1
[local]rock1200(config-link-pvc)#bind interface PVC_1 local
[local]rock1200(config-link-pvc)#exit
[local]rock1200(config-link-group)#dot1q pvc 2
[local]rock1200(config-link-pvc)#bind interface PVC_2 local
[local]rock1200(config-link-pvc)#end

! Add the first port to the LAG and configure QoS policies on that port:
[local]rock1200#configure
[local]rock1200(config)#port ethernet 1/2
[local]rock1200(config-port)#link-group DOT1Q_EXAMPLE
[local]rock1200(config-port)#qos policy metering meter_policy inherit
[local]rock1200(config-port)#qos policy policing policing_policy inherit
[local]rock1200(config-port)#qos policy queuing queuing_policy
[local]rock1200(config-port)#end

! Add the second port to the LAG and configure QoS policies on that port:
[local]rock1200#configure
[local]rock1200(config)#port ethernet 2/2
[local]rock1200(config-port)#link-group DOT1Q_EXAMPLE
[local]rock1200(config-port)#qos policy metering meter_policy inherit
[local]rock1200(config-port)#qos policy policing policing_policy inherit
[local]rock1200(config-port)#qos policy queuing queuing_policy
[local]rock1200(config-port)#end

```

### 3.5.3 Example: Provision QoS on an Access LAG

The following example shows how to provision QoS policies on an access LAG. For access LAGs, QoS policies are applied to the LAG itself or to the individual PVCs or subscribers configured under the LAG.

```

! Create a LAG that has three PVCs. Configure QoS policies on each PVC, and bind each PVC to an interface.
[local]rock1200#configure
[local]rock1200(config)#link-group ACCESS_EXAMPLE access
[local]rock1200(config-link-group)#encapsulation dot1q
[local]rock1200(config-link-group)#qos policy queuing queuing_policy
[local]rock1200(config-link-group)#dot1q pvc 1
[local]rock1200(config-dot1q-pvc)#qos policy metering meter_policy
[local]rock1200(config-dot1q-pvc)#bind interface PVC_1 local
[local]rock1200(config-dot1q-pvc)#exit
[local]rock1200(config-link-group)#dot1q pvc 2 encapsulation lqtunnel
[local]rock1200(config-dot1q-pvc)#qos policy policing policing_policy inherit
[local]rock1200(config-dot1q-pvc)#exit
[local]rock1200(config-link-group)#dot1q pvc 2:1
[local]rock1200(config-dot1q-pvc)#qos policy metering meter_policy
[local]rock1200(config-link-pvc)#bind interface PVC_2 local
[local]rock1200(config-dot1q-pvc)#exit
[local]rock1200(config-link-group)#dot1q pvc 2:2
[local]rock1200(config-dot1q-pvc)#qos policy metering meter_policy inherit
[local]rock1200(config-link-pvc)#bind authentication pap chap local
[local]rock1200(config-link-pvc)#end

! Configure the first port with dot1q encapsulation and add it to the LAG:
[local]rock1200#configure
[local]rock1200(config)#port ethernet 1/3
[local]rock1200(config-port)#encapsulation dot1q
[local]rock1200(config-port)#link-group ACCESS_EXAMPLE
[local]rock1200(config-port)#end

! Configure the second port with dot1q encapsulation and add it to the LAG:
[local]rock1200#configure
[local]rock1200(config)#port ethernet 2/3
[local]rock1200(config-port)#encapsulation dot1q
[local]rock1200(config-port)#link-group ACCESS_EXAMPLE
[local]rock1200(config-port)#end

```